A CATEGORICAL INTRODUCTION TO SHEAVES

DAPING WENG

Abstract. Sheaf is a very useful notion when defining and computing many different cohomology theories over topological spaces. There are several ways to build up sheaf theory with different axioms; however, some of the axioms are a little bit hard to remember. In this paper, we are going to present a "natural" approach from a categorical viewpoint, with some remarks of applications of sheaf theory at the end. Some familiarity with basic category notions is assumed for the readers.

Contents

1. Motivation

In many occasions, we may be interested in algebraic structures defined over local neighborhoods. For example, a theory of cohomology of a topological space often concerns with sets of maps from a local neighborhood to some abelian groups, which possesses a natural $\mathbb{Z}$-module struture. Another example is line bundles (either real or complex): since $\mathbb{R}$ or $\mathbb{C}$ are themselves rings, the set of sections over a local neighborhood forms an $\mathbb{R}$ or $\mathbb{C}$-module.

To analyze this local algebraic information, mathematians came up with the notion of sheaves, which accommodate local and global data in a natural way. However, there are many fashion of introducing sheaves; Tennison [2] and Bredon [1] have done it in two very different styles in their seperate books, though both of which bear the name "Sheaf Theory". In this paper, we would use category theory as a tool (which is closer to Tennison yet some proofs of this paper may be more categorical) to give an introductory survey to this useful notion, sheaves.

2. Definitions and Constructions

2.1. **Presheaf.** Before we define sheaves, we first want to introduce the notion of presheaves, which is simpler and yet very helpful in understanding sheaf theory. The idea of a presheaf over a space is to associate each open set with an algebraic object, which often carries data about the open set itself, in such a way that we can establish a map from a bigger open set to a smaller open set inside it. We can think of it as there are layers of open sets where smaller ones are sitting above the bigger ones, and we want to assign each layer some algebraic object in a compactible way from bottom to top. For simplicity, throughout this paper, we are going to use R for a commutative ring and X for a topological space unless otherwise specified.

For any space X, we want to define a category called the category of open sets $\mathfrak{Opn}_X$. The objects in $\mathfrak{Opn}_X$ are open sets of X and morphisms are inclusions. Then we define presheaf as following

Definition 2.1. A *presheaf* of R-modules on a space X is a contravariant functor

$$A : \mathfrak{Opn}_X^{op} \to \mathfrak{Mod}_R.$$

Elements in each such R-module are called *sections* of the presheaf over a particular open set.

By this definition, for any inclusion map $V \subset U$, we get an R-module homomorphism $j : A(U) \to A(V)$. To distinguish this R-module homomorphism j from others, we name this particular functorial one the *restriction* from $A(U)$ to $A(V)$, and by convention write

$$j(s) = s|_V$$

for any $s \in A(U)$.

Remark 2.2. We can define presheaf of many other categories using this definition as well. In particular, since abelian groups are $\mathbb{Z}$-modules, presheaf of abelian groups fits into this definition.

The name "restriction" maps may not be very meaningful in this context, since in presheaves restriction maps may not even be surjective. However, as we will see in the case of sheaves, the name "restriction" is in fact our familiar restriction in the common sense.

Notice that there can be more than one possible presheaf on a space X, and between two presheaves A and B there can be natural transformation η such that the following diagram commutes

$$\begin{array}{ccc} A(U) & \xrightarrow[\eta_U]{} & B(U) \\ \downarrow{\scriptstyle j_A} & & \downarrow{\scriptstyle j_B} \\ A(V) & \xrightarrow[\eta_V]{} & B(V) \end{array} \tag{2.3}$$

for any open sets $V \subset U$.

If we regard these presheaf functors as objects and natural transformations between them as morphisms, we get a new category called the category of presheaves over X, denoted by $\mathfrak{Mod}_R^{\mathfrak{Opn}_X^{op}}$ or simply by $\mathfrak{Prshf}_X$. By convention, the morphisms (natural transformation between presheaves) are called presheaf homomorphisms.

Examples 2.4.

(1) If we associate every open subset of space X to the trivial R-module and let the restrictions be the trivial homomorphism, we get the *zero presheaf* (the trivial presheaf), denoted by 0. Notice that this is both the initial and terminal object in $\mathfrak{Prshf}_X$, and hence the zero object in $\mathfrak{Prshf}_X$.

(2) Let X be a singleton set $\{*\}$. Then $\mathfrak{Prshf}_{\{*\}}$ is the same as the category $\mathfrak{Mod}_R$.

(3) Consider a smooth manifold M. For each open subset U, consider the set $\mathbb{C}_M(U)$ of all complex-valued smooth maps over U. It is easy to see that $\mathbb{C}_M(U)$ is a complex vector space (a $\mathbb{C}$-module). For two open subsets $V \subset U$, we can define the restriction map to be the restriction of maps in $\mathbb{C}_M(U)$ to the smaller open subset. This gives us a perfectly good presheaf, which will appear many times in this paper.

One fact about category of functors is that most constructions can be carried out componentwise in the codomain category. In the case of $\mathfrak{Prshf}_X$, since the codomain category is the algebraic category $\mathfrak{Mod}_R$, we will expect some basic algebraic notions to be well-defined for presheaves.

(1) A *subpresheaf* of a presheaf A is a presheaf B together with a presheaf monomorphism $\eta : B \to A$. In particular, the zero presheaf is a subpresheaf of any presheaf.

(2) The *kernel* of a presheaf homomorphism $\eta : A \to B$ is categorically defined to be the unique object $\ker(\eta)$ with a presheaf homomorphism into A that makes the following diagram on the left commute:

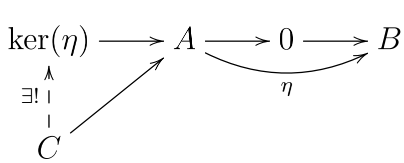

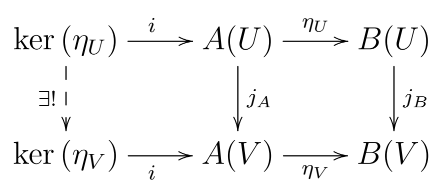

To show existence, we just need to take kernels of R-module homomorphism at every open set level, and by the categorical property of kernels we get all the restriction maps naturally, which fit into the definition of a presheaf (diagram on the right).

(3) The *product presheaf* $A \times B$ of two presheaves A and B is categorically defined as usual with the diagram below. One can verify that open set level-wise product in the category of $\mathfrak{Mod}_R$ gives the right construction for product presheaf, i.e., setting $A \times B(U) = A(U) \times B(U)$ for each open set U.

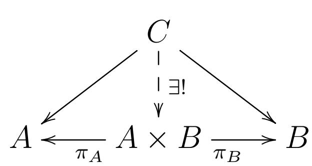

In fact, this definition of sheaf products can be generalized to products among a collection of presheaves of any cardinality.

(4) The *cokernel* of a presheaf homomorphism is the dual notion of kernel, as in the diagram on the left. Construction is also carried out open set level-wise, and restrictions

are natural result of cokernels in the category $\mathfrak{Mod}_R$.

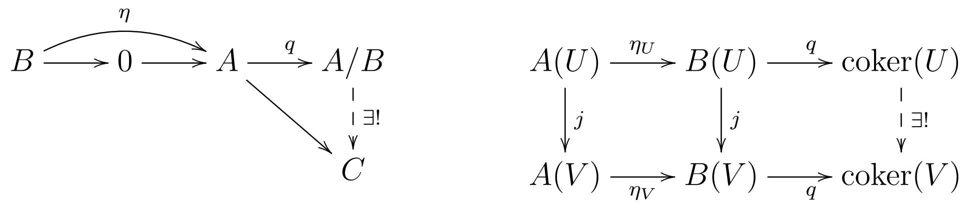

(5) In particular, a quotient presheaf B/A is a cokernel of a monomorphism $A \to B$.

(6) By a categorical convention, we define the image of a homomorphism $\eta : A \to B$ to be the kernel of the cokernel of η, denoted by $\text{img}(\eta)$ [2]. (For the readers who know more category theory, this definition can be used in any abelian category, to which $\mathfrak{Prshf}_X$ belongs.) Notice that a unique homomorphism can be obtained from A to $\text{img}(\eta)$ via the universal property of kernels:

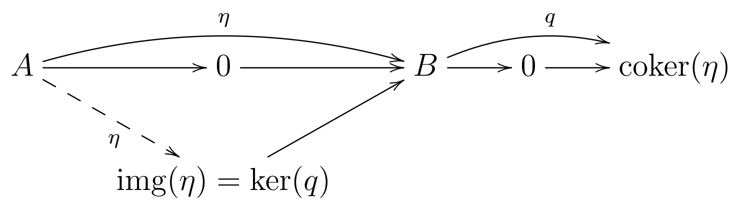

Readers can easily verify that $\text{img}(\eta)$ is infact a presheaf consisting of $(\text{img}(\eta))_U = \eta_U(A(U))$. By convention, this resulting homomorphism from A to $\text{img}(\eta)$ is also named η.

2.2. **Sheaf.** From the discussion above, we can see that a presheaf may carry a lot of data and we can always find out data about smaller open sets via restriction maps. However, ideally we would also want to process information in the opposite direction, namely obtaining data about a bigger open set by just looking at its open cover. Thus we would like to impose one condition to specify this particular kind of presheaf, which are the ones that allow us to "glue" the pieces over every open set in an open cover, and this is the notion of a sheaf.

Definition 2.5. A presheaf A is a *sheaf* if it satisfies the following equalizer diagram

$$A(U) \xrightarrow{f} \prod_\alpha A(U_\alpha) \underset{h}{\overset{g}{\rightrightarrows}} \prod_{\beta,\gamma} A(U_\beta \cap U_\gamma)$$

whenever $U = \bigcup_\alpha U_\alpha$. The map f is the product of restrictions, whereas g and h are defined by

$$g\left(\prod_\alpha s_\alpha\right) = \prod_{\alpha,\beta} s_\alpha|_{U_\alpha \cap U_\beta}, \qquad h\left(\prod_\alpha s_\alpha\right) = \prod_{\beta,\alpha} s_\alpha|_{U_\beta \cap U_\alpha}$$

One may wonder how this definition helps us "glue" pieces together. For demonstration, let A be a sheaf over X. Suppose $\{U_\alpha\}$ is an open cover of an open set U. Upon each of the U_α, suppose we pick a section s_α in a compatible way, namely whenever U_α intersect U_β non-trivially, we require that

$$s_\alpha\big|_{U_\alpha \cap U_\beta} = s_\beta\big|_{U_\alpha \cap U_\beta}.$$

Then this definition basically says that there must be a unique section $s \in A(U)$ such that

$$s|_{U_\alpha} = s_\alpha$$

for each α.

An important property of sheaves arises from this definition. If we consider the empty set as an open set in X and cover it with the empty covering, then the equalizer diagram above becomes

$$A(\emptyset) \xrightarrow{f} \{0\} \rightrightarrows \{0\}$$

where $\{0\}$ is the terminal object of the category $\mathfrak{Mod}_R$. (Recall that the 0-fold categorical product gives the terminal object.) Thus for any sheaf A, $A(\emptyset)$ is always the trivial R-module.

There are alternative definitions for sheaves, some involves a topology defined on the sheaf a priori; right now we just want to stick with our algebraic definition, and later we will show how this gives a natural topology on a sheaf. Another common algebraic definition uses two sheaf axioms called "monopresheaf" and "conjunctive" [1], which is essentially equivalent to the equalizer diagram above.

Since any sheaf is a presheaf, we have morphisms between sheaves in the same sense as presheaves, i.e., natural transformations between functors. Notice that any presheaf homomorphism between sheaves automatically commutes with the equalizer diagrams as illustrated below

$$\begin{array}{ccccc} A(U) & \xrightarrow{f} & \prod_\alpha A(U_\alpha) & \overset{g}{\underset{h}{\rightrightarrows}} & \prod_{\beta,\gamma} A(U_\beta \cap U_\gamma) \\ \downarrow{\scriptstyle \eta_U} & & \downarrow{\scriptstyle \prod_\alpha \eta_{U_\alpha}} & & \downarrow{\scriptstyle \prod_{\beta,\gamma} \eta_{U_\beta \cap U_\gamma}} \\ B(U) & \xrightarrow{f} & \prod_\alpha B(U_\alpha) & \overset{g}{\underset{h}{\rightrightarrows}} & \prod_{\beta,\gamma} B(U_\beta \cap U_\gamma) \end{array}$$

Hence we can define a new category called the category of sheaves over X, denoted by $\mathfrak{Shf}_X$, to be the full subcategory of $\mathfrak{Prshf}_X$ whose objects satisfies Definition 2.5. Also by convention, we would call the morphisms in this category *sheaf homomorphism.*

Examples 2.6. (1) The zero presheaf obviously satisfies the equalizer diagram and hence is a sheaf as well.

(2) An example of a presheaf that is not a sheaf. Take a two-point space $\{x, y\}$ with the discrete topology. Let A be a presheaf defined as follows:

$$A(\{x, y\}) = \mathbb{Z}, A(\{x\}) = \mathbb{Z}_2, A(\{y\}) = \mathbb{Z}_2, A(\emptyset) = 0$$

with the obvious R-module homomorphisms. However, this is not a sheaf since the map

$$A(\{x, y\}) \to A(\{x\}) \times A(\{y\})$$

is not injective.

3. Sheafification

3.1. **Direct Limit and Stalks.** We have not yet explained why this mathematical notion defined above is called "sheaf", or in other words, what these sheaves actually consist of. Thus in this section we want to tell the readers what the "stalks" inside a sheaf are so that we have some intuitive picture to make sense of this term. We would like to start with the notion of direct limit.

Definition 3.1. In a category $\mathfrak{C}$, a *directed system* is a set of objects

$$\{C_i \mid i \in I, I \text{ has a preorder} \leq \}$$

together with morphisms $f_{ij} : C_i \to C_j$ for $i \leq j$ such that

(1) $f_{ii} = \mathrm{Id}_{A_i}$;
(2) $f_{ik} = f_{jk} \circ f_{ij}$ for $i \leq j \leq k$.

Example 3.2. An example of a directed system can be found in presheaves. Fix a point $x \in X$ and consider all open sets U containing x. Notice that inclusion is a preorder on the collection of open sets (say $U \geq V$ if $U \subset V$). Thus for a presheaf A over X, we form a directed system $\{A(U) \mid x \in U\}$ in $\mathfrak{Mod}_R$, with the morphisms being the presheaf restriction maps.

Since a directed system is a small category, sometimes we want to know whether a colimit exists for its diagram. Thus we have the following definition:

Definition 3.3. A direct limit of a directed system in the category $\mathfrak{C}$ is the colimit of the directed system, i.e., an object C together with morphisms $\phi_i : C_i \to C$ such that the following diagram commutes for every $i \leq j$.

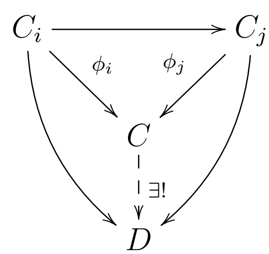

For a category in which every directed system has a colimit, we say that this category possesses direct limits. (It may sound confusing that the direct limit is actually a colimit. The dual notion of the direct limit is called "inverse limit", which is defined to be the limit of a directed system.)

The following is an important proposition for the category of R-modules.

Proposition 3.4. *The category $\mathfrak{Mod}_R$ possesses direct limits.*

Proof. Let $\{M_i \mid i \in (I, \leq)\}$ be a directed system in $\mathfrak{Mod}_R$. Take the direct sum $\bigoplus_{i \in I} M_i$ in $\mathfrak{Mod}_R$ and consider the submodule N generated by the elements in the form $m_i - f_{ij}(m_i)$ where $i \leq j$. Now we can take the quotient R-module $M = \left(\bigoplus_{i \in I} M_i\right) / N$. For each $i \in I$, define a morphism $\phi_i : M_i \to M$ to be the composition of the inclusion map from M_i into the direct sum followed by the quotient map by N. It is easy to see that M satisfies the commutative diagram above and hence is the direct limit of the directed system $\{M_i\}$. $\square$

As we described in Example 3.2, $\{A(U) \mid x \in U\}$ is a directed system in $\mathfrak{Mod}_R$ for any presheaf A, and hence we are welcome to take its direct limit:

Definition 3.5. The *stalk* A_x of presheaf A at x is the direct limit of directed system $\{A(U_i) \mid i \in I\}$, where $\{U_i \mid i \in I\}$ is a direct set of open neighborhoods of x.

Elements in a stalk are called *germs*. In particular, if a germ comes out as the direct limit of a section $s \in A(U)$, then we denote it by s_x.

We denote natural homomorphism that sends $s \in A(U)$ to s_x by θ_x.

The following diagram comes from the definition of direct limit ($V \subset U$):

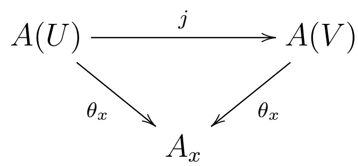

Remark 3.6. Notice that stalks are actually defined for all presheaves; nevertheless, which will be clear by the end of the next section, sheaves are exactly the presheaves that successfully tie its stalks into a bundle and that is how the name comes out.

The following is a useful fact that we should keep in mind:

Proposition 3.7. *If A is a sheaf and s, t are two sections in $A(U)$, then $s = t$ if and only if $s_x = t_x$ for all $x \in U$.*

Proof. The forward direction is trivial. In the backward direction, we know that $s_x = t_x$ if and only if there exists a neighborhood U_x of x, $U_x \subset U$, such that

$$s|_{U_x} = t|_{U_x}.$$

However, if we take all such neighborhoods U_x for all x, we get an open cover for U, and by the definition of sheaves,

$$A(U) \to \prod_{x \in U} A(U_x)$$

should be injective, and hence we have $s = t$. □

Another natural question to ask is how stalks from two presheaves relate. Fix a point $x \in X$ and suppose η is a homomorphism from presheaf A to presheaf B. Immediately from the definition of presheaf homomorphism we know that η gives R-module homomorphisms $\eta_U : A(U) \to B(U)$ for each open set U, and Diagram 2.3 gives the compatibility of directed systems $\{A(U) \mid x \in U\}$ and $\{B(U) \mid x \in U\}$. Now follow the universal property of direct limit, we obtain a unique R-module homomorphism η_x from A_x to B_x such that the following diagram commutes for any $x \in V \subset U$:

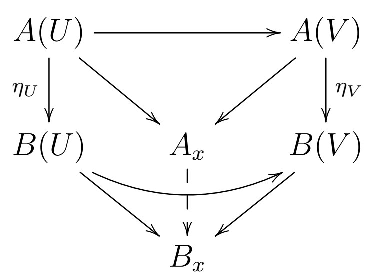

This seems to have some categorical implication: for each presheaf A over X, we can take stalks at each point in X, and there are R-module homomorphisms on corresponding stalks whenever there is a homomorphism between two presheaves - all of these look like that we have produced a new category! To make it precise, for any presheaf A over X, we define an object LA by taking

$$LA = \bigcup_{x \in X} A_x.$$

Between two objects LA and LB, if there is a morphism $\eta : A \to B$ in $\mathfrak{Prshf}_X$, we can define a map

$$L\eta : LA \to LB, \qquad L\eta(s_x) = \eta_x(s_x).$$

It is easy to see that

$$\begin{aligned} L\,\mathrm{Id}(s_x) =& s_x \\ \text{and} \qquad L(\eta \circ \xi)(s_x) =& (\eta \circ \xi)_x(s_x) = \eta_x \circ \xi_x(s_x) = L\eta \circ L\xi(s_x) \end{aligned}$$

Thus L can be view as a functor from $\mathfrak{Prshf}_X$ to this new category.

However, these objects and maps are not exciting enough. At first glance, LA looks like a principal bundle, except for the lack of a topology. Thus, one way to make LA interesting is to put a topology onto it, in such a way that the maps $L\eta$ becomes continuous. As a result, we have the following proposal: consider an open set U in X and define

$$s_U = \{s_x \mid x \in U, s \in A(U)\}.$$

It is not hard to verify that $\{s_U \mid U \text{ open}, s \in A(U)\}$ gives a topological basis for LA. Now we just need to show the continuity of $L\eta : LA \to LB$:

Proposition 3.8. *$L\eta$ is continuous.*

Proof. It is enough to show that the preimage of t_V is open for any $t \in B(V)$. Suppose $L\eta\,(s_x) = \eta_x\,(s_x) = t_x$. By the definition of direct limit, this means there exists some neighborhood $W \subset V$ such that $\eta_W : s|_W \mapsto t|_W$. But then this implies that $\eta_y\,(s_y) = t_y$ for all $y \in W$, i.e., $L\eta\,(s_W) \subset t_V$. □

Now since in this new category, objects are topological spaces and morphisms are continuous, it is a subcategory of $\mathfrak{Top}$, the category of topological spaces. We shall name this new category *the sheaf spaces over X*, denoted by $\mathfrak{ShfSpc}_X$. We need to point out that when L is regarded as a functor from $\mathfrak{Prshf}_X$ to $\mathfrak{ShfSpc}_X$.

We can see that by applying the functor L to any presheaf A, what it does is sorting out the stalks over X and putting a topology such that each cross-section over an open set is open and homeomorphic to the underlying open set. A sheaf space almost looks like a sheaf; in fact, as we will see in the next section, all we need to do is just "tie up" the stalks naturally and the sheaf space will automatically turns into a sheaf. This is precisely the idea of "sheafifying" a presheaf into a sheaf.

3.2. **Sheafification in Action.** As indicated in the previous section, there seems to be a way to make every presheaf into a sheaf: we have already defined a functor L that can turn a presheaf into something close to be a sheaf; we just need to finish the "tie up" process, which will be realized by a functor $\Gamma : \mathfrak{ShfSpc}_X \to \mathfrak{Shf}_X$.

For each object LA in $\mathfrak{ShfSpc}_X$, we define ΓLA to be the presheaf that associates an open set U with

$$\Gamma LA(U) = \{\mu_U \subset LA \mid \mu_U \text{ is homeomorphic to } U \text{ via the map } p : s_x \mapsto x\}$$

.

This construction may seem to be too weak and rather arbitrary. Elements in $\Gamma LA(U)$ are homeomorphic sets in the sheaf space LA; but what we need is an R-module structure on $\Gamma LA(U)$, so how are addition and scalor multiplication defined over homeomorphic open sets?

Recall that the topology on LA is defined by the basis $\{s_U \mid s \in A(U)\}$. Therefore a typical element $\mu_U \in \Gamma LA(U)$ would be a union of some s_{U_α} in which $\{U_\alpha\}$ is a open cover of U. Moreover, since we demand the map $p : s_x \mapsto x$ to be a homeomorphism, there must be one and only one germ over each point from the corresponding stalk. Therefore wherever U_β intersects U_γ non-trivially, s_{U_β} must agree with s_{U_γ} on the germs over the intersection, i.e., for $x \in U_\beta \cap U_\gamma$,

$$s_x \in s_{U_\beta} \iff s_x \in s_{U_\gamma}.$$

Thus, μ_U can be understood as many pieces of s_{U_α} "gluing" together at the intersections.

Since in μ_U, there is one and only one germ for each point, we can define addition and scalor multiplication stalk-wise. The only thing we need to verify is that we still get a homeomorphic image of U after addition and scalor multiplication. Notice that locally at the point x, μ_U is just s_V for some neighborhood V of x. Observe that for any $r \in R$,

$$rs_V + t_V = \{rs_x + t_x \mid x \in V\} = (rs + t)_V.$$

Therefore both addition and scalor multiplication are well-defined.

Between an open set and one of its open subsets, the restriction map is just truncating each section to the homeomoprhic part corresponding to the open subset, i.e., for open sets $V \subset U$,

$$(\mu_U)\big|_V = \{s_x \in \mu_U \mid x \in V\}.$$

It is not hard to see that this restriction map is an R-module homomorphism.

Up to this point, we have shown that ΓLA is a presheaf; now the task is to verify the sheaf condition (Definition 2.5):

$$\Gamma LA(U) \longrightarrow \prod_\alpha \Gamma LA(U_\alpha) \rightrightarrows \prod_{\beta,\gamma} \Gamma LA(U_\beta \cap U_\gamma)$$

Suppose we have another R-module M such that the following two composite maps are the same

$$M \xrightarrow{f} \prod_\alpha \Gamma LA(U_\alpha) \rightrightarrows \prod_{\beta,\gamma} \Gamma LA(U_\beta \cap U_\gamma)$$

Then we obtain a well defined map $\tilde{f} : M \to \Gamma LA(U)$ by sending

$$\tilde{f} : m \mapsto \bigcup_\alpha \pi_\alpha \circ f(m)$$

(where π_α is the projection map of the product R-module). Notice that the commutativity of the above diagram guarantees that this map is well-defined. Furthermore, $\tilde{f}$ is the unique map that makes the following diagram commute:

$$\begin{array}{ccccc} \Gamma LA(U) & \longrightarrow & \prod_\alpha \Gamma LA(U_\alpha) & \rightrightarrows & \prod_{\beta,\gamma} \Gamma LA(U_\beta \cap U_\gamma) \\ \tilde{f} \uparrow & \nearrow_{f} & & & \\ M & & & & \end{array}$$

(Uniqueness follows from the fact that the map f already determines the germ s_x over every $x \in U$ in $f(m)$.) Thus, we can conclude that ΓLA does in fact form a sheaf for each presheaf A.

To finish the definition of the functor Γ, we still need to define $\Gamma L\eta : \Gamma LA \to \Gamma LB$ for a presheaf homomorphism $L\eta : LA \to LB$. The most obvious thing to do is

$$\Gamma L\eta_U\,(\mu_U) = L\eta\,(\mu_U) = \{\eta_x(s_x) \mid s_x \in \mu_U\}\,.$$

Lastly, it is easy to see the following two functorial properties of Γ:

$$\Gamma L\,\mathrm{Id}_U(\mu_U) = L\,\mathrm{Id}\,(\mu_U) = \mu_U$$

and $$\Gamma L(\eta \circ \xi)_U(\mu_U) = L(\eta \circ \xi)\,(\mu_U) = L\eta \circ L\xi\,(\mu_U) = \Gamma L\eta \circ \Gamma L\xi\,(\mu_U)$$

In conclusion, we have the following diagram showing the functors that we have so far (I is the inclusion functor):

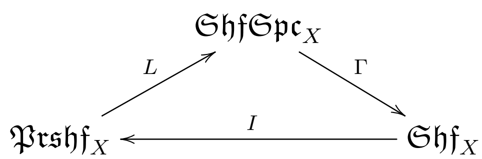

One interesting question to ask is what if we apply this functor to a sheaf? Is the new sheaf any different from the original one?

The answer is no. If we recall that for each section $s \in A(U)$, there is a corresponding basis element $s_U \in \Gamma LA(U)$, and the map $\sigma_U : s \mapsto s_U$ is an R-module homomorphism. We claim that this homomorphism is bijective and hence we have the following:

Proposition 3.9. *If A is a sheaf over X, then $A(U)$ and $\Gamma LA(U)$ are isomorphic (via θ_U).*

Proof. For surjectivity, let μ_U be an element in $\Gamma LA(U)$, which can be thought of as a union of s_{U_α}, in which $\{U_\alpha\}$ is an open cover of U. But then wherever U_β and U_γ intersect non-trivially, the germs in s'_{U_β} and s''_{U_γ} over these points must be identical. By Proposition 3.7, we can deduce that

$$s'|_{(U_\beta \cap U_\gamma)} = s''|_{(U_\beta \cap U_\gamma)}.$$

Now by Definition 2.5, we know that there must be a section $s \in A(U)$ such that

$$s|_{U_\beta} = s_{U_\beta}$$

and therefore $\theta_U(s) = \mu_U$.

For injectivity, suppose $s_U = 0_U$. Then at each stalk, we know that $s_x = 0$ for all $x \in U$. Proposition 3.7 yields immediately $s = 0$, and we are done. ☐

This shows that ΓLA is essentially the same sheaf as A; for simplicity, we are not going to distinguish the sheafification of a sheaf from the sheaf itself. Moreover, on the morphism level, we also have the following:

Proposition 3.10. *If $\eta : A \to B$ is a sheaf homomorphism, then $\Gamma L\eta = \eta$.*

Proof. If η_U sends $s \in A(U)$ to $t \in B(U)$, then by the definition of $\Gamma L\eta$,

$$\Gamma L\eta_U\,(s_U) = \{\eta_x(s_x) \mid x \in U\} = \{t_x \mid x \in U\} = t_U.$$ ☐

We now know that the sheafification functor is the same as the identity functor when restricted to the subcategory $\mathfrak{Shf}_X$. Moreover, since when A is a sheaf, the sections in $\Gamma LA(U)$ are just $s_U = \{s_x \mid x \in U\}$, with restriction map between $V \subset U$ being $s_U|_V = s_V$, it is easy to see that $(s_U)_x = s_x$. Thus we also have the following proposition

Proposition 3.11. *The stalk A_x of a presheaf A is naturally isomorphic to the stalk $(\Gamma LA)_x$ for each point x.*

Remark 3.12. Notice that midway through sheafification, we get a sheaf space as a union of stalks. Thus we can identify the sheaf space with the original sheaf and thereby get a nice topology on a sheaf, in which each section is a basis open set that is homeomorphic to the associated open set in the underlying space. Therefore in many texts, a sheaf is identified with the corresponding sheaf space and is defined with topological properties a priori [1]. Next we are going to rediscover the necessary and sufficient condition for a topological space to be a sheaf over some subspace.

Theorem 3.13. *For a topological space A and a subspace X with a continuous map $p : A \to X$, we say that (A, p) is a sheaf over X if it satisfies the following:*

(1) p is a local homeomorphism onto X;
(2) each "stalk" $A_x = p^{-1}(x)$ is an R-module;
(3) the R-module operations are continuous in the following manner: define

$$A\Delta A = \{(\alpha, \beta) \in A \times A \mid p(\alpha) = p(\beta),$$

and we demand that both addition and scalor multiplication by r

$$\begin{aligned} + : A\Delta A &\to A \\ (\alpha,\beta) &\mapsto \alpha+\beta \end{aligned} \qquad \begin{aligned} r : A &\to A \\ \alpha &\mapsto r\alpha \end{aligned}$$

to be continuous.

Proof. To show that they are sufficient, we just need to set $A(U) = p^{-1}(U)$ for each open set $U \subset X$ and define the R-module structure by addition and scalor multiplication point-wise. Check with Definition 2.5, this does give rise to a sheaf.

To show that they are necessary, let A be a sheaf over X as in Definition 2.5. By the remark above, we know that A is equiped naturally with a topology. Now let the map p be defined as

$$\begin{aligned} p : A = LA &\to X \\ \alpha \in A_x &\mapsto x \end{aligned}$$

Since the sections of A are basis open sets and homeomorphic to the open sets that they are over, the map p is obviously local homeomorphic.

Stalks A_x are R-modules by construction.

Lastly, for $r \in R$ and $s, t \in A(U)$, we have

$$rs + t = rs_U + t_U = (rs + t)_U.$$

Since s_U, t_U and $(rs + t)_U$ are all basis open sets, addition and scalor multiplication defined in (3) are continuous. □

We would also want some equivalent topological definition of a sheaf homomorphism.

Corollary 3.14. *A map η from a sheaf A to a sheaf B over X is a sheaf homomorphism if and only if η is open and the restriction of η to A_x (denoted by η_x) is an R-module homomorphism into B_x.*

Proof. The forward direction is clear. In the backward direction, fix any open set U in X. Since both A and B are sheaves, we know that

$$A(U) = \bigcup_{x \in U} A_x \qquad \text{and} \qquad B(U) = \bigcup_{x \in U} B_x$$

and hence the restriction of η to $A(U)$ (denoted by η_U) is a map into $B(U)$. Now we need to show that η_U is an R-module homomorphism with respect to the R-module structures on $A(U)$ and $B(U)$.

Pick a section $s \in A(U)$. We know that

$$\eta_U(s) = \eta_U(s_U) = \{\eta_x(s_x) \mid x \in U\}.$$

Since η is an open map, $\eta_U(s_U) = \eta(s_U)$ must be open. Thus there exists some basis open set t_{V_x} in B such that

$$\eta_x(s_x) \in t_{V_x} \subset \eta_U(s_U).$$

Notice that $\{V_x \mid x \in U\}$ form an open cover for U, and there is one and only one germ over every point $x \in U$. This forces the union of all t_{V_x} to be a section in $B(U)$, i.e.,

$$\bigcup_{x \in U} t_{V_x} = t_U = t \in B(U).$$

Thus, we can conclude that at least η_U maps sections to sections. But then it is not hard to see that η_U preseves the algebraic operations since η_x preserves them stalk-wise. Therefore the conditions stated are sufficient for η to be a sheaf homomorphism. $\square$

The above theorem and corollary provide an alternative way of defining a sheaf and sheaf homomorphism. Though maybe not as natural as Definition 2.5, sometimes it is easier to just give out the topology and stalk-wise R-module struture. From now on, we are going to use both definitions of sheaves in this paper interchangeably, without further explanation.

3.3. **Sheafification as an Adjoint Functor.** Now between the two categories $\mathfrak{Prshf}_X$ and $\mathfrak{Shf}_X$, we have the sheafification functor ΓL going in one direction and the inclusion functor I going in the other. If we consider a presheaf A and a sheaf B on some space X, by convention we can denote the set of presheaf homomorphisms between A and B (regarded as a presheaf) by $\mathfrak{Prshf}_X(A, IB)$, and similarly denote the set of sheaf homomorphisms between ΓLA and B by $\mathfrak{Shf}_X(\Gamma LA, B)$.

First notice that any section $s \in A(U)$ in the presheaf A is turned into a basis open set s_U, which is an element in $\Gamma LA(U)$ of the generated sheaf. It is easy to see that this association $\sigma_U : s \mapsto s_U$ is indeed an R-module homomorphism from $A(U)$ to $\Gamma LA(U)$, which is compatible with restriction maps. Thereby we obtain a presheaf homomorphism $\sigma : A \to \Gamma LA$. By easy computation, we can also see that $\Gamma L\sigma = \mathrm{Id}$.

$$\begin{array}{ccc} A(U) & \xrightarrow{\sigma_U} & \Gamma LA(U) \\ \downarrow & & \downarrow \\ A(V) & \xrightarrow[\sigma_V]{} & \Gamma LA(V) \end{array}$$

Now for any sheaf homomorphism $\eta : \Gamma LA \to B$, we can pre-compose σ and get a presheaf homomorphism $\sigma \circ \eta : A \to B$. Sheafify the composition map, we get

$$\Gamma L(\sigma \circ \eta) = \Gamma L\sigma \circ \Gamma L\eta = \Gamma L\eta = \eta.$$

From this, we can see that any sheaf homomorphism in $\mathfrak{Shf}_X(\Gamma LA, B)$ comes from some presheaf homomorphism in $\mathfrak{Prshf}_X(A, IB)$. Therefore the map

$$\begin{aligned}\Gamma L : \mathfrak{Prshf}_X(A, IB) &\to \mathfrak{Shf}_X(\Gamma LA, B)\\ \eta &\mapsto \Gamma L\eta\end{aligned}$$

is at least surjective.

We further claim that this map is also injective. Suppose η and ξ are two presheaf morphisms between A and B such that $\Gamma L\eta = \Gamma L\xi$. This implies that for any open set U and $s \in A(U)$, we have

$$\Gamma L\eta_U(s_U) = \Gamma L\xi_U(s_U).$$

By the definition of $\Gamma L\eta$ and properties of direct limits, we have

$$\Gamma L\eta_U(s_U) = \{\eta_x(s_x) \mid x \in U\} = (\eta_U(s))_U.$$

Thus, there must be

$$(\eta_U(s))_U = (\xi_U(s))_U.$$

But we just showed in Proposition 3.9 that $B(U) \cong \Gamma LB(U)$ for any sheaf B. Thus we can conclude that $\eta_U(s) = \xi_U(s)$, and hence $\eta = \xi$. This assertion gives the following conclusion:

Theorem 3.15. *The sheafification functor is left adjoint to the inclusion functor and the inclusion functor is right adjoint to the sheafification functor, i.e.,* $\Gamma L \dashv I$.

One important categorical fact about adjunction of functors is the following theorem, which we will not prove in this paper but readers can find easily accessible proof in texts such as Awodey [3]; we will apply it to notions that we have introduced before in the category $\mathfrak{Prshf}_X$ and once for all define and prove the existence of equivalent notions in the category $\mathfrak{Shf}_X$.

Theorem 3.16 (RAPL & LAPC)**.** *Right adjoints preserve limits and left adjoints preserve colimits.*

Corollary 3.17. *There exist products and kernels in the category* $\mathfrak{Shf}_X$*, which are isomorphic to the presheaf products and presheaf kernels. A sheaf monomorphism is also a presheaf monomorphism, and hence a subsheaf is automatically a subpresheaf.*

Proof. Notice that products and kernels are limits; therefore they commute with the right adjoint functor I. Thus, if sheaf products and sheaf kernels exist, they must be isomorphic to their presheaf counterparts. To show existence, we just need to check whether the presheaf products and kernels actually satisfy Definition 2.5; and fortunately, they do.

For monomorphisms, notice that a monomorphism $\eta : A \to B$ in any category can be realized as a map such that the following pull-back diagram yields A:

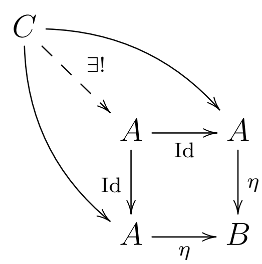

Since pull-backs are limits, they are preserved under the right adjoint functors, such as I. Thus, sheaf monomorphisms are also presheaf monomorphisms when regarded as presheaf homomorphisms. $\square$

Now we should look at the notion of cokernels. Suppose $\eta : A \to B$ is a presheaf homomorphism. Since the sheafification functor is a left adjoint functor, colimts such as cokernels are preserved under it. Thus we can say that the *sheaf cokernel* of the map $\Gamma L\eta$ is nothing but $\Gamma L(\mathrm{coker}(\eta))$. Notice that sheafification is exactly the identity functor when restricted to the subcategory of sheaves. Therefore we have the following:

Corollary 3.18. *If $\eta : A \to B$ is a sheaf homomorphism, then the sheaf cokernel is isomorphic to the sheafification of the presheaf cokernel obtained when regarding η as a presheaf homomorphism. Recall that a quotient is the cokernel of a monomorphism, and a sheaf monomorphism is also a presheaf monomorphism. Thus, for two sheaves $B \subset A$, the quotient sheaf is given by $\Gamma L(A/B)$.*

Without all the categorical machinaries, the results above would seem rather arbitrary; however, if we look at them from a categorical viewpoint, they are in fact natural constructions! Just to convince the readers that sheafification is necessary in producing categorical cokernels ans quotients in $\mathfrak{Shf}_X$, the following is an example in which the presheaf quotient of two sheaves is not a sheaf.

Example 3.19. Let X be the interval $[0,1]$ with the usual topology. Let A be the sheaf defined by

$$A(U) = \{f : U \to \mathbb{R} \mid f \text{ is continuous}\}$$

(Notice that $A(U)$ possesses an $\mathbb{R}$-module structure) with the restriction maps being truncation. On the other hand, let B be the subsheaf of A defined by

$$B(U) = \left\{ f \in A(U) \;\middle|\; \lim_{x=1/2} f(x) = 0 \text{ if } 1/2 \text{ is a limit point of } U \right\}.$$

If we just carry out the quotient presheaf construction, we have

$$A/B(U) = \begin{cases} \mathbb{R} & \text{if } 1/2 \text{ is a limit point of } U \\ 0 & \text{otherwise} \end{cases}$$

Now consider the open set $[0,1/2)$ and its open cover

$$\{[0, 1/2 - 1/n) \mid n > 2, n \in \mathbb{Z}\}.$$

By construction, $A/B[0,1/2) \cong \mathbb{R}$, but $A/B[0,1/2-1/n) = 0$ for all $n > 2, n \in \mathbb{Z}$, which cannot satisfy the sheaf condition in Definition 2.5.

Remarks 3.20. As in the case of monomorphism, epimorphism can be defined to be a map $\eta : A \to B$ such that the following push-out diagram yields B. Since push-outs are colimits, the sheafification of η remains epimorphic from ΓLA to ΓLB.

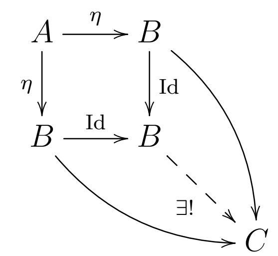

In particular, an immediate result is that a sheaf homomorphism that is epimorphic as a presheaf homomorphism is automatically epimorphic as a sheaf homomorphism.

Similar to presheaves, the image of a sheaf homomorphism η is defined to be the kernel of the cokernel of η. Similar to presheaf images, for every sheaf homomorphism $\eta : A \to B$, we get a unique sheaf homomorphism from A to $\operatorname{img}(\eta)$ via the universal property of kernels, which we shall name η as well by convention. Since sheaf cokernel may not agree with presheaf cokernel, the sheaf image may also not coincide with the presheaf image.

From now on, we are going to use $(A/B)_s$, $\operatorname{coker}_s(\eta)$ and $\operatorname{img}_s(\eta)$ for sheaf quotients, cokernels and images, just to distinguish from their presheaf counterparts $(A/B)_p$, $\operatorname{coker}_p(\eta)$ and $\operatorname{img}_p(\eta)$. As for the case of sheaf kernels and other notions that are isomorphic to their presheaf counterparts, we are going to just use $\ker(\eta)$ and so on without further distinguishment.

4. Exact Sequence

As in other categories with kernels and images, we can define the notion of an exact sequence of sheaves, and there are interesting properties that we need to pay attention to before we go into further investigation of sheaf theory.

Definition 4.1. A sequence of presheaves (sheaves) with presheaf (sheaf) homomorphisms

$$\longrightarrow A \xrightarrow{\eta} B \xrightarrow{\xi} C \longrightarrow$$

is said to be *exact* if $\ker(\xi) \cong \operatorname{img}_p(\eta)$ ($\ker(\xi) \cong \operatorname{img}_s(\eta)$) for every place B. In particular, the following exact sequence is call a *short exact sequence*:

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

Readers may sometimes think of short exact sequences as in the form $0 \to A \to B \to A/B \to 0$. However, we should remark here that this is not generally true in all categories. In fact, trivial kernels and cokernels are a little bit weaker than being monomorphic and epimorphic. Nontheless, fortunately we can show that these notions are equivalent in our interest categories $\mathfrak{Prshf}_X$ and $\mathfrak{Shf}_X$, and we don't have to worry about their distinction here.

Proposition 4.2. *Let $\eta : A \to B$ be a presheaf homomorphism. Then*

(1) η is monomorphic if and only if $\ker(\eta) = 0$*;*

(2) η is epimorphic if and only if $\operatorname{coker}_p(\eta) = 0$.

In particular, both statements also hold for sheaf homomorphisms (of course, we need to replace $\operatorname{coker}_p(\eta)$ *by* $\operatorname{coker}_s(\eta)$*).*

Proof. The forward direction of both statements are true in any categories with a zero, kernels and cokernels. Consider the following commutative diagrams:

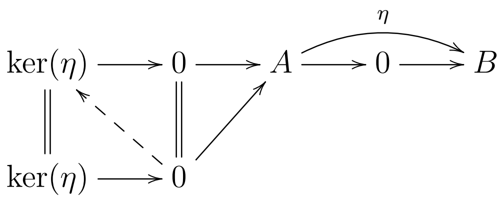

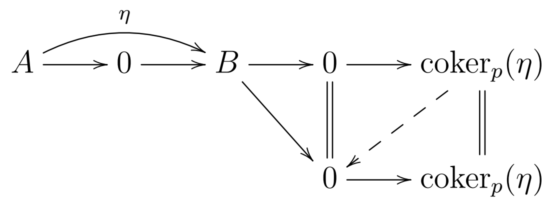

Since these two diagrams are dual to each other, we will just reason for the diagram on the left. The fact that η is a monomorphism guarantees the map $\ker(\eta) \to A$ factors through 0. But then by the universal property of kernels, we get a map from 0 back to $\ker(\eta)$, which can be easily shown to be an isomorphism. Since this proof is purely categorical, it automatically applies to the case of sheaf homomorphisms.

The backward direction deserves more careful treatment. Suppose ξ and ζ are two presheaf homomorphisms from presheaf C to presheaf A such that $\eta\xi = \eta\zeta : C \to B$. We can define a presheaf homomorphism $\xi - \zeta : C \to A$ by

$$(\xi - \zeta)_U = \eta_U - \xi_U.$$

It is not hard to verify that this is in fact a presheaf homomorphism. Moreover, it can be shown that $\eta(\xi - \zeta) = 0$. Therefore we get the commutative diagram below on the left from the universal property of kernels. But then this diagram tells us that $\xi - \zeta = 0$, i.e., $\xi_U - \zeta_U = 0$. Thus we can conclude that $\xi = \zeta$.

Since sheaf kernels coincides with presheaf kernels, the same argument and diagram works for the case of sheaf homomorphisms.

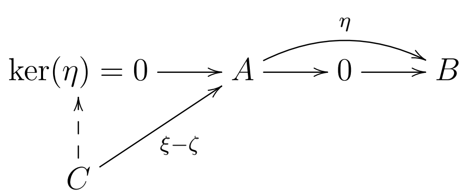

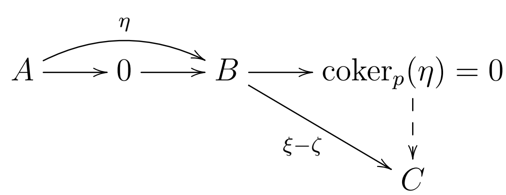

A similar argument can be formulated for the case of epimorphism. Let ξ and ζ be two presheaf homomorphisms from B to C such that $\xi\eta = \zeta\eta$. Define the presheaf homomorphism $\xi - \zeta$ as above and we can see that $\xi - \zeta = 0$ from the above commutative diagram on the right, which implies $\xi = \zeta$. Replacing $\text{coker}_p(\eta)$ with $\text{coker}_s(\eta)$ gives the argument to prove the backward direction for the case of sheaf homomorphism. ☐

So what does this proposition tell us? Suppose $0 \to A \xrightarrow{\eta} B \xrightarrow{\xi} C \to 0$ is a short exact sequence of presheaves. By definition, this implies that $\ker(\eta) = 0$, and by the proposition above, η is hence monomorphic, i.e., A is in fact a subpresheaf of B. Now we have another short exact sequence from the definition of quotient presheaf (cokernel of a monomorphism):

$$0 \longrightarrow A \xrightarrow{\eta} B \xrightarrow{q} (A/B)_p \longrightarrow 0 \, .$$

Notice that

$$\ker(q) = \text{img}_p(\eta) = \ker(\xi),$$

which implies that at each open set level, we have $\ker(q_U) = \ker(\xi_U)$. But then if we now compare the R-module homomorphism q_U and ξ_U, their images must be isomorphic! This shows that $C(U)$ is in fact isomorphic to $(A/B)_p$ at each open set level, and therefore we can conclude that $C \cong (A/B)_p$.

The same argument works for the case of short exact sequence of sheaves, except that we need to use $\mathrm{img}_s(\eta)$ instead of $\mathrm{img}_p(\eta)$.

Thus we can now safely say that short exact sequece of presheaves (sheaves) does take the form

$$0 \longrightarrow A \longrightarrow B \longrightarrow (A/B)_* \longrightarrow 0$$

($*$ is either p or s, depending on the context). The following corollary is another way of stating the same thing:

Corollary 4.3. *If* $0 \to A \xrightarrow{\eta} B \xrightarrow{\xi} C \to 0$ *is a short exact sequence of presheaves (sheaves), then we have* $A = \ker(\xi)$ *and* $C = \mathrm{coker}_p(\eta)$ *(*$C = \mathrm{coker}_s(\eta)$*).*

In many occasions when we have two categories where exactness is defined in both, we may want to ask the question how well a functor preserves exactness. Hence mathematicians came up with the following definition:

Definition 4.4. Let $\mathfrak{C}$ and $\mathfrak{D}$ be two categories with zero and exactness and let $0 \to A \to B \to C \to 0$ be a short exact sequence in the category $\mathfrak{C}$. We say a functor $F : \mathfrak{C} \to \mathfrak{D}$ (zero preserving) to be

(1) *left-exact* if the sequence $0 \to FA \to FB \to FC$ is exact;
(2) *right-exact* if the sequence $FA \to FB \to FC \to 0$ is exact;
(3) *exact* if F is both left-exact and right-exact.

Let's try to apply this definition to the two functors between $\mathfrak{Prshf}_X$ and $\mathfrak{Shf}_X$ that we know, namely the sheafification functor ΓL and the inclusion functor I, and see if they preserve exactness or not.

Examples 4.5. (1) Let $0 \to A \xrightarrow{\eta} B \xrightarrow{\xi} C \to 0$ be a short exact sequence of presheaves, then by Corollary 4.3, we know that $C = \mathrm{coker}_p(\eta)$. Sheafifying $A \xrightarrow{\eta} B \to \mathrm{coker}_p(\eta) \to 0$, we get

$$\Gamma LA \xrightarrow{\Gamma L\eta} \Gamma LB \longrightarrow \mathrm{coker}_s(\Gamma L\eta) \longrightarrow 0$$

in which $\mathrm{coker}_s(\Gamma L\eta) = \Gamma L\,\mathrm{coker}_p(\eta) = \Gamma LC$ since ΓL is a left adjoint functor. Thus the sheafification functor is at least right-exact.

How about left-exactness? Well, it turns out that the sheafification functor does also preserve left-exactness and hence it is exact. However, this fact is more of an algebraic result than a categorical result, and we would omit the proof here.

(2) On the other hand, if we start with a short exact sequence of sheaves $0 \to A \xrightarrow{\eta} B \xrightarrow{\xi} C \to 0$, Corollary 4.3 tells us that $A = \ker(\xi)$. Including $0 \to \ker(\xi) \to B \xrightarrow{\xi} C$ it into the category $\mathfrak{Prshf}_X$, nothing changes (recall that presheaf kernel and sheaf kernel are identical for a sheaf homomorphism). Thus the inclusion functor I is left-exact.

Remark 4.6. In the examples above, the only categorical property we used is Theorem 3.16. Thus the proof can be generalized to any adjoint pair of functors between categories of sheaves or presheaves, and gives the conclusion that the left adjoint functor being right-exact and the right adjoint functor being left-exact.

5. Induced Sheaf

So far we have been staying in a single topological space X and working on presheaves or sheaves over it. But since the motivation behind sheaf theory is to compute cohomology, we should be able to related sheaves over different topological spaces, especially the ones where there are continuous maps between them. Throughout this section, we will use X and Y for topological spaces and $f : X \to Y$ for a continuous map between them.

5.1. **Direct Image.** We should start with a simple observation: since f is continuous, the preimage of any open set in Y is still a open set in X, and taking preimages should preserve inclusion. Thus, f naturally induces a functor f^{-1} from $\mathfrak{Opn}_Y$ to $\mathfrak{Opn}_X$. Moreover, we can flip the arrows in both categories and f^{-1} retains its functoriality. Now let A be a presheaf over X and consider the following

$$\mathfrak{Opn}_Y^{op} \xrightarrow{f^{-1}} \mathfrak{Opn}_X^{op} \xrightarrow{A} \mathfrak{Mod}_R$$

It seems to be natural to just pre-compose f^{-1} with a presheaf A over X to get a presheaf $f^{-1} \circ A$ over Y! Furthermore, for a presheaf homomorphism (natural transformation) η from presheaf A over X to a presheaf B over X, we have the following for open sets $U \subset V \subset Y$:

$$\begin{array}{ccc} f^{-1} \circ A(U) = A\left(f^{-1}(U)\right) & \xrightarrow{\eta_{f^{-1}(U)}} & B\left(f^{-1}(U)\right) = f^{-1} \circ B(U) \\ \downarrow & & \downarrow \\ f^{-1} \circ A(V) = A\left(f^{-1}(V)\right) & \xrightarrow[\eta_{f^{-1}(V)}]{} & B\left(f^{-1}(V)\right) = f^{-1} \circ B(V) \end{array}$$

Thus $\eta_{f^{-1}(-)}$ gives a good natural transformation from $f^{-1} \circ A$ to $f^{-1} \circ B$. We can now thereby define a functor $f_* : \mathfrak{Prshf}_X \to \mathfrak{Prshf}_Y$ by assigning $f_*(A) = f^{-1} \circ A$ and $f_*\eta = \eta_{f^{-1}(-)}$. This functor is called the *direct image* of presheaves over X.

It is also not hard to verify that the direct image of a sheaf is still a sheaf, for the reason that the preimage of an open cover in the codomain is still an open cover in the domain. Therefore f_* is also a functor from $\mathfrak{Shf}_X$ to $\mathfrak{Shf}_Y$.

Example 5.1. One particular case is when we map the whole space X into a singleton space $\{*\}$. The preimage of the only non-trivial open set in $\{*\}$ is the whole set X, and therefore the functor f_* actually takes a sheaf A and gives out the "global section" $A(X)$ (on the empty set level, any sheaf is trivial, and therefore we may as well disregard that). Equivalently speaking, this functor can be regarded as a functor from $\mathfrak{Shf}_X$ to $\mathfrak{Mod}_R$, which we would denote it by G.

5.2. **Inverse Image.** There is actually another way to induce a sheaf. Suppose instead, we only know a presheaf B over the codomain space Y, and we would want to know whether it is possible to induce a presheaf over the domain X. Obviously, we don't have a natural way of associating an open set in X to a R-module just by looking at the map f, since the continuous image of an open set may not be open. Nonetheless, we can mimic the construction of stalks and get a direct limit over each open set in X, and hopefully this will give us some fruitful result.

The detailed procedure goes as follows. For an open set $U \subset X$, we can look at its image $f(U) \subset Y$; all open sets in Y containing $f(U)$ will give rise to a directed system

$$\{B(W) \mid f(U) \subset W\}$$

Now take the direct limit (which we know to be exist from Section 3) and name it $f^p B(U)$. Notice that if $U \subset V$, then any open set containing $f(V)$ will automatically contain $f(U)$. Thus, we obtain a natural restriction map from the universal property of direct limit (suppose $f(V) \subset f(U) \subset W_2 \subset W_1$):

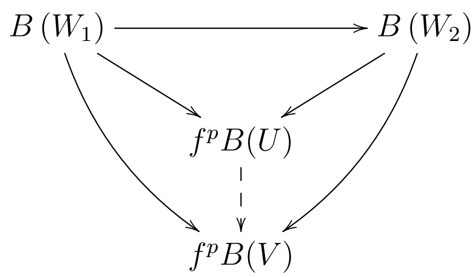

Thus from this construction, we can see that $f^p B$ is a presheaf over X.

Moreover, for a presheaf homomorphism $\eta : B \to B'$ where B and B' are both presheaves over Y, we can also get a map from $f^p B(U)$ to $f^p B'(U)$ from the universal property of direct limit, which we shall name $f^p \eta$ (suppose $f(U) \subset W$):

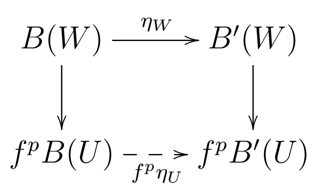

It is not hard to verify that $f^p \operatorname{Id}_U = \operatorname{Id}_U$ and $f^p(\eta \circ \xi)_U = f^p \eta_U \circ f^p \xi_U$ by using the universal property of direct limit again. Therefore we can conclude that f^p is in fact a functor from the category $\mathfrak{Prshf}_Y$ to $\mathfrak{Prshf}_X$.

Now the next question is whether we still get a sheaf over X if we apply f^p to a sheaf over Y. Unfortunately, the answer is not positive as we may expect. Consider the following example:

Example 5.2. Let X and Y both be the circle S^1 with the usual topology. Upon Y, construct

$$B = [0,1] \times \mathbb{Z} \Big/ \{(0,n) \sim (1,-n)\}$$

and define a map $p : B \to Y$ by sending (m,n) to $2m\pi$. If we check with Theorem 3.13, this defines a sheaf B over Y. Notice that by simple calculation, $B(Y) = 0$ and $B(\varphi, \psi) = \mathbb{Z}$ for any open arc (φ, ψ).

Now consider the map $f : X \to Y$ that wraps the circle twice around itself, i.e., $\varphi \mapsto 2\varphi$. The induced presheaf $f^p B$ over X has the following properties:

(1) $B(X) = 0$, since the only open set that contains $f(X)$ is Y.

(2) $B(\varphi,\psi)=\mathbb{Z}$ whenever the length of (φ,ψ) is less than π, since it will be mapped to an open arc in Y under f.

Now we can devide the circle X into three open arcs as shown in the picture. These three arcs form an open cover of X. However, the sheaf condition is not satisfied by f^pB: the equalizer of the following map is $\mathbb{Z}$, not $B(X)=0$.

$$\underset{0}{\mathbb{Z}}\times\underset{1}{\mathbb{Z}}\times\underset{2}{\mathbb{Z}}\rightrightarrows\underset{0\cap 0}{\mathbb{Z}}\times\underset{0\cap 1}{\mathbb{Z}}\times\underset{0\cap 2}{\mathbb{Z}}\times\cdots\times\underset{2\cap 2}{\mathbb{Z}}$$

Thus, if we want a functor that goes from $\mathfrak{Shf}_Y$ to $\mathfrak{Shf}_X$, we need more than just f^p. One easy way to get to the correct destination is to apply the sheafification functor right after f^p, and we would call this the "inverse image" functor f^*, i.e., $f^*=\Gamma L f^p$. The following is an example as well as a new way of looking at one of our old friends:

Example 5.3. If we let X be the singleton space $\{*\}$ and map this single point to a point $y\in Y$, then the inverse image of a presheaf B over Y is just the direct limit of the directed system

$$\{B(W)\mid y\in W\}\,,$$

which is exactly the stalk B_y! Therefore stalks are just a special kind of inverse image.

5.3. **Adjunction.** We have introduced two induced sheaf functors f_* and f^* for any continuous map $f:X\to Y$. Since these two functors are going forward and backward between $\mathfrak{Shf}_X$ and $\mathfrak{Shf}_Y$, we may have a feeling that it would be great if they were adjoints. In the rest of this section, we are going to examine this conjecture and see if it is truly the case.

Let A and B be two sheaves over X and Y respectively. Essentially we would like to show that there is an isomorphism between $\mathfrak{Shf}_X(f^*B,A)$ and $\mathfrak{Shf}_Y(B,f_*A)$. Notice that $f^*=\Gamma L f^p$, and by the $\Gamma L\dashv I$ adjunction, $\mathfrak{Shf}_X(f^*B,A)\cong\mathfrak{Prshf}_X(f^pB,A)$. Thus we just need an isomorphism between $\mathfrak{Prshf}_X(f^pB,A)$ and $\mathfrak{Shf}_Y(B,f^*A)$.

First suppose η is a map in $\mathfrak{Shf}_Y(B,f_*A)$. This means that η_W is an R-module homomorphism from $B(W)$ to $f_*A(W)=A\left(f^{-1}(W)\right)$, for an open set $W\subset Y$. By definition, $f^pB(U)$ is the direct limit of

$$\{B(W)\mid f(U)\subset W\}.$$

Thus we have the following commutative diagram, where the arrow on the right is the restriction map and the arrow at the bottom is given by the universal property of direct limits, which we would name $(\Phi(\eta))_U$.

$$\begin{array}{ccc} B(W) & \xrightarrow{\eta_W} & A\left(f^{-1}(W)\right) \\ {\scriptstyle\theta_U}\downarrow & & \downarrow \\ f^pB(U) & \underset{(\Phi(\eta))_U}{\dashrightarrow} & A(U) \end{array}$$

It is not hard to see that $(\Phi(\eta))_U$ is compatible with restriction maps in presheaves over X. Therefore $\Phi(\eta)$ is in fact a presheaf homomorphism from f^pB to A. Here Φ serves as a map from $\mathfrak{Shf}_Y(B,f_*A)$ to $\mathfrak{Prshf}_X(f^pB,A)$.

Next we would like to find an inverse map for Φ. Notice that if we take the direct image of f^pB,

$$f_*f^pB(W)=f^pB\left(f^{-1}(W)\right)=B(W).$$

Thus we have the identity $f_* f^p = \mathrm{Id}_{\mathfrak{Shf}_Y}$. This implies that for any presheaf homomorphism $\xi : f^p B \to A$, we can get a sheaf homomorphism

$$f_*\xi : f_* f^p B = B \to f_* A.$$

We claim that Φ and f_* are inverse of each other (here f_* is regarded as a map from $\mathfrak{Prshf}_X(f^p B, A)$ to $\mathfrak{Shf}_Y(B, f_* A)$). It is easy to see that $f_* \circ \Phi(\eta) = \eta$ by just looking at the following diagram:

$$\begin{array}{ccc} B(W) & \xrightarrow{\eta_W} & A\left(f^{-1}(W)\right) \\ \| & & \| \\ f^p B\left(f^{-1}(W)\right) & \underset{(\Phi(\eta))_{f^{-1}(W)}}{\dashrightarrow} & A\left(f^{-1}(W)\right) \end{array}$$

On the other hand, $\Phi \circ f_*(\xi) = \xi$ by the uniqueness from the universal property of direct limit (the map on the right is the sheaf restriction map in A)

$$\begin{array}{ccc} f_* f^p B(W) & \xrightarrow{f_*\xi_W} & f_* A(W) \\ \theta_U \downarrow & & \downarrow \\ f^p B(U) & \underset{(\Phi \circ f_*(\xi))_U}{\overset{\xi_U}{\rightrightarrows}} & A(U) \end{array}$$

Up to this point, we have successfully demonstrated that

$$\mathfrak{Prshf}_X(f^p B, A) \cong \mathfrak{Shf})_Y(B, f_* A).$$

As remarked previously, it follows immediately that

$$\mathfrak{Shf}_X(f^* B, A) \cong \mathfrak{Shf}_Y(B, f_* A).$$

In other words, we have the following theorem:

Theorem 5.4. *The inverse image functor $f^* : \mathfrak{Shf}_Y \to \mathfrak{Shf}_X$ and the direct image functor $f_* : \mathfrak{Shf}_X \to \mathfrak{Shf}_Y$ form an adjunction, i.e.,*

$$f^* \dashv f_*.$$

Since now we have another pair of adjoint functors between categories of sheaves, we can invoke Remark 4.6 from last section and get the following easy corollary:

Corollary 5.5. *The direct image functor f_* is left-exact and the inverse image functor f^* is right-exact. In particular, the global section functor G is left-exact and stalks are right-exact. (With algebraic tools, one can actually show that stalks are left-exact as well, but this is out of the scope of this paper and hence is omitted here.)*

6. A Brief Introduction to Sheaf Cohomology

We would like to conclude our paper with an important and some what categorical application of sheaf theory: sheaf cohomology. Since the existence of sheaf does not require simplicial or cellular structures on the topological space a priori, it is more general and natural. We will skip all the detailed proofs here, and yet try our best to give the general idea about what the construction does to formulate such a theory.

Categorically speaking, a cohomology theory can be constructed by first hitting an injective resolution with a left-exact functor(ideally not right-exact at the same time) and then measure the failure of resulting sequence being exact with the right-derived functor. Of course, we should start with giving the definitions of the related terms.

Definition 6.1. In a category $\mathfrak{C}$, an object I is said to be *injective* if whenever there is a monomorphism $A \to B$ together with an arrow $A \to I$, there always exists (not necessary unique) an arrow $B \to I$ such that the following diagram commutes:

$$\begin{array}{ccc} A & \longrightarrow & B \\ \downarrow & \swarrow & \\ I & & \end{array}$$

Definition 6.2. In a category $\mathfrak{C}$, a *resolution* of the object A is a graded long exact sequence

$$0 \longrightarrow A \longrightarrow K^0 \longrightarrow K^1 \longrightarrow K^2 \longrightarrow \dots$$

We say K^* is an *injective resolution* if all K^* are injective.

One good thing about sheaves is that the category $\mathfrak{Shf}_X$ over any topological space X possesses an abundant amount of injective objects, which always allow us to form an injective resolution (not necessary unique) started with any sheaf A:

$$0 \longrightarrow A \longrightarrow I^0 \longrightarrow I^1 \longrightarrow I^2 \longrightarrow \dots$$

Now in order to produce a cohomology theory, we need a left-exact functor. Fortunately we have already encountered one: recall from the last corollary in the previous section, the global section functor G is left-exact! Therefore after hitting the above sequence with G, we get a sequence

$$0 \longrightarrow A(X) \xrightarrow{d^{-1}} I^0(X) \xrightarrow{d^0} I^1(X) \xrightarrow{d^1} I^2(X) \xrightarrow{d^2} \dots$$

with the property that $d^k \circ d^{k-1} = 0$ for any $k \geq 0$. Thus as in the case of homology theories, the failure of being exact can be measured by

$$H^k_A(X) = \ker\left(d^k\right) \big/ \operatorname{img}_s\left(d^{k-1}\right).$$

As a result, we call $H^*_A(X)$ the *sheaf cohomology of X (with respect to the sheaf A)*. Of course, since injective resolution is not unique, one also need to show the well-defined-ness of this construction. The idea lies in the fact that the all objects I^* are injective. Starting with a sheaf homomorphism $\eta : A \to B$ in $\mathfrak{Shf}_X$, we can build a compatible set of maps between injective resolutions $0 \to A \to I^*$ and $0 \to B \to J^*$, which is unique up to homotopy. Detailed proof can be found in Bredon [1]:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \longrightarrow & I^0 & \longrightarrow & I^1 & \longrightarrow & \dots \\ & & \downarrow{\scriptstyle \eta} & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & B & \longrightarrow & J^0 & \longrightarrow & J^1 & \longrightarrow & \dots \end{array}$$

Just like singular cohomology, which is also naturally defined, sheaf cohomology may sometimes be very difficult to compute, since the construction of an injective resolution is generally non-trivial. Nevertheless, this construction is very important and useful. For example, the smooth maps from open sets on a manifold into $\mathbb{R}$ or $\mathbb{C}$ form a sheaf, and the sheaf cohomology defined above gives a natural cohomology theory on the particular manifold. We can also use the induced sheaf tools developed in the previous section to analyze cohomologies over spaces with continuous maps in between.

Conclusion and Acknowlegdment

Sheaf theory also has more advanced application in other topics such as the classification of line bundles, and now has become a very useful language to use in algebraic topology and algebraic geometry. We hope that through this paper, readers can see the most naturality of this language from a categorical perspective.

I need to thank Professor May and Emily Riehl for the inspirational talk on category theory, which leads to the idea of formulating sheaf theory using categorical language, and also for the help on clarifying problems that I have encountered.

I also want to thank Yiwei She for her advice on the paper topic and her help on revising my paper, without which my paper would be much inferior.

References

[1] Bredon, Glen E. *Sheaf Theory.* Springer, 1997.
[2] Tennison, B. R. *Sheaf Theory.* Cambridge University Press, 1975.
[3] Awodey, Steve. *Category Theory.* Oxford University Press, 2006