Colimits and Sheaves

Alex Massarenti

March 19, 2008

Abstract

These notes introduce colimits and sheaves from a category-theoretic point of view, with an eye toward the geometric and cohomological applications that motivate sheaf theory in modern mathematics. We review filtered colimits, direct limits, and basic universal properties, then explain how presheaves and sheaves organize local-to-global constructions on a topological space. Special attention is given to stalks as filtered colimits over neighborhoods, the sheaf condition as an equalizer diagram, and standard examples illustrating that natural presheaf constructions (such as images) need not be sheaves.

Contents

1 Introduction

Sheaf theory is one of the central languages through which "local" data are assembled into "global" objects. It provides a systematic framework for gluing functions, sections, or solutions of geometric and analytic problems, and it packages this local-to-global principle into a robust functorial formalism. From its classical origins in complex analysis and topology to its pervasive role in algebraic geometry, homological algebra, and the theory of topoi, the concept of a sheaf has evolved into a foundational tool. Classical treatments emphasizing the topological and analytic intuition can be found in Godement's monograph [God58] and in Tennison's introduction [Ten75], while Hartshorne's *Algebraic Geometry* [Har77] develops sheaves as the basic technical infrastructure for schemes and their cohomology. On the more categorical side, the sheaf condition and the fundamental constructions are naturally expressed in terms of limits and colimits, a point of view emphasized in Mac Lane [ML98] and developed extensively in Kashiwara–Schapira [KS06].

A guiding theme of these notes is that many everyday sheaf-theoretic operations are, at heart, statements about *universal properties.* The category-theoretic perspective clarifies not only *what* sheaves do, but also *why* standard definitions are the correct ones: the sheaf axioms can be recast as an equalizer diagram, stalks arise as filtered colimits over neighborhoods, and the passage from local sections to global sections is a functor whose exactness properties are best understood via derived functors and cohomology. This philosophy culminates, at a much deeper level, in Grothendieck's theory of topoi, where a topos is conceived as a generalized universe of sheaves and sheaf cohomology becomes a primary invariant [GAV71]. Even in introductory settings, adopting

categorical language early makes later developments (e.g. Čech cohomology, derived categories, and étale cohomology) feel far less ad hoc; for an accessible entry point into étale cohomology built on sheaf-theoretic foundations, see Tamme [Tam94].

Colimits as a computational and conceptual tool

Colimits occur whenever one forms an object by "gluing together" a diagram of objects. Among them, *filtered colimits* (often called *direct limits*) are especially important because they behave well with respect to finite limits in many algebraic categories, and because they model stabilization under refinement. A filtered colimit is the categorical way to encode the idea that we have a directed system of objects and compatible transition maps, and we want the object that receives compatible maps from all stages. The universal property packages this construction efficiently, but the real gain is conceptual: once a universal property is identified, many proofs reduce to checking that a candidate object satisfies the appropriate mapping property, rather than chasing elements.

In these notes, filtered colimits appear in two complementary ways. First, they occur explicitly as limits of directed systems in familiar algebraic categories (sets, groups, rings, modules), where one can still compute them by hand. Second, and more importantly, they appear *geometrically* in the definition of stalks. The stalk $\mathcal{F}_x$ of a presheaf (or sheaf) at a point x is defined as the colimit of the sections over all neighborhoods of x, ordered by inclusion. This definition is not merely convenient: it is the correct mechanism for extracting the "infinitesimal" local behavior of sections near x and for explaining why stalks turn sheaf axioms into local statements. The ubiquitous slogan "a sheaf is determined by its stalks" is, in large part, a reflection of the fact that stalks are built from filtered colimits.

Presheaves, sheaves, and the sheaf condition

Given a topological space X, a presheaf $\mathcal{F}$ associates to each open set $U \subseteq X$ a set (or group, or ring, etc.) $\mathcal{F}(U)$, together with restriction maps $\mathcal{F}(U) \to \mathcal{F}(V)$ for inclusions $V \subseteq U$. This is simply a contravariant functor from the poset category of opens to a target category. The passage from topology to category theory is already visible here: the structure is functorial, and the axioms of a presheaf say precisely that restrictions compose correctly and that restricting to the same open set does nothing. From this point of view, the definition is completely formal (cf. [ML98]).

A *sheaf* is a presheaf that satisfies an additional *gluing* condition: sections that agree on overlaps must glue uniquely to a section on the union. Concretely, for a cover $U = \bigcup_i U_i$, the usual sheaf axiom says that a compatible family $(s_i \in \mathcal{F}(U_i))_i$ with $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ arises from a unique $s \in \mathcal{F}(U)$. Categorically, this can be expressed as an equalizer condition

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

making the relationship with limits completely explicit; see, for instance, [Har77, §II.1] for the standard formulation in algebraic geometry, and [KS06] for a systematic categorical treatment.

This equalizer formulation is not just elegant: it is operational. It allows one to understand sheafification as a left adjoint to the inclusion of sheaves into presheaves, and it makes clear why certain naturally occurring presheaves are not sheaves. A particularly instructive example, discussed in the notes, is that the *image presheaf* of a morphism of sheaves need not satisfy the sheaf condition. The obstruction is genuinely local-to-global: locally one may lift a section, but globally there may be a monodromy or cohomological obstruction that prevents gluing within the image. This phenomenon anticipates the role of exact sequences of sheaves and the need for derived functors and cohomology in measuring such obstructions—one of the reasons sheaves became indispensable in geometry.

Why this matters: geometry, cohomology, and beyond

In algebraic geometry, the sheaf viewpoint is not optional: schemes are locally ringed spaces, and their structure is encoded in the structure sheaf. Properties of morphisms, divisors, and line bundles are naturally expressed as sheaf-theoretic statements. Hartshorne's text [Har77] is a canonical reference for this viewpoint,

developing cohomology as a tool to study global sections, embeddings, and intersection-theoretic phenomena. In a broader categorical direction, Grothendieck's insight was that the notion of a sheaf should be separated from ordinary topology and instead attached to a *site*, yielding the concept of a topos [GAV71]. This abstraction is not a mere generalization for its own sake: it is precisely what allows one to define étale cohomology for schemes, replacing classical topology by the étale site and recovering powerful invariants with arithmetic meaning (see [Tam94] for an introductory account).

The present notes stay at the level of topological spaces, but they are written so that the transition to sites and topoi feels natural: the key ideas are already visible. Covers, refinement, and compatibility are categorical notions; stalks are filtered colimits; and the sheaf axiom is a limit condition. In this sense, even elementary sheaf theory is a meeting ground between topology, algebra, and category theory, and the category-theoretic point of view is the common language that keeps the constructions coherent [ML98, KS06].

Organization of the notes

We begin by recalling directed and filtered categories and the construction of colimits in basic algebraic settings. We then introduce presheaves and sheaves, discuss stalks and germs via filtered colimits, and formulate the sheaf condition both in its classical gluing form and in its categorical equalizer form. We close with examples and standard pitfalls (such as images not being sheaves), which motivate the need for sheafification and foreshadow the cohomological viewpoint developed in the classical and modern references cited above [God58, Ten75, Har77, GAV71, Tam94, KS06].

2   Limits and colimits

Throughout this section $\mathcal{C}$ denotes a *preadditive* category (i.e. each Hom-set is an abelian group and composition is bilinear), and I denotes a small category.

Definition 1 (Compatible family)**.** Let $F : I \to \mathcal{C}$ be a covariant functor and let $X \in \mathcal{C}$. A family of morphisms $(f_i : X \to F(i))_{i \in I}$ is called *compatible* if for every morphism $g : i \to j$ in I one has

$$f_j = F(g) \circ f_i.$$

Equivalently, the triangles

commute for all $g : i \to j$.

Definition 2 (Projective limit)**.** Let $F : I \to \mathcal{C}$ be a covariant functor. A *(projective) limit* of F is a pair $(\varprojlim F, (f_i))$ consisting of an object $\varprojlim F \in \mathcal{C}$ and a compatible family of morphisms $f_i : \varprojlim F \to F(i)$ such that for every object $X \in \mathcal{C}$ and every compatible family $(g_i : X \to F(i))$ there exists a unique morphism $h : X \to \varprojlim F$ satisfying $f_i \circ h = g_i$ for all i.

Definition 3 (Inductive limit / colimit)**.** Let $F : I \to \mathcal{C}$ be a covariant functor. A *colimit* (or *inductive limit*) of F is a pair $(\varinjlim F, (f_i))$ consisting of an object $\varinjlim F \in \mathcal{C}$ and a family of morphisms $f_i : F(i) \to \varinjlim F$ which is compatible in the sense that for every morphism $\lambda : i \to j$ in I one has $f_j \circ F(\lambda) = f_i$, i.e. the diagram

commutes. Moreover, for every object $X \in \mathcal{C}$ and every compatible family $(g_i : F(i) \to X)$ there exists a unique morphism $h : \varinjlim F \to X$ such that $g_i = h \circ f_i$ for all i.

Definition 4. A preadditive category $\mathcal{C}$ is *complete* if every covariant functor $F : I \to \mathcal{C}$ admits a limit. It is *cocomplete* if every covariant functor $F : I \to \mathcal{C}$ admits a colimit.

Proposition 1 (Uniqueness up to isomorphism)**.** *If the colimit of a functor $F : I \to \mathcal{C}$ exists, then it is unique up to a unique isomorphism.*

Proof. Let $(X, (f_i))$ and $(Y, (g_i))$ be two colimits of F. Since (g_i) is compatible, the universal property of $(X, (f_i))$ yields a unique morphism $h : X \to Y$ with $h \circ f_i = g_i$. Similarly, the universal property of $(Y, (g_i))$ yields a unique morphism $q : Y \to X$ with $q \circ g_i = f_i$. Then

$$(hq) \circ g_i = h \circ (qg_i) = h \circ f_i = g_i \qquad \forall i,$$

so by uniqueness we get $hq = \mathrm{id}_Y$. Likewise $qh = \mathrm{id}_X$, hence h is an isomorphism with inverse q. $\qquad\square$

2.1 Examples of colimits

Example (Coproducts are colimits)*.* If I is a discrete category and $F : I \to \mathcal{C}$ is a functor, then a colimit of F is precisely a coproduct

$$\varinjlim F \cong \coprod_{i \in I} F(i),$$

with its canonical injections $f_i : F(i) \to \coprod_i F(i)$.

Example (Pushouts as colimits)*.* Suppose I is the "span" category $i \leftarrow k \to j$ (equivalently, a partially ordered set with $k \leq i$ and $k \leq j$ and no other relations). Then a colimit of $F : I \to \mathcal{C}$ is the pushout of the diagram

$$F(i) \xleftarrow{\ F(u)\ } F(k) \xrightarrow{\ F(v)\ } F(j).$$

More explicitly, if P is a pushout with structure maps $\alpha_i : F(i) \to P$ and $\alpha_j : F(j) \to P$, then $\alpha_i \circ F(u) = \alpha_j \circ F(v)$, and the universal property of the pushout is exactly the universal property of $\varinjlim F$ in this indexing shape.

Example (Cokernels as pushouts)*.* Let $f : C \to D$ be a morphism in a preadditive category with a zero object 0. Consider the pushout square

$$
\begin{array}{ccc}
C & \xrightarrow{\ f\ } & D \\
\downarrow & & \downarrow{\scriptstyle g} \\
0 & \longrightarrow & Q.
\end{array}
$$

Then $g : D \to Q$ is a cokernel of f. Indeed, the commutativity means $g \circ f = 0$, and the pushout universal property says that any $h_2 : D \to Y$ with $h_2 f = 0$ factors uniquely through g. Thus cokernels are special pushouts, hence special colimits.

Theorem 1. *Let $\mathcal{C}$ be a preadditive category with a zero object. Then $\mathcal{C}$ is cocomplete if and only if it admits all coproducts and all cokernels.*

Proof. We already observed that coproducts and cokernels are particular colimits, hence they exist in any cocomplete category.

Conversely, assume that $\mathcal{C}$ has all coproducts and all cokernels, and let $F : I \to \mathcal{C}$ be a functor. Form the coproducts

$$\left(\coprod_{i \in I} F(i),\ (f_i : F(i) \to \coprod_i F(i)) \right), \qquad \left(\coprod_{\lambda : i \to j} F(i),\ (g_\lambda : F(i) \to \coprod_\lambda F(i)) \right),$$

where in the second coproduct the index runs over all morphisms $\lambda : i \to j$ of I (so the summand is $F(s(\lambda)) = F(i)$).

For each morphism $\lambda : i \to j$ define a map

$$\ell_\lambda := f_i - f_j \circ F(\lambda) : \ F(i) \longrightarrow \coprod_{i \in I} F(i).$$

By the universal property of the second coproduct there exists a unique morphism

$$L : \coprod_{\lambda : i \to j} F(i) \longrightarrow \coprod_{i \in I} F(i)$$

such that $L \circ g_\lambda = \ell_\lambda$ for all λ. Let $p : \coprod_i F(i) \to \operatorname{coker}(L)$ be the cokernel of L. Set $\tilde{f}_i := p \circ f_i : F(i) \to \operatorname{coker}(L)$.

Compatibility. Since $pL = 0$, for each $\lambda : i \to j$ we have

$$0 = p\ell_\lambda = p(f_i - f_j F(\lambda)) = \tilde{f}_i - \tilde{f}_j \circ F(\lambda),$$

hence $\tilde{f}_j \circ F(\lambda) = \tilde{f}_i$.

Universal property. Let $(h_i : F(i) \to Y)$ be any compatible family, i.e. $h_j F(\lambda) = h_i$ for all $\lambda : i \to j$. By the coproduct property there exists a unique $\alpha : \coprod_i F(i) \to Y$ with $\alpha f_i = h_i$. We claim $\alpha L = 0$. Indeed, for each $\lambda : i \to j$,

$$\alpha \ell_\lambda = \alpha(f_i - f_j F(\lambda)) = h_i - h_j F(\lambda) = 0,$$

so $\alpha L g_\lambda = 0$ for all λ, hence $\alpha L = 0$. Therefore α factors uniquely through the cokernel p, i.e. there exists a unique $\chi : \operatorname{coker}(L) \to Y$ such that $\chi p = \alpha$. Then $\chi \tilde{f}_i = \chi p f_i = \alpha f_i = h_i$ for all i, proving that $(\operatorname{coker}(L), (\tilde{f}_i))$ is a colimit of F. □

2.2 Functoriality of the colimit

Assume now that $\mathcal{C}$ is cocomplete. Let $F, G : I \to \mathcal{C}$ be functors and $\varphi : F \Rightarrow G$ a natural transformation. Write $(\varinjlim F, (f_i))$ and $(\varinjlim G, (g_i))$ for colimits. Naturality gives $g_j \varphi_j F(\lambda) = g_j G(\lambda) \varphi_i = g_i \varphi_i$, so the family $(g_i \varphi_i)$ is compatible for F. Hence there exists a unique morphism

$$\varinjlim(\varphi) : \varinjlim F \longrightarrow \varinjlim G$$

such that $\varinjlim(\varphi) \circ f_i = g_i \circ \varphi_i$ for all i.

Proposition 2. *The assignment*

$$\varinjlim : \operatorname{Hom}(I, \mathcal{C}) \longrightarrow \mathcal{C}, \qquad F \longmapsto \varinjlim F, \quad (\varphi : F \Rightarrow G) \longmapsto \varinjlim(\varphi)$$

defines a covariant functor.

Proof. Let $F \overset{\varphi}{\Rightarrow} G \overset{\psi}{\Rightarrow} H$ be natural transformations. Let $(\varinjlim F, (f_i))$, $(\varinjlim G, (g_i))$, $(\varinjlim H, (h_i))$ be colimits. By definition, $\varinjlim(\psi \circ \varphi)$ is the unique map satisfying $\varinjlim(\psi \circ \varphi) f_i = h_i(\psi_i \circ \varphi_i)$. On the other hand,

$$\big(\varinjlim(\psi) \circ \varinjlim(\varphi)\big) f_i = \varinjlim(\psi) (g_i \varphi_i) = h_i \psi_i \varphi_i = h_i(\psi_i \circ \varphi_i),$$

so by uniqueness $\varinjlim(\psi \circ \varphi) = \varinjlim(\psi) \circ \varinjlim(\varphi)$. Similarly, $\varinjlim(\operatorname{id}_F) = \operatorname{id}_{\varinjlim F}$. □

Theorem 2 (Right exactness)**.** *Let $\mathcal{C}$ be a preadditive category in which cokernels exist (e.g. an abelian category), and assume $\mathcal{C}$ is cocomplete. Then for every small category I the functor $\varinjlim : \operatorname{Hom}(I, \mathcal{C}) \to \mathcal{C}$ is right exact: if*

$$F \overset{\varphi}{\Rightarrow} G \overset{\psi}{\Rightarrow} H \longrightarrow 0$$

is pointwise exact (i.e. $F(i) \to G(i) \to H(i) \to 0$ is exact for all i), then

$$\varinjlim F \xrightarrow{\varinjlim(\varphi)} \varinjlim G \xrightarrow{\varinjlim(\psi)} \varinjlim H \longrightarrow 0$$

is exact at the last two terms.

Proof. We first show that $\varinjlim(\psi) : \varinjlim G \to \varinjlim H$ is an epimorphism. Let $v : \varinjlim H \to X$ satisfy $v \circ \varinjlim(\psi) = 0$. For each i we have $v\, h_i\, \psi_i = v\, \varinjlim(\psi)\, g_i = 0$. Since ψ_i is epi, this implies $v h_i = 0$ for all i. By the universal property of $\varinjlim H$, the family $(v h_i)$ determines v uniquely, and the zero family corresponds to 0, hence $v = 0$. Thus $\varinjlim(\psi)$ is epi.

Next we show $\mathrm{Im}(\varinjlim(\varphi)) = \mathrm{Ker}(\varinjlim(\psi))$. Since $\psi \circ \varphi = 0$ pointwise, we have $\varinjlim(\psi) \circ \varinjlim(\varphi) = \varinjlim(\psi \circ \varphi) = 0$, hence $\mathrm{Im}(\varinjlim(\varphi)) \subseteq \mathrm{Ker}(\varinjlim(\psi))$.

Conversely, let $\xi : \varinjlim G \to X$ satisfy $\xi \circ \varinjlim(\varphi) = 0$. Then $\xi g_i \varphi_i = \xi\, \varinjlim(\varphi)\, f_i = 0$, so ξg_i kills $\mathrm{Im}(\varphi_i) = \mathrm{Ker}(\psi_i)$. By pointwise exactness, ψ_i is the cokernel of φ_i, hence there exists a unique $\xi_i : H(i) \to X$ with $\xi_i \psi_i = \xi g_i$.

We claim the family (ξ_i) is compatible for H. If $\lambda : i \to j$, naturality of ψ gives $H(\lambda)\psi_i = \psi_j G(\lambda)$, hence

$$\xi_j H(\lambda)\psi_i = \xi_j \psi_j G(\lambda) = \xi g_j G(\lambda) = \xi g_i = \xi_i \psi_i.$$

Since ψ_i is epi, we get $\xi_j H(\lambda) = \xi_i$. Therefore, by the universal property of $\varinjlim H$, there exists a unique $\bar{\xi} : \varinjlim H \to X$ with $\bar{\xi} h_i = \xi_i$. Then for each i,

$$\bar{\xi}\, \varinjlim(\psi)\, g_i = \bar{\xi}\, h_i\, \psi_i = \xi_i \psi_i = \xi g_i,$$

and by the universal property of $\varinjlim G$ we conclude $\bar{\xi}\, \varinjlim(\psi) = \xi$. This shows that $\varinjlim(\psi)$ is the cokernel of $\varinjlim(\varphi)$, hence the sequence is right exact. $\qquad\square$

2.3 Directed systems and an explicit construction in modules

Definition 5 (Directed sets). A partially ordered set $(I, \leq)$ is *directed* if for every $i, j \in I$ there exists $k \in I$ such that $i \leq k$ and $j \leq k$. We view I as a small category by declaring that there is a unique morphism $i \to j$ precisely when $i \leq j$.

Definition 6 (Direct system). Let $(I, \leq)$ be a directed set and let $\mathcal{C}$ be a category. A *direct system* in $\mathcal{C}$ indexed by I consists of objects $(C_i)_{i \in I}$ and morphisms $f_i^j : C_i \to C_j$ for all $i \leq j$ such that $f_i^i = \mathrm{id}_{C_i}$ and $f_i^k = f_j^k \circ f_i^j$ whenever $i \leq j \leq k$. Equivalently, a direct system is a covariant functor $F : I \to \mathcal{C}$. Its colimit is called the *direct limit* of the system. Similarly, an *inverse system* is a functor $I^{\mathrm{op}} \to \mathcal{C}$, and its limit is called an *inverse limit*.

We now describe an explicit construction of direct limits in the category $A\text{-}\mathbf{Mod}$ of modules over a ring A.

Proposition 3. *Let (X_i, f_i^j) be a direct system of A-modules indexed by a directed set $(I, \leq)$. Then the direct limit $\varinjlim X_i$ can be realized as follows. Let $X = \bigsqcup_{i \in I} X_i$ be the disjoint union and define an equivalence relation $\sim$ on X by*

$$x_i \in X_i \ \sim \ x_j \in X_j \quad \Longleftrightarrow \quad \exists k \in I \text{ with } i \leq k,\ j \leq k \text{ and } f_i^k(x_i) = f_j^k(x_j).$$

Then $X/\sim$ carries a natural structure of A-module, and the canonical maps $\varepsilon_i : X_i \to X/\sim$ form a colimiting cocone:

$$\varinjlim_{i \in I} X_i \ \cong \ X/\sim.$$

Proof. We sketch the construction, emphasizing the universal property.

Step 1: the equivalence relation. Reflexivity is clear (take $k = i$). Symmetry is immediate. For transitivity, assume $x_i \sim x_j$ and $x_j \sim x_t$. Then there exist k, h with $i, j \leq k$ and $j, t \leq h$ such that $f_i^k(x_i) = f_j^k(x_j)$ and $f_j^h(x_j) = f_t^h(x_t)$. Since I is directed, choose ℓ with $k \leq \ell$ and $h \leq \ell$. Then $f_i^\ell(x_i) = f_j^\ell(x_j) = f_t^\ell(x_t)$, hence $x_i \sim x_t$.

Step 2: module structure. Let $\widehat{x}_i$ denote the class of $x_i \in X_i$ in $X/\sim$. Given $\widehat{x}_i, \widehat{y}_j$, choose k with $i \leq k$ and $j \leq k$, and set

$$\widehat{x}_i + \widehat{y}_j \ := \ \widehat{f_i^k(x_i) + f_j^k(y_j)} \ \in X/\sim.$$

Similarly define scalar multiplication by $a\widehat{x}_i := \widehat{ax}_i$. One checks (using the directedness and the compatibility of the maps f_i^j) that these operations are well defined and turn $X/\sim$ into an A-module.

6

Step 3: the universal property. Define $\varepsilon_i : X_i \to X/\!\sim$ by $\varepsilon_i(x_i) = \widehat{x}_i$. Compatibility $\varepsilon_j \circ f_i^j = \varepsilon_i$ follows from the definition of $\sim$.

Now let $(\xi_i : X_i \to Y)$ be any compatible family of A-linear maps, i.e. $\xi_j f_i^j = \xi_i$. Define $\xi : X/\!\sim \to Y$ by $\xi(\widehat{x}_i) = \xi_i(x_i)$. If $x_i \sim x_j$, choose k as in the definition and compute

$$\xi_i(x_i) = \xi_k(f_i^k(x_i)) = \xi_k(f_j^k(x_j)) = \xi_j(x_j),$$

so ξ is well defined. It is A-linear by construction, and satisfies $\xi\varepsilon_i = \xi_i$. Uniqueness is immediate because the classes $\widehat{x}_i$ generate $X/\!\sim$. Therefore $(X/\!\sim, (\varepsilon_i))$ is a colimit, hence $X/\!\sim \cong \varinjlim X_i$. $\qquad\square$

3 Sheaves

Let X be a topological space. We work primarily with sheaves of abelian groups, although the formalism extends to presheaves valued in other categories.

Definition 7 (Presheaf)**.** A *presheaf* $\mathcal{F}$ of abelian groups on X consists of:

- an abelian group $\mathcal{F}(U)$ for every open set $U \subseteq X$;

- for every inclusion of open sets $V \subseteq U$ a homomorphism (restriction map) $\rho_{U,V} : \mathcal{F}(U) \to \mathcal{F}(V)$,

such that:

1. $\mathcal{F}(\varnothing) = 0$;

2. $\rho_{U,U} = \mathrm{id}_{\mathcal{F}(U)}$ for all U;

3. if $W \subseteq V \subseteq U$ then $\rho_{U,W} = \rho_{V,W} \circ \rho_{U,V}$.

We often write $s|_V$ for $\rho_{U,V}(s)$ when $s \in \mathcal{F}(U)$.

Remark. Let $\mathbf{Top}(X)$ be the category whose objects are open subsets of X and whose morphisms are inclusions. Then a presheaf of abelian groups is equivalently a contravariant functor $\mathbf{Top}(X)^{\mathrm{op}} \to \mathbf{Ab}$. More generally, a presheaf with values in a category $\mathcal{C}$ is a contravariant functor $\mathbf{Top}(X)^{\mathrm{op}} \to \mathcal{C}$.

Definition 8 (Sheaf)**.** A presheaf $\mathcal{F}$ on X is a *sheaf* if it satisfies:

1. (*locality*) if U is open and $\{V_i\}$ is an open cover of U and $s \in \mathcal{F}(U)$ satisfies $s|_{V_i} = 0$ for all i, then $s = 0$;

2. (*gluing*) if U is open and $\{V_i\}$ is an open cover of U and for each i there is $s_i \in \mathcal{F}(V_i)$ such that $s_i|_{V_i \cap V_j} = s_j|_{V_i \cap V_j}$ for all i, j, then there exists $s \in \mathcal{F}(U)$ with $s|_{V_i} = s_i$ for all i.

The element s in (2) is automatically unique by (1).

Example (Regular functions)*.* Let X be an algebraic variety over a field K. For each open $U \subseteq X$ let $\mathcal{O}_X(U)$ be the ring of regular functions $U \to K$, and define restriction by restriction of functions. Then $\mathcal{O}_X$ is a sheaf (the *structure sheaf*).

Example (A presheaf that is not a sheaf)*.* Let $\mathcal{L}$ be the presheaf on $\mathbb{R}$ assigning to each open set U the abelian group of bounded functions $U \to \mathbb{R}$, with restriction given by restriction of functions. This presheaf is not a sheaf. Indeed, take $U = (0, 1)$ and the open cover $V_n = (\frac{1}{n}, 1)$. The functions $s_n(x) = \frac{1}{x}$ are bounded on V_n and agree on overlaps, but they do not glue to a bounded function on U.

3.1 Stalks and germs

Fix a point $x \in X$. For every open neighbourhood $U \ni x$ there is a restriction map to smaller neighbourhoods. Thus $(\mathcal{F}(U))_{U \ni x}$ is a direct system in $\mathbf{Ab}$.

Definition 9 (Stalk)**.** Let $\mathcal{F}$ be a presheaf on X and $x \in X$. The *stalk* of $\mathcal{F}$ at x is the direct limit

$$\mathcal{F}_x := \varinjlim_{x \in U} \mathcal{F}(U)$$

taken over the directed system of open neighbourhoods U of x ordered by reverse inclusion. Elements of $\mathcal{F}_x$ are called *germs* at x.

Concretely, a germ is represented by a pair (U, s) with $x \in U$ and $s \in \mathcal{F}(U)$, where (U, s) and (V, t) represent the same germ if there exists an open $W \subseteq U \cap V$ with $x \in W$ such that $s|_W = t|_W$. We write $s_x \in \mathcal{F}_x$ for the germ of s at x.

Example. If $\mathcal{F} = \mathcal{O}_X$ is the structure sheaf of a variety, then $\mathcal{O}_{X,x}$ is the local ring at x.

3.2 Morphisms of (pre)sheaves

Definition 10 (Morphism of presheaves)**.** Let $\mathcal{F}$ and $\mathcal{G}$ be presheaves on X. A *morphism of presheaves* $\varphi : \mathcal{F} \to \mathcal{G}$ consists of homomorphisms $\varphi(U) : \mathcal{F}(U) \to \mathcal{G}(U)$ for all open $U \subseteq X$ such that for every inclusion $V \subseteq U$ the diagram

$$
\begin{array}{ccc}
\mathcal{F}(U) & \xrightarrow{\varphi(U)} & \mathcal{G}(U) \\
\rho^{\mathcal{F}}_{U,V} \downarrow & & \downarrow \rho^{\mathcal{G}}_{U,V} \\
\mathcal{F}(V) & \xrightarrow[\varphi(V)]{} & \mathcal{G}(V)
\end{array}
$$

commutes. The same definition applies to sheaves.

Every morphism of presheaves induces a morphism on stalks.

Proposition 4. *Let* $\varphi : \mathcal{F} \to \mathcal{G}$ *be a morphism of presheaves and* $x \in X$. *Then there is an induced homomorphism of stalks*

$$\varphi_x : \mathcal{F}_x \longrightarrow \mathcal{G}_x, \qquad s_x \longmapsto (\varphi(s))_x,$$

and this is precisely the map $\varinjlim \varphi(U)$ *on direct limits.*

3.3 Subsheaves, kernels, images, cokernels

Definition 11 (Subsheaf)**.** A *subsheaf* $\mathcal{F}' \subseteq \mathcal{F}$ is a sheaf such that $\mathcal{F}'(U) \subseteq \mathcal{F}(U)$ for all open U, and the restriction maps of $\mathcal{F}'$ are the restrictions of those of $\mathcal{F}$.

Let $\varphi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves.

- The presheaf $\mathrm{Ker}(\varphi)$ given by $U \mapsto \mathrm{Ker}(\varphi(U))$ is a subsheaf of $\mathcal{F}$.

- The presheaf $\mathrm{Im}(\varphi)$ given by $U \mapsto \mathrm{Im}(\varphi(U))$ is in general *not* a sheaf.

- The presheaf $\mathrm{coker}(\varphi)$ given by $U \mapsto \mathrm{coker}(\varphi(U))$ is in general *not* a sheaf.

Proposition 5. *If* $\varphi : \mathcal{F} \to \mathcal{G}$ *is a morphism of sheaves, then the presheaf* $\mathrm{Ker}(\varphi)$ *is a sheaf.*

Proof. Let U be open and $\{V_i\}$ an open cover of U.

(1) If $s \in \mathrm{Ker}(\varphi)(U) \subseteq \mathcal{F}(U)$ restricts to 0 in each $\mathrm{Ker}(\varphi)(V_i)$, then $s|_{V_i} = 0$ in $\mathcal{F}(V_i)$ for all i. Since $\mathcal{F}$ is a sheaf, $s = 0$.

(2) Suppose $s_i \in \mathrm{Ker}(\varphi)(V_i)$ satisfy $s_i|_{V_i \cap V_j} = s_j|_{V_i \cap V_j}$. Since $\mathcal{F}$ is a sheaf, there exists $s \in \mathcal{F}(U)$ with $s|_{V_i} = s_i$. Applying φ and using commutativity with restriction, we get $\varphi(s)|_{V_i} = \varphi(s_i) = 0$ for all i. Since $\mathcal{G}$ is a sheaf, $\varphi(s) = 0$, hence $s \in \mathrm{Ker}(\varphi)(U)$. $\qquad\square$

Example (The image presheaf need not be a sheaf)*.* Let $X = \mathcal{C}$ with its usual topology, $\mathcal{F} = \mathcal{O}_\mathcal{C}$ the sheaf of holomorphic functions, and $\mathcal{G} = \mathcal{O}_\mathcal{C}^\times$ the sheaf of nowhere vanishing holomorphic functions. The exponential map defines a morphism of sheaves $\exp : \mathcal{F} \to \mathcal{G}$. On the open set $U = \mathcal{C}^\times$ take the cover $U = V \cup W$ where $V = \mathcal{C} \setminus \mathbb{R}_{\leq 0}$ and $W = \mathcal{C} \setminus \mathbb{R}_{\geq 0}$. On each of V and W there exists a holomorphic logarithm, so the identity section on V and on W lies in the image presheaf. These two identity sections agree on $V \cap W$, but they do *not* glue to a global section in the image presheaf on U, because there is no holomorphic logarithm on all of $\mathcal{C}^\times$. Therefore $\mathrm{Im}(\exp)$ is not a sheaf.
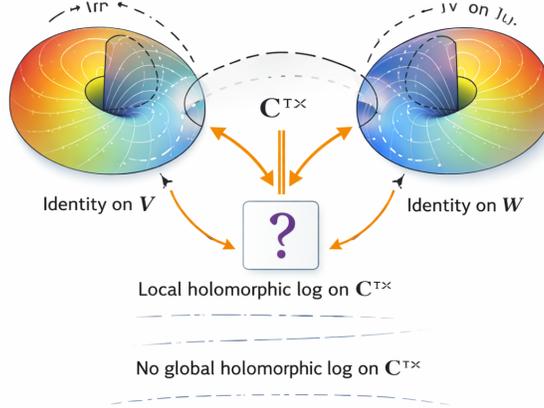
The figure above illustrates why the image presheaf $\mathrm{Im}(\exp) \subseteq \mathcal{O}_\mathbb{C}^\times$ fails to be a sheaf on $U = \mathbb{C}^\times$. The open sets $V = \mathbb{C} \setminus \mathbb{R}_{\leq 0}$ and $W = \mathbb{C} \setminus \mathbb{R}_{\geq 0}$ remove a branch cut, so on each of them there exists a holomorphic logarithm and hence the identity section $1 \in \mathcal{O}_\mathbb{C}^\times$ lies in the image of $\exp$. On the overlap $V \cap W$ the two local logarithms differ by the constant $2\pi i$, so the corresponding identity sections agree in $\mathcal{O}_\mathbb{C}^\times(V \cap W)$. However, there is no holomorphic logarithm on all of $\mathbb{C}^\times$, reflecting the nontrivial monodromy around the origin; therefore the identity section on U does not belong to $\mathrm{Im}(\exp)(U)$, and the local data do not glue inside the image presheaf.

3.4 Sheafification

Given a presheaf $\mathcal{F}$, one can construct a sheaf $\mathcal{F}^+$ together with a morphism $\mathcal{F} \to \mathcal{F}^+$ that is universal among maps from $\mathcal{F}$ to a sheaf (the *sheafification*).

Definition 12 (Associated sheaf)**.** Let $\mathcal{F}$ be a presheaf on X. Define a presheaf $\mathcal{F}^+$ by

$$\mathcal{F}^+(U) = \left\{ f : U \to \bigsqcup_{x \in U} \mathcal{F}_x \;\middle|\; \begin{array}{l} f(x) \in \mathcal{F}_x \text{ for all } x \in U, \\ \text{and locally } f \text{ is represented by a section of } \mathcal{F} \end{array} \right\}.$$

More precisely, for each $x \in U$ there exist an open neighbourhood $V \subseteq U$ of x and a section $s \in \mathcal{F}(V)$ such that $f(y) = s_y$ for all $y \in V$. Then $\mathcal{F}^+$ is a sheaf, called the *sheaf associated to* $\mathcal{F}$.

Remark (Étale space viewpoint)*.* Let $S = \bigsqcup_{x \in X} \mathcal{F}_x$ and define a projection $p : S \to X$ by $p(s_x) = x$. One can topologize S so that p becomes a local homeomorphism; the resulting space is the *étalé space* of $\mathcal{F}$. Then $\mathcal{F}^+(U)$ can be identified with the set of continuous sections of p over U.

3.5 Local nature of morphisms of sheaves

Let $\mathbf{Sh}(X)$ denote the category of sheaves of abelian groups on X.

Proposition 6. *Let* $\varphi : \mathcal{F} \to \mathcal{G}$ *be a morphism of sheaves on* X.

1. For every $x \in X$ *one has* $\mathrm{Ker}(\varphi_x) = (\mathrm{Ker}\,\varphi)_x$ *and* $\mathrm{Im}(\varphi_x) = (\mathrm{Im}\,\varphi)_x$.

2. φ *is injective if and only if* φ_x *is injective for all* x.

3. φ is surjective if and only if φ_x is surjective for all x.

4. A sequence of sheaves is exact if and only if the induced sequence on stalks is exact at every point.

Proof. (1) follows from the description of stalks as direct limits and from the fact that direct limits in **Ab** commute with kernels and preserve images.

(2) If φ is injective then $\mathrm{Ker}(\varphi) = 0$, hence $(\mathrm{Ker}\,\varphi)_x = 0$ for all x, so $\mathrm{Ker}(\varphi_x) = 0$ and φ_x is injective. Conversely, assume all φ_x are injective. Let $s \in \mathcal{F}(U)$ with $\varphi(s) = 0$. Then $(\varphi(s))_x = 0$ for all $x \in U$, hence $\varphi_x(s_x) = 0$. By injectivity of φ_x, $s_x = 0$ for all $x \in U$. Thus for each x there is a neighbourhood $V_x \subseteq U$ with $s|_{V_x} = 0$. By locality for $\mathcal{F}$, $s = 0$.

(3) If φ is surjective then clearly all φ_x are surjective. Conversely, assume all φ_x are surjective. Let $t \in \mathcal{G}(U)$. For each $x \in U$ choose a germ $s_x \in \mathcal{F}_x$ with $\varphi_x(s_x) = t_x$. Pick a representative $s^{(x)} \in \mathcal{F}(V_x)$ of s_x on some neighbourhood $V_x \subseteq U$. Then $\varphi(s^{(x)})$ and $t|_{V_x}$ have the same germ at x, so after shrinking V_x we may assume $\varphi(s^{(x)}) = t|_{V_x}$. On overlaps $V_x \cap V_y$ the sections $s^{(x)}$ and $s^{(y)}$ map to the same section of $\mathcal{G}$, and by injectivity on stalks (applied to the difference) they agree on overlaps after shrinking. By gluing for $\mathcal{F}$, these local sections glue to $s \in \mathcal{F}(U)$ with $\varphi(s) = t$.

(4) is obtained by combining (1)–(3) with the fact that exactness can be tested via kernels and images. $\square$

Proposition 7. *A morphism of sheaves $\varphi : \mathcal{F} \to \mathcal{G}$ is an isomorphism if and only if each $\varphi_x : \mathcal{F}_x \to \mathcal{G}_x$ is an isomorphism.*

Proof. If φ is an isomorphism then each φ_x is an isomorphism. Conversely, if all φ_x are isomorphisms, then φ is both injective and surjective by Proposition 6(2)–(3), hence an isomorphism in **Sh**(X). $\square$

Proposition 8 (Sheaf-theoretic isomorphism theorems)**.** *Let $\varphi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves.*

1. There is a natural isomorphism of sheaves $\mathcal{F}/\mathrm{Ker}(\varphi) \cong (\mathrm{Im}\,\varphi)^+$.

2. There is a natural isomorphism of sheaves $\mathcal{G}/(\mathrm{Im}\,\varphi)^+ \cong (\mathrm{coker}\,\varphi)^+$.

In particular, after sheafification, images and cokernels behave as expected from the classical isomorphism theorems.

Proof. For each open U the classical group isomorphism theorem yields

$$\mathcal{F}(U)/\mathrm{Ker}(\varphi(U)) \;\cong\; \mathrm{Im}(\varphi(U)).$$

These isomorphisms are compatible with restrictions, hence define an isomorphism of presheaves $\mathcal{F}/\mathrm{Ker}(\varphi) \to \mathrm{Im}(\varphi)$. Since the left side is a sheaf and the right side need not be, sheafifying the right side yields (1).

For (2), note that for each U one has $\mathrm{coker}(\varphi(U)) \cong \mathcal{G}(U)/\mathrm{Im}(\varphi(U))$, compatibly with restrictions, hence after sheafification we obtain $\mathcal{G}/(\mathrm{Im}\,\varphi)^+ \cong (\mathrm{coker}\,\varphi)^+$. $\square$

References

[GAV71] A. Grothendieck, M. Artin, and J. L. Verdier, *Théorie des topos et cohomologie étale des schémas*, Springer, Berlin, 1971.

[God58] Roger Godement, *Théorie des faisceaux*, Hermann, Paris, 1958.

[Har77] Robin Hartshorne, *Algebraic geometry*, Springer-Verlag, New York, 1977, Chapter II deals specifically with Sheaves.

[KS06] Masaki Kashiwara and Pierre Schapira, *Categories and sheaves*, Springer, Berlin, 2006.

[ML98] Saunders Mac Lane, *Categories for the working mathematician*, 2nd ed., Graduate Texts in Mathematics, vol. 5, Springer-Verlag, New York, 1998.

[Tam94] Günter Tamme, *Introduction to étale cohomology*, Springer, Berlin, 1994.

[Ten75] B. R. Tennison, *Sheaf theory*, Cambridge University Press, Cambridge, 1975.