Topos theory in Logic and Algebraic Geometry

Eliot Hecky

August 2022

Introduction

This text is a report written after my master's first year internship from May to July 2022 with Olivia Caramello, whom I warmly thank.

The purpose is here to summarize most of the things I studied in topos theory, keeping it self-contained and relatively brief. We begin by studying sheaves and Grothendieck toposes, and then their link with theories in first-order logic. Almost all statements have proofs, the ones being omitted are explicitely specified.

The aim of this text is not to give an introduction to category theory, which is why the well-known definitions and results in general category theory are relegated to the end, in the appendix. The last sections give examples of classifying toposes coming from algebraic geometry. Since only the topos-theoretic part is relevant here, readers unfamiliar with basic algebraic geometry can read [15] or [14].

By default, every category is assumed locally small (see definition A.1).

Contents

1 Sheaves in topology

In areas where topology is involved, like algebraic or differential geometry, sheaves are a powerful tool to study the spaces in play. For example, schemes and manifolds are defined as sheaves - this section's aim is to study some of their properties.

Definition 1.1. Let $\mathcal{C}$ be a small category. Denote by $\mathbf{Psh}(\mathcal{C})$ the category $[\mathcal{C}^{\text{op}}, \mathbf{Set}]$ of contravariants set-valued functors on $\mathcal{C}$. An object $F : \mathcal{C}^{\text{op}} \longrightarrow \mathbf{Set}$ of $\mathbf{Psh}(\mathcal{C})$ is a *presheaf* on $\mathcal{C}$ and a natural transformation $F \longrightarrow G$ is a *morphism of presheaves.* If X is a topological space, write $\mathcal{O}(X)$ for the category of open subsets of X with inclusions as morphisms, and $\mathbf{Psh}(X)$ for the category $\mathbf{Psh}(\mathcal{O}(X))$. If F is a presheaf, $i : U \subset V$ is an inclusion of open subsets and $s \in F(V)$, write $s|_U$ for $F(i)(s)$.

Definition 1.2. Let X be a topological space. A presheaf F on X is a *sheaf* when for every open subset V covered by open subsets $(V_i)_{i\in I}$ and for every family $(s_i \in F(V_i))_{i\in I}$ such that $s_i|_{V_i\cap V_j} = s_j|_{V_i\cap V_j}$ for all i, j there is a unique $s \in F(V)$ such that $s|_{V_i} = s_i$ for all i. Write $\mathbf{Sh}(X)$ for the full subcategory of $\mathbf{Psh}(X)$ on sheaves.

It will be shown in the next section that $\mathbf{Psh}(X)$ and $\mathbf{Sh}(X)$ are toposes, which means that they come with several nice categorical properties.

Proposition 1.3. *Let $\mathcal{C}$ be a small category. Then $\mathbf{Psh}(\mathcal{C})$ has all small limits, exponentials (that is, every product functor $X \times -$ has a right adjoint $(-)^X$) and has a subobject classifier. In particular, $\mathbf{Psh}(\mathcal{C})$ is cartesian closed. The same is true for $\mathbf{Sh}(X)$ where X is a topological space, and will be shown later, in theorem 2.8.*

Proof. $\mathbf{Psh}(\mathcal{C})$ is a functor category to $\mathbf{Set}$ which has all small limits, so it has all small limits. Precisely, limits are computed pointwise : if $D : J \longrightarrow \mathbf{Psh}(\mathcal{C})$ is a diagram of small shape J, $(\lim D)(c)$ is the limit (in $\mathbf{Set}$) of the diagram $D_c : J \longrightarrow \mathbf{Set}$, sending i to $D_c(i) = D(i)(c)$ and $i \longrightarrow j$ to $D_c(i \longrightarrow j) = D(i \longrightarrow j)_c$.

Let $P, Q \in \mathbf{Psh}(\mathcal{C})$ be presheaves. Define $Q^P(c) = \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(\mathbf{y}_{\mathcal{C}}(c) \times P, Q)$. This defines a presheaf Q^P, which comes with a natural transformation $e : Q^P \times P \longrightarrow Q$, where $e_c(\theta, y) = \theta_c(1_c, y)$. It is straightforward to see that we indeed get an adjunction : to every $\phi \in \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(R \times P, Q)$, define a unique corresponding $\phi' \in \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(R, Q^P)$ as follows. For $c \in \mathcal{C}$ and $u \in Rc$, let $\phi'_c(u) : \mathbf{y}_{\mathcal{C}}(c) \times P \longrightarrow Q$ be the natural transformation with components :

$$(\phi'_c(u))_d \colon \begin{aligned} \mathrm{Hom}_{\mathcal{C}}(d, c) \times Pd &\longrightarrow Qd \\ (f, x) &\longmapsto \phi_d(R(f)(u), x). \end{aligned}$$

This definition is natural in d, so we get $\phi' \in \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(R, Q^P)$, and moreover $e_c(\phi'_c(u), y) = (\phi'_c(u))_c(1_c, y) = \phi_c(u, y)$ so that $\phi = e \circ (\phi' \times 1)$. The bijective assignment $\phi \longrightarrow \phi'$ is natural in R : $(-)^P$ is right adjoint to $- \times P$.

If $\mathbf{Psh}(\mathcal{C})$ has a subobject classifier Ω, it must classify the representables presheaves :

$$\mathrm{Sub}_{\mathbf{Psh}(\mathcal{C})}(\mathbf{y}_{\mathcal{C}}(c)) \cong \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(\mathbf{y}_{\mathcal{C}}(c), \Omega).$$

By the Yoneda lemma, the right-hand side is $\Omega(c)$. This shows that Ω must send every object c to the set of subpresheaves of $\mathrm{Hom}_{\mathcal{C}}(-, c)$. This is how one defines the functor Ω. Then let true : $1 \longrightarrow \Omega$ be the obvious $\mathrm{true}_c : \{*\} \longrightarrow \Omega(c)$ pointing at $\mathbf{y}_{\mathcal{C}}(c)$. To show that this is indeed a subobject classifier, let $Q \subset P$ be a subfunctor of a presheaf on $\mathcal{C}$. For $c \in \mathcal{C}$ and $x \in Pc$, set $\phi_c(x) = \{f \mid P(f)(x) \in Q(\mathrm{dom}(f))\}$. This gives a natural transformation $\phi : P \longrightarrow \Omega$. The diagram

$$\begin{array}{ccc} Q & \xrightarrow{\;!\;} & 1 \\ \big\downarrow & & \big\downarrow{\scriptstyle \text{true}} \\ P & \xrightarrow[\;\phi\;]{} & \Omega \end{array}$$

is then a pullback since $\phi_c(x) = \mathbf{y}_{\mathcal{C}}(c)$ if and only if $x \in Qc$. This ϕ is moreover unique : if $\theta : P \longrightarrow \Omega$ yields the same pullback diagram, then for every $f : a \longrightarrow c$ and $x \in Pc$, $P(f)(x) \in Qa$ if and only if $\theta_a(P(f)(x)) = \mathrm{true}_a$, namely $P(f)(\theta_c(x)) = \mathrm{true}_a$, that is $f \in \theta_c(x)$. We finally get $\theta_c(x) = \phi_c(x)$ for every c and x : $\theta = \phi$. □

When studying sheaves on topological spaces, there is a geometric mirror turning sheaves into spaces which can shed light on both sides. Let X be a topological space.

Definition 1.4. The category of continuous maps (bundles on X) $Y \longrightarrow X$ and commutative triangles $\mathbf{Top}/X$ is denoted $\mathrm{Bund}(X)$. The full subcategory on the étale bundles (those $Y \longrightarrow X$ which are local homeomorphisms) is denoted $\text{Étale}(X)$.

The end of this section is dedicated to construct a pair of adjoint functors $\Gamma \dashv \Lambda$ between $\mathbf{Psh}(X)$ and $\mathrm{Bund}(X)$ restricting to an equivalence $\mathbf{Sh}(X) \simeq \text{Étale}(X)$.

Definition 1.5. Let $p : Y \longrightarrow X$ be a bundle. A *cross-section* of p on an open $U \subset X$ is a (continuous) map $s : U \longrightarrow Y$ such that ps is the inclusion of U in X. Denote by Γ_p the presheaf sending U to the set $\Gamma_p U$ of all cross-sections of p on U.

Let F be a presheaf on X, $x \in X$ and U, V two open neighborhoods of x. If $s \in FU$ and $t \in FV$ are such that there exists an open $W \subset U \cap V$ containing x such that $s|_W = t|_W$, we say that s and t have the same *germ* at x. Define $\mathrm{germ}_x(s)$ as the equivalence class of sections of F under this relation, and denote $F_x = \{\mathrm{germ}_x(s)\}$ the *stalk* of F at x. Equivalently, $F_x = \varinjlim_{U \ni x} FU$. The set $\Lambda_F = \{(x, r) : x \in X, r \in F_x\} \cong \coprod_{x \in X} F_x$ comes with a projection $p : \Lambda_F \longrightarrow X$ of which each $s \in FU$ determines a cross-section $\dot{s} : U \longrightarrow \Lambda_F$ sending x to $\mathrm{germ}_x(s)$. p is made into a bundle by giving Λ_F the topology whose basis is open sets of the form $\dot{s}(U)$.

It is straightforward to see that Γ_p is always a sheaf, and that Λ_F is always étale. Both constructions are functorial, so we get two functors :

$$\Gamma : \mathrm{Bund}(X) \rightleftarrows \mathbf{Psh}(X) : \Lambda$$

Theorem 1.6. *There is an adjunction $\Lambda \dashv \Gamma$, whose unit (called* sheafification*) $\eta_P : P \longrightarrow \Gamma\Lambda P$ is an isomorphism when P is a sheaf, and whose counit $\epsilon_Y : \Lambda\Gamma Y \longrightarrow Y$ is an isomorphism when Y is étale. In particular, every sheaf is a sheaf of cross-sections, and every étale bundle is a bundle of germs.*

Proof. Define $\eta_P : P \longrightarrow \Gamma\Lambda P$ to send an element $s \in P(U)$ to the cross-section $\dot{s} : U \longrightarrow \Lambda_P$. If P is a sheaf, then $\dot{s} = \dot{t}$ (that is, $\mathrm{germ}_x(s) = \mathrm{germ}_x(t)$ for each $x \in U$) implies that for each x there is an open $V_x \subset U$ with $s|_{V_x} = t|_{V_x}$. Since the V_x's cover U, we get $s = t$. To show that η_P is surjective, let $h : U \longrightarrow \Lambda_P$ be a cross-section. Each $x \in U$ has an open neighborhood U_x such that $hx = \mathrm{germ}_x(s_x)$ for some $s_x \in P(U_x)$. By continuity of h there is an open neighborhood V_x of x in U_x such that $h(V_x) \subset \dot{s_x}(U_x)$, that is $h = \dot{s_x}$ on V_x. Since P is a sheaf there is an $s \in P(U)$ with $s|_{V_x} = s_x$ whence $h = \dot{s}$ and η_P is indeed an isomorphism.

Now define $\epsilon_Y : \Lambda\Gamma Y \longrightarrow Y$ as follows. Each point of $\Lambda\Gamma Y$ is of the form $\dot{s}x$ with $s : U \longrightarrow Y$ a cross-section and $x \in X$. Define $\epsilon_Y(\dot{s}x) = sx$, which is independent of the choice of s (two sections having the same germ implies that they agree on x). If $p : Y \longrightarrow X$ is étale, ϵ_Y has inverse θ_Y sending each $y \in Y$ to $\dot{s}(py)$ where s is a cross-section on a neighborhood of py.

To show the adjunction $\Lambda \dashv \Gamma$, it suffices to show the triangle identities :

$$\begin{array}{ccc} \Gamma & \xrightarrow{\eta\Gamma} & \Gamma\Lambda\Gamma \\ & \searrow^{1} & \downarrow{\scriptstyle \Gamma\epsilon} \\ & & \Gamma \end{array} \qquad \begin{array}{ccc} \Lambda & & \\ \downarrow{\scriptstyle \Lambda\eta} & \searrow^{1} & \\ \Lambda\Gamma\Lambda & \xrightarrow{\epsilon\Lambda} & \Lambda \end{array}$$

If Y is a bundle over X and $s \in \Gamma_Y U$, the composite on the left sends s to $\dot{s} \in \Gamma\Lambda\Gamma_Y U$ then to $s \in \Gamma_Y U$. Similarly, if P is a presheaf on X, $s \in PU$ and $x \in X$, the composite on the right sends $\mathrm{germ}_x(s)$ to $\mathrm{germ}_x(\dot{s})$ then to $\dot{s}x = \mathrm{germ}_x(s)$. $\square$

It is a fundamental fact in algebraic geometry that a continuous map $f : X \longrightarrow Y$ induces two functors between $\mathbf{Sh}(X)$ and $\mathbf{Sh}(Y)$. The first $f^* : \mathbf{Sh}(Y) \longrightarrow \mathbf{Sh}(X)$ sends a sheaf G on Y to the sheafification of $(f^*G)(U) = \varinjlim_{V \supset f(U)} G(V)$ and the second $f_* : \mathbf{Sh}(X) \longrightarrow \mathbf{Sh}(Y)$ sends a sheaf F on X to the sheaf $(f_*F)(V) = F(f^{-1}V)$. There is an adjunction $f^* \dashv f_*$ and the left adjoint f^* preserves finite limits ; this will constitute an example of a geometric morphism, see example 2.13.

2 Grothendieck toposes

2.1 Introduction

In this section we introduce the concept of Grothendieck toposes, which generalizes the construction of $\mathbf{Sh}(X)$ from $\mathbf{Psh}(\mathcal{O}(X))$, replacing $\mathcal{O}(X)$ with any small category.

Definition 2.1. Let $\mathcal{C}$ be a category and $c \in \mathcal{C}$.

- A *presieve* on c is a collection of arrows with codomain c ;
- A *sieve* on c is a presieve S on c such that for any $f \in S$ and any arrow g composable with f, we have $f \circ g \in S$. If $\mathcal{C}$ is small, this is equivalently a subpresheaf $S \rightarrowtail \mathrm{Hom}_{\mathcal{C}}(-, c)$.

A presieve S *generates* the sieve $\overline{S}$ consisting of arrows with codomain c factoring through an arrow in S.

Definition 2.2. Let $\mathcal{C}$ be a category. A *Grothendieck topology* on $\mathcal{C}$ is a function J assigning to any $c \in \mathcal{C}$ a collection of sieves on c such that :

(i) For any $c \in \mathcal{C}$, the *maximal sieve* $M_c = \{f \mid \mathrm{cod}(f) = c\}$ is in $J(c)$;
(ii) For any $f : d \longrightarrow c$ and $S \in J(c)$, the *pullback sieve* $f^*S = \{g : e \longrightarrow d \mid f \circ g \in S\}$ is in $J(d)$;
(iii) For any sieve S on $c \in \mathcal{C}$ and any $T \in J(c)$, if $f^*S \in J(\mathrm{dom}(f))$ for all $f \in T$ then $S \in J(c)$.

If J is a topology on $\mathcal{C}$ and S is a sieve of J, we say that S is (J-)covering. A *site* is a pair $(\mathcal{C}, J)$ with $\mathcal{C}$ a category and J a Grothendieck topology on $\mathcal{C}$.

Example 2.3. Let $\mathcal{C}$ be a category.

(i) The *trivial topology* on $\mathcal{C}$ is given by $J(c) = \{M_c\}$: on any given object, only the maximal sieve is covering ;
(ii) Suppose $\mathcal{C}$ satisfies the right Ore condition : every pair of arrows with common codomain fits in a commutative square. The *atomic topology* J_{at} on $\mathcal{C}$ has as covering sieves all the non-empty ones. If $\mathcal{C}$ dit not satisfy this condition, the axiom 2.2.(ii) would not be fulfilled ;
(iii) If X is a topological space, define a Grothendieck topology on $\mathcal{O}(X)$ for which covering sieves on an open $U \subset X$ are those generated by (small) families $(U_i)_{i \in I}$ which are open coverings of U.

Definition 2.4. Let $(\mathcal{C}, J)$ be a small site. Recall that a *presheaf* on $\mathcal{C}$ is an object of the functor category $\mathbf{Psh}(\mathcal{C}) = [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$. Mimicking the definition of sheaves on a topological space, a *sheaf* on $(\mathcal{C}, J)$ is a presheaf $P \in \mathbf{Psh}(\mathcal{C})$ such that the following condition is satisfied. For any $c \in \mathcal{C}$, $S \in J(c)$, and any family $(x_f \in P(\mathrm{dom}(f)))_{f \in S}$ such that for any arrow g composable with f, $P(g)(x_f) = x_{f \circ g}$, there exists a unique *amalgamation* $x \in P(c)$ such that $x_f = P(f)(x)$ for all $f \in S$. Equivalently, for each $c \in \mathcal{C}$ and $S \in J(c)$, the evident diagram

$$P(c) \longrightarrow \prod_{f \in S} P(\mathrm{dom}(f)) \rightrightarrows \prod_{f \in S, \mathrm{dom}(f) = \mathrm{cod}(g)} P(\mathrm{dom}(g))$$

is an equalizer (in **Set**).

Such a family $(x_f \in P(\mathrm{dom}(f)))_{f\in S}$ is called a *matching family*. In a weaker manner, P is a *separated presheaf* if all matching families have at most one amalgamation (that is, if the first arrow in the diagram is injective).

Define $\mathbf{Sh}(\mathcal{C}, J)$ to be the full subcategory of $\mathbf{Psh}(\mathcal{C})$ on sheaves. A *Grothendieck topos* is a category $\mathcal{E}$ which is equivalent to some $\mathbf{Sh}(\mathcal{C}, J)$ (for a small site $(\mathcal{C}, J)$) ; in this case we say that $(\mathcal{C}, J)$ is a *site of definition* of $\mathcal{E}$.

If $\mathcal{C}$ is locally small, a Grothendieck topology J is *subcanonical* if all representable presheaves $\mathrm{Hom}_{\mathcal{C}}(-, c)$ are sheaves. In this case, the Yoneda embedding factors through the inclusion $\mathbf{Sh}(\mathcal{C}, J) \hookrightarrow \mathbf{Psh}(\mathcal{C})$ into a full embedding $\mathbf{y} : \mathcal{C} \longrightarrow \mathbf{Sh}(\mathcal{C}, J)$.

Example 2.5. (i) **Set** is the topos of sheaves on a point ;

(ii) Any presheaf category is a Grothendieck topos, since $\mathbf{Psh}(\mathcal{C}) = \mathbf{Sh}(\mathcal{C}, J)$ where J is the trivial topology on $\mathcal{C}$;

(iii) If X is a topological space and J is the topology described in 2.3.(iii), then $\mathbf{Sh}(\mathcal{O}(X), J) = \mathbf{Sh}(X)$;

One reason for which Grothendieck toposes are studied is that they enjoy many nice categorical properties. Some of them are depicted in the next two theorems.

Theorem 2.6. *Let $(\mathcal{C}, J)$ be a small site. The inclusion $i : \mathbf{Sh}(\mathcal{C}, J) \hookrightarrow \mathbf{Psh}(\mathcal{C})$ has a left adjoint called* sheafification *which preserves finite limits.*

Proof. Let P be a presheaf on $\mathcal{C}$ and for $R \in J(c)$, let $\mathrm{Match}(R, P)$ be the set of matching families indexed by R. Define :

$$P^+(c) = \varinjlim_{R\in J(c)} \mathrm{Match}(R, P).$$

An element of $P^+(c)$ is then an equivalence class of matching families, with $(x_f)_{f\in R}$ and $(y_g)_{g\in S}$ being equivalent when there is a covering sieve $T \subset R \cap S$ such that $x_h = y_h$ for $h \in T$. P^+ is a presheaf, with $P^+(h)((x_f)_{f\in R}) = (x_{hg})_{g\in h^*R}$. This definition is functorial in P, and we get a canonical natural transformation $\eta : P \longrightarrow P^+$ given by $\eta_c(x) = \{P(f)(x) \colon f \in M_c\}$.
Notice that η is a monomorphism if and only if P is a separated presheaf, and it is an isomorphism if and only if P is a sheaf. We will show that P^+ is always separated, and that if P is separated then P^+ is a sheaf.

Let P be a presheaf, we will show that P^+ is separated. Let $x = (x_f)_{f\in R}, y = (y_g)_{g\in S} \in P^+(c)$ such that $P^+(h)(x) = P^+(h)(y)$ for some $Q \in J(c)$ and all $h \in Q$. This means that for all $h \in Q$, there is a covering sieve $T_h \subset h^*R \cap h^*S$ such that $x_{ht} = y_{ht}$ for all $t \in T_h$. By 2.2.(iii), $T = \{ht : h \in Q, t \in T_h\}$ is in $J(c)$ and $T \subset R \cap S$, so $x = y$.

Let P be a separated presheaf, we will show that P^+ is a sheaf. Let $(x_f \in P^+(\mathrm{dom}(f)))_{f\in R}$ be a matching family for P^+, with $R \in J(c)$. Each x_f is the equivalence class of some $(x_{f,g} \in P(\mathrm{dom}(g)))_{g\in S_f}$ matching family for P. This means that for all $f : d \longrightarrow c$ in R and $h : d' \longrightarrow d$, there is a covering sieve $T_{f,h} \subset h^*(S_f) \cap S_{fh}$ of d' such that for all $g \in T_{f,h}$, $x_{f,hg} = x_{fh,g}$. By axiom 2.2.(iii), $Q = \{fg \colon f \in R, g \in S_f\}$ is in $J(c)$, and indexes a matching family $y = (x_{f,g})_{fg\in Q}$. This definition does not depend on the factorization fg, because if $fg = f'g'$ then for any $k \in T_{f,g} \cap T_{f',g'}$, $P(k)(x_{f,g}) = x_{f,gk} = x_{fg,k} = x_{f',g'k} = P(k)(x_{f',g'})$. Since P is separated, $x_{f,g} = x_{f',g'}$. There is left to prove that y is an amalgamation of $(x_f)_{f\in R}$, namely that $(y_{fh})_{f\in h^*Q}$ and $x_f = (x_{f,g})_{g\in S_f}$ are equivalent. But $S_f \subset f^*Q$ and for any $g \in S_f$, $y_{fg} = x_{f,g}$. This proves that P^+ is a sheaf.

Now let $a = (-)^{++}$ be the sheafification functor candidate. We have to show the adjunction claimed in the theorem.

Let F be a sheaf, and P a presheaf. Any natural transformation $\phi : P \longrightarrow F$ factors uniquely through η into $\psi : P^+ \longrightarrow F$. Indeed, let $(x_f)_{f\in R}$ be a matching family for some $R \in J(c)$. Then for any $h : d \longrightarrow c$ in R, $\eta_d(x_h) = \{P(k)(x_h) \colon k \in M_d\}$ and $P(h)((x_f)_{f\in R}) = (x_{hg})_{g\in h^*R}$. Since $h \in R$, $h^*R = M_d$, and since the family is matching, we get the equality :

$$\eta_d(x_h) = P(h)((x_f)_{f\in R}).$$

If ψ exists, then $\psi((x_f)_{f\in R})$ is the unique $y \in F(c)$ with :

$$P(h)(y) = P(h)(\psi((x_f)_{f\in R})) = \psi(P(h)((x_f)_{f\in R}) = \psi(\eta_d(x_h)) = \phi(x_h).$$

Since $(\phi(x_h))_{h\in R}$ is matching for the sheaf F, this y indeed exists.

This means that $\eta'_P : P \xrightarrow{\eta_P} P^+ \xrightarrow{\eta_{P^+}} a(P) = P^{++}$ defines the unit of the desired adjunction.

Finally, notice that for a presheaf P and a covering sieve R, $\mathrm{Match}(R, P) \cong \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(R, P)$. Since $\mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(R, -)$ preserves limits and since filtered colimits commute with finite limits in **Set**, $(-)^+$ preserves finite limits, so a does as well. ☐

Definition 2.7. An *elementary topos* is a category which has finite limits, is cartesian closed and has a subobject classifier.

Theorem 2.8. *Any Grothendieck topos is an elementary topos and has small (co)limits.*

Proof. Let $(\mathcal{C}, J)$ be a (small) site. (Small, but in particular) finite limits in $\mathbf{Sh}(\mathcal{C}, J)$ are computed pointwise like in $\mathbf{Psh}(\mathcal{C})$, because the inclusion $i : \mathbf{Sh}(\mathcal{C}, J) \hookrightarrow \mathbf{Psh}(\mathcal{C})$ has a left adjoint so it preserves limits. In the same manner, since sheafification has a right adjoint it preserves colimits, so computing colimits in $\mathbf{Psh}(\mathcal{C})$ then sheafifying them is a way to compute (small) colimits in $\mathbf{Sh}(\mathcal{C}, J)$.

If exponentials exist in $\mathbf{Sh}(\mathcal{C}, J)$ then for any sheaves F and G, and any presheaf P, we have (naturally in P) :

$$\begin{aligned}
\mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(P, i(G^F)) &\cong \mathrm{Hom}_{\mathbf{Sh}(\mathcal{C},J)}(a(P), G^F) \\
&\cong \mathrm{Hom}_{\mathbf{Sh}(\mathcal{C},J)}(a(P) \times F, G) \\
&\cong \mathrm{Hom}_{\mathbf{Sh}(\mathcal{C},J)}(a(P \times i(F)), G) \\
&\cong \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(P \times i(F), i(G)) \\
&\cong \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(P, i(G)^{i(F)}).
\end{aligned}$$

This means that if they exist, exponentials in $\mathbf{Sh}(\mathcal{C}, J)$ are to be computed in $\mathbf{Psh}(\mathcal{C})$. There is left to show that if F and G are sheaves then G^F is a sheaf (and in fact, only G needs to be a sheaf). Recall from proposition 1.3 that $G^F(c) = \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(\mathbf{y}_\mathcal{C}(c) \times F, G)$. Let $\tau, \sigma \in G^F(c)$ be two such natural transformations and $S \in J(c)$ such that $G^F(f)(\tau) = G^F(f)(\sigma)$ for all $f : c' \longrightarrow c$ in S. This implies that $\tau(f, x) = \sigma(f, x)$ for all $x \in F(c')$. Let $k : c' \longrightarrow c$ be any arrow and $g \in k^*S$ (that is, $kg \in S$). Then $G^F(g)(\tau(k, x)) = \tau(kg, xg) = \sigma(kg, xg) = G^F(g)(\sigma(k, x))$. Since $k^*S \in J(c')$ and G is separated, we get $\tau = \sigma$ so G^F is separated as well. Now let $S \in J(c)$ and suppose we are given a natural transformation $\tau_f : \mathbf{y}_\mathcal{C}(d) \times F \longrightarrow G$ for every $f : d \longrightarrow c$ in S, forming a matching family. Let us construct $\tau' : \mathbf{y}_\mathcal{C} \times F \longrightarrow G^+$ such that for each $f \in S$, the diagram

$$\begin{array}{ccc}
\mathbf{y}_\mathcal{C}(d) \times F & \xrightarrow{\tau_f} & G \\
{\scriptstyle \mathbf{y}_\mathcal{C}(f)\times 1}\downarrow & & \downarrow{\scriptstyle \eta_G} \\
\mathbf{y}_\mathcal{C}(c) \times F & \xrightarrow[\tau']{} & G^+
\end{array}$$

commutes. If τ' exists, we get an amalgamation $(\eta_G)^{-1} \circ \tau'$ (η_G is an isomorphism since G is a sheaf), which is what we want. Let b be an object of $\mathcal{C}$, $k : b \longrightarrow c$ and $x \in F(b)$. Set :

$$\tau'_b(k, x) = \{\tau_{kh}(1, F(h)(x)) \colon h \in k^*S\}.$$

This family is matching for G for the $k^*S \in J(b)$, since for any suitable m we have $G(m)(\tau_{kh}(1, F(h)(x))) = \tau_{kh}(m, F(hm)(x)) = \tau_{khm}(1, F(hm)(x))$: $\tau' : \mathbf{y}_{\mathcal{C}}(c) \times F \longrightarrow G^+$ is well defined.

To show that the square indeed commutes, let $f : d \longrightarrow c$ in S ; so that f^*S is the maximal sieve M_d on d. For $k : b \longrightarrow d$ and $x \in F(b)$, we have on one side :

$$(\tau' \circ (\mathbf{y}_{\mathcal{C}}(f) \times 1))(k, x) = \tau'(fk, x) = \{\tau_{fkh}(1, F(h)(x)) \colon h \in (fk)^*S = M_b\}$$

and on the other side :

$$\eta_G \tau_f(k, x) = \eta_G(\tau_{fk}(1, x)) = \{\tau_{fkh}(1, F(h)(x)) \colon h \in M_b\}.$$

Now there is left to show that $\mathbf{Sh}(\mathcal{C}, J)$ has a subobject classifier. Say that a sieve S on an object c is *J-closed* if the arrows $f : d \longrightarrow c$ such that $f^*S \in J(d)$ are all in S. Now define the subobject classifier $(\Omega \in \mathbf{Sh}(\mathcal{C}, J), \mathrm{true} : 1_{\mathbf{Sh}(\mathcal{C},J)} \rightarrowtail \Omega)$ as :

$$\begin{aligned} \Omega(c) &= \{J\text{-closed sieves on } c\} \\ \Omega(f) &= f^*(-) \\ \mathrm{true}(*)(c) &= M_c. \end{aligned}$$

It is straightforward to prove that Ω is indeed a sheaf. To show that it is a subobject classifier, for F a sheaf and $A \subset F$ a subsheaf define the characteristic morphism $\chi_A : F \longrightarrow \Omega$ as :

$$(\chi_A)_c(x) = \{f : d \longrightarrow c \mid F(f)(x) \in A(d)\}.$$

The corresponding square

$$\begin{array}{ccc} A & \longrightarrow & 1 \\ \big\downarrow & & \big\downarrow{\scriptstyle \mathrm{true}} \\ F & \xrightarrow[\chi_A]{} & \Omega \end{array}$$

is then a pullback. Indeed it is pointwise a pullback since for all c in $\mathcal{C}$ and $x \in F(c)$, $x \in A(c)$ if and only if $(\chi_A)_c(x) = M_c$. This also shows the uniqueness of χ_A since for any $f : d \longrightarrow c$, $f \in (\chi_A)_c(x)$ if and only if $\mathrm{id}_d \in f^*(\chi_A)_c(x) = (\chi_A)_d(F(f)(x))$, if and only if $F(f)(x) \in A(d)$. ☐

Definition 2.9. Let $(\mathcal{C}, J)$ be a small site. A subcategory $\mathcal{D}$ of $\mathcal{C}$ is *J-dense* if

(i) every object $c \in \mathcal{C}$ has a covering sieve $R \in J(c)$ generated by morphisms whose domains are in $\mathcal{D}$ (or equivalently, the sieve generated by all morphisms of codomain c is J-covering) ;
(ii) for any $f : c \longrightarrow d$ in $\mathcal{C}$ with $d \in \mathcal{D}$, there is $R \in J(c)$ generated by morphisms $g : b \longrightarrow c$ with fg in $\mathcal{D}$ (or equivalently, the family of those morphisms generates a J-covering sieve).

Given any subcategory $\mathcal{D}$ of $\mathcal{C}$, define the *restriction* $J|_{\mathcal{D}}$ on $\mathcal{D}$ by letting $J|_{\mathcal{D}}(d)$ be the collection of sieves $R|_{\mathcal{D}} = R \cap \mathrm{Morphisms}(\mathcal{D})$ with $R \in J(d)$.

Lemma 2.10. *Let $(\mathcal{C}, J)$ be a small site and $\mathcal{D}$ a subcategory satisfying the second condition of the definition above. Then*

(i) a sieve S is $J|_{\mathcal{D}}$-covering if and only if the generated sieve $\overline{S}$ in $\mathcal{C}$ is J-covering ;
(ii) for any sheaf A on $(\mathcal{C}, J)$, the restriction of A to $\mathcal{D}$ is a $J|_{\mathcal{D}}$-sheaf.

In particular, if $\mathcal{D}$ is a J-dense subcategory of $\mathcal{C}$ then the restriction $\mathbf{Psh}(\mathcal{C}) \longrightarrow \mathbf{Psh}(\mathcal{D})$ *restricts to a functor* $\mathbf{Sh}(\mathcal{C}, J) \longrightarrow \mathbf{Sh}(\mathcal{D}, J|_{\mathcal{D}})$.

Proof. (i) Since $\overline{S}|_{\mathcal{D}} = \overline{S} \cap \text{Morphisms}(\mathcal{D}) = S$, the *if* part is obvious. Conversely, if $S = R|_{\mathcal{D}}$ for some $R \in J(d)$, then for any $f : c \longrightarrow d$ in R the sieve $f^*\overline{S}$ contains all morphisms $g : b \longrightarrow c$ with fg in $\mathcal{D}$ so it is J-covering, and thus $\overline{S} \in J(d)$.

(ii) Let $S \in J|_{\mathcal{D}}(d)$ and $(s_f \in A(\text{dom}(f)))_{f \in S}$ be a matching family. By the first part of the lemma, it is enough to show that this family extends to a matching family for $\overline{S}$, that is if there is a commutative diagram

$$\begin{array}{ccc} b & \xrightarrow{g_1} & c_1 \\ {\scriptstyle g_2}\downarrow & & \downarrow{\scriptstyle f_1 \in S} \\ c_2 & \xrightarrow[f_2 \in S]{} & d \end{array}$$

then $A(g_1)(s_{f_1}) = A(g_2)(s_{f_2})$. Since b can be covered by morphisms $h : a \longrightarrow b$ such that $g_1 h$ and $g_2 h$ are in $\mathcal{D}$, the images of the $A(g_i)(s_{f_i})$ under $A(h)$ are both $s_{f_1 g_1 h}$. Since A is a sheaf, they are equal.

□

Theorem 2.11 (Comparison lemma)**.** *Let $(\mathcal{C}, J)$ be a small site and $\mathcal{D}$ a dense subcategory of $\mathcal{C}$. Then the restriction $\mathbf{Sh}(\mathcal{C}, J) \longrightarrow \mathbf{Sh}(\mathcal{D}, J|_{\mathcal{D}})$ is an equivalence of categories.*

Proof. Let $B \in \mathbf{Psh}(\mathcal{D})$. For an object $c \in \mathcal{C}$, let $A(c)$ be the limit of the composite

$$(\mathcal{D}/c)^{\text{op}} \longrightarrow \mathcal{D}^{\text{op}} \xrightarrow{B} \mathbf{Set}.$$

We clearly get a presheaf $A \in \mathbf{Psh}(\mathcal{C})$, and the assignment $\tau : B \mapsto A$ is clearly a functor $\mathbf{Psh}(\mathcal{D}) \longrightarrow \mathbf{Psh}(\mathcal{C})$.

If A is a sheaf on $(\mathcal{C}, J)$, then for any matching family $(s_f \in A|_{\mathcal{D}}(\text{dom}(f)))_{f \in \mathcal{D}/c}$ the same argument as in lemma 2.10.(ii) ensures that it extends to a matching family for the sieve generated by all morphisms in $\mathcal{D}$ to c, hence there is a unique $s \in A(c)$ with $A(f)(s) = s_f$ for all f.

Conversely, if B is a sheaf on $(\mathcal{D}, J|_{\mathcal{D}})$ and we are given an element $(s_f \in B(\text{dom}(f)))_{f \in \mathcal{D}/d}$ of $\tau(B)(d)$, then for each $f : e \longrightarrow d$ the morphisms $g : e' \longrightarrow e$ for which fg is in $\mathcal{D}$ generate a $J|_{\mathcal{D}}$-covering sieve. Thus, s_f is uniquely determined by the s_{fg} for all such g, the latter being determined by s_{id_d} (since id_d is the terminal object in $\mathcal{D}/d$).

It just remains to show that if B is a sheaf on $(\mathcal{D}, J|_{\mathcal{D}})$ then $\tau(B)$ is a sheaf on $(\mathcal{C}, J)$. We prove that for each $c \in \mathcal{C}$ and $S \in J(c)$, the natural transformations $S \longrightarrow \tau(B)$ factor uniquely through $S \rightarrowtail \mathbf{y}_{\mathcal{C}}(c)$. Let $\alpha : S \longrightarrow \tau(B)$ be such a natural transformation : we have to find an element of $\tau(B)(c)$. By definition this is a matching family $(x_f \in B(d))_{f : d \longrightarrow c, d \in \mathcal{D}}$. For such f, the sieve $(f^*R)|_{\mathcal{D}}$ is $J|_{\mathcal{D}}$-covering on d and α defines a matching family of elements of B for this covering sieve. Since B is a sheaf, this gives an element x_f as required. □

As for any other structure in mathematics, let us define a notion of morphism between Grothendieck toposes. The usual notion of a functor does not carry enough structure for $\mathfrak{Topos}$ to be an interesting category :

Definition 2.12. Let $\mathcal{E}$ and $\mathcal{F}$ be Grothendieck toposes. A *geometric morphism* $f : \mathcal{E} \longrightarrow \mathcal{F}$ is a pair of adjoint functors $f^* \dashv f_*$ such that the left adjoint $f^* : \mathcal{F} \longrightarrow \mathcal{E}$ preserves finite limits. If moreover f^* has a left adjoint, then f is an *essential* geometric morphism. A *point* of a topos $\mathcal{E}$ is a geometric morphism $\mathbf{Set} \longrightarrow \mathcal{E}$. A *geometric transformation* $\eta : f \longrightarrow g$ between geometric morphisms is a natural transformation $f^* \longrightarrow g^*$. Denote by $\mathfrak{Geom}(\mathcal{E}, \mathcal{F})$ the category of geometric morphisms $\mathcal{E} \longrightarrow \mathcal{F}$ and geometric transformations between them.

Example 2.13. (i) Let $f : X \longrightarrow Y$ be a continuous map of topological spaces. Then the inverse and direct image functors introduced in the end of the first section yield a geometric morphism $f : \mathbf{Sh}(X) \longrightarrow \mathbf{Sh}(Y)$. If Y is Hausdorff (or more generally, sober) then the corresponding map $\mathrm{Hom}_{\mathbf{Top}}(X, Y) \longrightarrow \mathfrak{Geom}(\mathbf{Sh}(X), \mathbf{Sh}(Y))$ is a bijection (see [12], p.348). In particular, under this assumption, the points of the topos $\mathbf{Sh}(Y)$ are in bijection with the points of the space Y.

(ii) Let $f : \mathcal{C} \longrightarrow \mathcal{D}$ be a functor. It induces an essential geometric morphism $f : [\mathcal{C}, \mathbf{Set}] \longrightarrow [\mathcal{D}, \mathbf{Set}]$, where $f^* = - \circ f$. The left and right adjoints of f^* are given by its Kan extensions. If $\mathcal{C}$ and $\mathcal{D}$ have finite limits and they are preserved by f, then the left extension $f_! : \mathbf{Psh}(\mathcal{C}) \longrightarrow \mathbf{Psh}(\mathcal{D})$ also preserves finite limits, and is thus the left adjoint of a geometric morphism $\mathbf{Psh}(\mathcal{D}) \longrightarrow \mathbf{Psh}(\mathcal{C})$.

(iii) There are non-trivial toposes without any points. In the light of the next sections, these are the toposes arising from consistent geometric theories which have no models in **Set**. It is discussed in [2], but an example found by Deligne is given in [1], p.243 : let K be a compact space and μ a measure on K. Endow the ordered set of measurable subsets of K modulo subsets of measure zero with the Grothendieck topology given by countable open covers (modulo subsets of measure zero). If μ is not trivial, the corresponding topos is not empty, but its points are given by those $x \in K$ such that $\mu(\{x\}) \neq 0$. Take $K = [0, 1]$ with the Lebesgue measure to find a pointless topos.

2.2 The link with local operators

There is another description of Grothendieck toposes in terms of Lawvere-Tierney topologies on presheaf toposes. This equivalence shortens some proofs and definitions, and we shall give it here.

Definition 2.14. Let $\mathcal{E}$ be an elementary topos. A *local operator*, or *Lawvere-Tierney topology* on $\mathcal{E}$ is an arrow $j : \Omega \longrightarrow \Omega$ such that these three diagrams commute :

$$\begin{array}{ccc} 1 & & \\ {\scriptstyle \text{true}}\downarrow & \searrow^{\text{true}} & \\ \Omega & \xrightarrow[j]{} & \Omega \end{array} \qquad \begin{array}{ccc} \Omega & & \\ {\scriptstyle j}\downarrow & \searrow^{j} & \\ \Omega & \xrightarrow[j]{} & \Omega \end{array} \qquad \begin{array}{ccc} \Omega \times \Omega & \xrightarrow{\wedge} & \Omega \\ {\scriptstyle j \times j}\downarrow & & \downarrow{\scriptstyle j} \\ \Omega \times \Omega & \xrightarrow{\wedge} & \Omega \end{array}$$

$\wedge : \Omega \times \Omega \longrightarrow \Omega$ being the meet operation in the internal Heyting algebra structure of Ω, that is the classifying arrow of the monomorphism $(\text{true}, \text{true}) : 1 \rightarrowtail \Omega \times \Omega$ (see remark 3.10).

A *closure operator* on $\mathcal{E}$ is a family of functions $c_X : \mathrm{Sub}_{\mathcal{E}}(X) \longrightarrow \mathrm{Sub}_{\mathcal{E}}(X)$ indexed by $X \in \mathcal{E}$ such that $m \leq c(m)$ and $c(c(m)) = c(m)$ for all m. c is called *universal* if it preserves pullbacks.

Example 2.15. In Heyting algebras, the intuitionistic identities $x \leq \neg\neg x$, $\neg\neg\neg\neg x = \neg\neg x$ and $\neg\neg(x \wedge y) = \neg\neg x \wedge \neg\neg y$ always hold (see definition 3.6) : $\neg\neg$ is a universal closure operator.

Theorem 2.16. *If $\mathcal{E}$ is an elementary topos, there is a bijection between Lawvere-Tierney topologies on $\mathcal{E}$ and universal closure operators.*

Proof. Send a closure operator c to $j_c : \Omega \longrightarrow \Omega$ the classifying morphism of $c(\text{true} : 1 \longrightarrow \Omega)$. It is a Lawvere-Tierney topology : the first diagram commutes by the universal property of the pullback (with the arrows $\text{true} : 1 \longrightarrow \Omega$ and $1 \longrightarrow 1$), the second by universality of c and the fact that $c(c(\text{true})) = c(\text{true})$.

For the third one, notice that by universality of c, every vertical face in this diagram is a pullback :

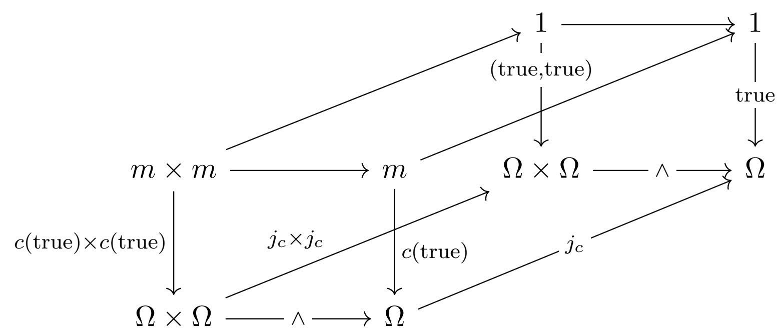

Conversely, send a local operator $j : \Omega \longrightarrow \Omega$ to c_j which takes a subobject $m \rightarrowtail b$ with classifying arrow $\chi_m : b \longrightarrow \Omega$ to the subobject m' with classifying arrow $j\chi_m$ ($c_j(m) \rightarrowtail b$ is the pullback of true along $j\chi_m$). Since $j \circ \text{true} = \text{true}$, we can use the universal property of the pullback to get a morphism $m \longrightarrow c_j(m)$. The identity $j^2 = j$ proves directly that $c_j(c_j(m)) = c_j(m)$, and the square defining j can be used to show that c_j preserves pullbacks. Finally, use the same arguments to get $c_{j_c}(m) = m$ for all m and $j = j_{c_j}$. □

Definition 2.17. Let $\mathcal{E}$ be an elementary topos and c a universal closure operator on $\mathcal{E}$. A monomorphism $m : Y' \rightarrowtail Y$ is *c-dense* if $c(m) = \mathrm{id}_Y$. An object X is a *c-sheaf* if for any c-dense monomorphism $m : Y' \rightarrowtail Y$ the arrows $Y' \longrightarrow X$ all factor uniquely through m ; that is

$$\mathrm{Hom}_{\mathcal{E}}(m, X) : \mathrm{Hom}_{\mathcal{E}}(Y, X) \longrightarrow \mathrm{Hom}_{\mathcal{E}}(Y', X)$$

is bijective. Write $\mathbf{sh}_c(\mathcal{E})$ (or $\mathbf{sh}_j(\mathcal{E})$ if j and c correspond under the equivalence of theorem 2.16) for the full subcategory of $\mathcal{E}$ on c-sheaves.

Theorem 2.18. *Let $\mathcal{C}$ be a small category. The Grothendieck topologies on $\mathcal{C}$ correspond bijectively to the Lawvere-Tierney topologies on $\mathbf{Psh}(\mathcal{C})$.*

Proof. It was shown in proposition 1.3 that the subobject classifier in $\mathbf{Psh}(\mathcal{C})$ is given by $\Omega(c)$ being the set of subpresheaves of $\mathbf{y}_{\mathcal{C}}(c)$, that is sieves on c. Given a Lawvere-Tierney topology j on $\mathbf{Psh}(\mathcal{C})$, the subobject $J \rightarrowtail \Omega$ whose classifying arrow is j verifies $S \in J(c)$ if and only if $j_c(S) = M_c$. Since $j \circ \text{true} = \text{true}$, $M_c \in J(c)$. Since j is a natural transformation, each arrow $f : c' \longrightarrow c$ induces $j_{c'}(f^*S) = f^*j_c(S)$; in particular $S \in J(c)$ implies $f^*S \in J(c')$. For the transitivity, if $S \in J(c)$ and T is a sieve on c such that $g^*T \in J(\mathrm{dom}(g))$ for all $g \in S$ then for such $g : d \longrightarrow c$ we have $g^*j_c(T) = j_d(g^*T) = M_d$. Since $\mathrm{id}_d \in g^*j_c(T)$ for all $g \in S$, $S \subset j_c(T)$ and $M_c = j_c(S) \subset j_cj_c(T)$. This shows that $j_c(T) = M_c$, that is $T \in J(c)$.

Conversely, given a Grothendieck topology J on $\mathcal{C}$ one can define $j : \Omega \longrightarrow \Omega$ by saying that $j_c(S)$ is the set of arrows $g : d \longrightarrow c$ such that $g^*S \in J(\mathrm{dom}(g))$. The naturality of j is obvious, as well as the fact that $j_c(M_c) = M_c$, which means that $j \circ \text{true} = \text{true}$. Since $S \subset T$ implies $j_c(S) \subset j_c(T)$, we get $j_c(S \cap T) = j_c(S) \cap j_c(T)$. Finally, it is clear that $S \subset j_c(S)$ so $j_c(S) \subset j_cj_c(S)$, and if $g \in j_cj_c(S)$ then $g^*j_c(S) \in J(\mathrm{dom}(g))$. Since for each $h \in j_c(S)$ one has $h^*S = M_c \in J(\mathrm{dom}(h))$, the transitivity axiom implies that $g \in j_c(S)$.

It is straightforward to check that these two assignments are inverse bijections. □

Theorem 2.19. *Let $\mathcal{C}$ be a small category. If a Lawvere-Tierney topology j on $\mathbf{Psh}(\mathcal{C})$ and a Grothendieck topology J on $\mathcal{C}$ correspond under the bijection of theorem 2.18, then the j-sheaves are exactly the J-sheaves : $\mathbf{sh}_j(\mathbf{Psh}(\mathcal{C})) = \mathbf{Sh}(\mathcal{C}, J)$.*

Proof. Let $P \in \mathbf{sh}_j(\mathbf{Psh}(\mathcal{C}))$. For each object c and $S \in J(c)$, one has

$$\mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(S, P) \cong \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(\mathbf{y}_{\mathcal{C}}(c), P) \cong P(c)$$

because by definition $S \rightarrowtail \mathbf{y}_{\mathcal{C}}(c)$ is a j-dense monomorphism. This clearly shows that P is a J-sheaf.

Conversely, if P is a J-sheaf and $A \rightarrowtail E$ is a j-dense monomorphism in $\mathbf{Psh}(\mathcal{C})$, we have to show that all natural transformations $\sigma : A \longrightarrow P$ extend uniquely to $E \longrightarrow P$. By the construction in theorem 2.18, we have for every $e \in E(c)$ that $e \in c_j(A)(c)$ if and only if $\chi_A(e) \in J(c)$. By the construction of χ_A in theorem 2.8, this is the case exactly when $\{f : d \longrightarrow c \mid E(f)(e) \in A(d)\} \in J(c)$. Since we already have a matching family $(\sigma_d(E(f)(e))$ for this covering sieve, there is an amalgamation $p \in P(c)$ which can define $\tau_c(e) = p$ as a natural extension of σ to E. □

Example 2.20. Let $\mathcal{E}$ be an elementary topos. The $\neg\neg$ closure from example 2.15 defines a subtopos $\mathbf{sh}_{\neg\neg}(\mathcal{E})$. This topos is always a boolean category, and this fact can be used to prove that the continuum hypothesis is not provable in ZFC. If A is a set strictly larger than $\mathbf{N}$, let P be the poset of maps $p : F_p \longrightarrow 2$ (with F_p a finite subset of $A \times \mathbf{N}$), where $q \leq p$ if $F_p \subset F_q$ and $q|_{F_p} = p$. Then it can be shown that in the internal logic of the *Cohen topos* $\mathbf{sh}_{\neg\neg}(\mathbf{Psh}(P))$ (which is boolean), the axiom of choice holds but not the continuum hypothesis (for the set A).

2.3 Diaconescu's equivalence

A fundamental theorem of Diaconescu gives a nice description of geometric morphisms with target a given topos $\mathbf{Sh}(\mathcal{C}, J)$ as certain functors coming out of $\mathcal{C}$. Let $\mathcal{C}$ be a small category.

Definition 2.21. For a presheaf $P \in \mathbf{Psh}(\mathcal{C})$, the *category of elements* $\int P$ has objects pairs (c, x) with $c \in \mathcal{C}$ and $x \in Pc$ and arrows $(c, x) \longrightarrow (d, y)$ for each $f : c \longrightarrow d$ such that $P(f)(y) = x$. There is a canonical projection functor $\pi_P : \int P \longrightarrow \mathcal{C}$.

Proposition 2.22. *Let $\mathcal{E}$ be a locally small and cocomplete category. Then for any functor $A : \mathcal{C} \longrightarrow \mathcal{E}$, the functor :*

$$R_A : \begin{array}{l} \mathcal{E} \longrightarrow \mathbf{Psh}(\mathcal{C}) \\ e \longmapsto \mathrm{Hom}_{\mathcal{E}}(A(-), e) \end{array}$$

has a left adjoint $- \otimes_{\mathcal{C}} A : \mathbf{Psh}(\mathcal{C}) \longrightarrow \mathcal{E}$.

Proof. The desired left adjoint sends $P \in \mathbf{Psh}(\mathcal{C})$ to $P \otimes_{\mathcal{C}} A = \mathrm{colim}(A \circ \pi_P)$. A morphism of presheaves $\tau : P \longrightarrow R_A(e)$ is a family of maps $(\tau_c : Pc \longrightarrow \mathrm{Hom}_{\mathcal{E}}(Ac, e))_{c \in \mathcal{C}}$ natural in c. We view this data as a family of arrows $(\tau_c(x) : Ac \longrightarrow e)_{(c,x) \in \int P}$, and the naturality of τ is viewed as the commutativity of

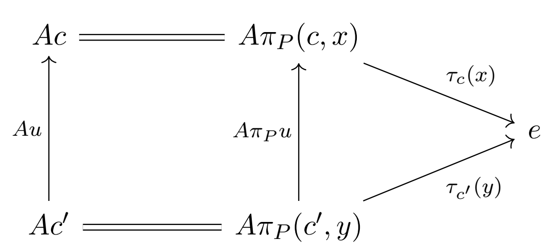

for each $u : c' \longrightarrow c$. This means that the arrows $\tau_c(x)$ form a cocone over the diagram $A\pi_P$. By the universal property of colimits, cocones with target e are in bijection with arrows $\mathrm{colim}(A\pi_P) \longrightarrow e$, which means that $\mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(P, R_A e) \cong \mathrm{Hom}_{\mathcal{E}}(\mathrm{colim}(A\pi_P), e)$. This bijection is natural in P and e, which concludes the proof. □

Definition 2.23. Let $\mathcal{E}$ be a Grothendieck topos. A functor $A : \mathcal{C} \longrightarrow \mathcal{E}$ is *flat* if $- \otimes_{\mathcal{C}} A$ preserves finite limits. Let J be a Grothendieck topology on $\mathcal{C}$. A is *J-continuous* if it sends J-covering sieves to epimorphic families. Let $\mathbf{Flat}(\mathcal{C}, \mathcal{E})$ be the full subcategory of $[\mathcal{C}, \mathcal{E}]$ on flat functors, and $\mathbf{Flat}_J(\mathcal{C}, \mathcal{E})$ be the full subcategory of $\mathbf{Flat}(\mathcal{C}, \mathcal{E})$ on J-continuous flat functors.

Theorem 2.24 (Weak Diaconescu's equivalence)**.** *Let $\mathcal{C}$ be a small category and $\mathcal{E}$ a Grothendieck topos. Then there is an equivalence of categories :*

$$\mathfrak{Geom}(\mathcal{E}, \mathbf{Psh}(\mathcal{C})) \simeq \mathbf{Flat}(\mathcal{C}, \mathcal{E})$$

sending a geometric morphism $f : \mathcal{E} \longrightarrow \mathbf{Psh}(\mathcal{C})$ to the flat functor $f^* \circ \mathbf{y}_{\mathcal{C}}$ and a flat functor $A : \mathcal{C} \longrightarrow \mathcal{E}$ to the geometric morphism $- \otimes_{\mathcal{C}} A \dashv R_A$.

Proof. Begin by noticing that by the Yoneda lemma, $R_{\mathbf{y}_{\mathcal{C}}}(E)(c) = \mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(\mathbf{y}_{\mathcal{C}}(c), E) \cong E(c)$, which means that $R_{\mathbf{y}_{\mathcal{C}}} \cong \mathrm{id}_{\mathbf{Psh}(\mathcal{C})}$ and by the uniqueness of the left adjoint, $P \cong P \otimes_{\mathcal{C}} \mathbf{y}_{\mathcal{C}}$: any presheaf is a colimit of representable presheaves. Furthermore, f^* and $- \otimes_{\mathcal{C}} (f^* \circ \mathbf{y}_{\mathcal{C}})$ agree on representables and both commute with colimits, whence they are isomorphic ; in particular $f^* \circ \mathbf{y}_{\mathcal{C}}$ is indeed flat. Proposition 2.22 shows that $- \otimes_{\mathcal{C}} A \dashv R_A$ is a geometric morphism when A is flat. Notice that the category $\int \mathbf{y}_{\mathcal{C}}(c)$ has $\mathrm{id} : c \longrightarrow c$ as a terminal object ; the colimit of $A\pi_{\mathbf{y}_{\mathcal{C}}(c)}$ is its value on this terminal object :

$$[(- \otimes_{\mathcal{C}} A) \circ \mathbf{y}_{\mathcal{C}}](c) \cong A\pi_{\mathbf{y}_{\mathcal{C}}(c)}(c, \mathrm{id}_c) = Ac,$$

whence the equivalence. □

Definition 2.25. A functor $A : \mathcal{C} \longrightarrow \mathcal{E}$ from a small category $\mathcal{C}$ into a topos $\mathcal{E}$ is *filtering* if :

(i) the family of maps $(Ac \longrightarrow 1)_{c \in \mathcal{C}}$ is epimorphic ;
(ii) for any two objects c and d, the family of maps

$$(Au, Av) : Ab \longrightarrow Ac \times Ad$$

indexed by the spans $c \xleftarrow{u} b \xrightarrow{v} d$ is epimorphic ;
(iii) for any two parallel arrows $u, v : c \longrightarrow d$, letting $e_{u,v}$ be the equalizer in $\mathcal{E}$ of Au and Av, the family of maps $Ab \longrightarrow e_{u,v}$ factoring the Aw for $w : b \longrightarrow c$ equalizing u and v ($uw = vw$) is epimorphic.

Lemma 2.26. *Let $\mathcal{C}$ be a small category with finite limits and $\mathcal{E}$ a Grothendieck topos. Then a functor $\mathcal{C} \longrightarrow \mathcal{E}$ is flat if and only if it preserves finite limits.*

Proof. For conciseness purposes, we admit the fact that a functor $A : \mathcal{C} \longrightarrow \mathcal{E}$ is flat if and only if it is filtering (see [12], theorem VII.9.1, pages 399 to 409. We only need here the *if* part, which is proven by showing that $- \otimes_{\mathcal{C}} A$ preserves the terminal object as well as pullbacks). Suppose A is flat. Since $\mathbf{y}_{\mathcal{C}}$ preserves limits, the composite $(- \otimes_{\mathcal{C}} A) \circ \mathbf{y}_{\mathcal{C}} : \mathcal{C} \longrightarrow \mathcal{E}$ preserves finite limits. It was shown in theorem 2.24 that this functor is naturally isomorphic to A, whence A preserves finite limits as well. Conversely, suppose that A preserves finite limits. We show that it is filtering (see definition 2.25). The first condition is fulfilled since $A(1) \longrightarrow 1$ is an isomorphism. For the second condition, notice that for two objects c and d, the span $c \xleftarrow{\pi_c} c \times d \xrightarrow{\pi_d} d$ gives an isomorphism $(A\pi_c, A\pi_d) : A(c \times d) \longrightarrow Ac \times Ad$. The third condition is evident as A preserves equalizers. □

Lemma 2.27. *Let $(\mathcal{C}, J)$ be a small site, $\mathcal{E}$ a Grothendieck topos and $f : \mathcal{E} \longrightarrow \mathbf{Psh}(\mathcal{C})$ a geometric morphism. The following are equivalent :*

(i) f factors through $\mathbf{Sh}(\mathcal{C}, J) \rightarrowtail \mathbf{Psh}(\mathcal{C})$;
(ii) $f^* \circ \mathbf{y}_{\mathcal{C}}$ maps J-covering sieves to colimits ;
(iii) $f^* \circ \mathbf{y}_{\mathcal{C}}$ is J-continuous.

Proof. $(i) \Rightarrow (ii)$: let S be a J-covering sieve on an object c : the monomorphism $u : S \rightarrowtail \mathbf{y}_{\mathcal{C}}(c)$ is J-dense, so $\mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(u, f_*F)$ is an isomorphism for all $F \in \mathcal{E}$ (see theorem 2.19). By the adjunction $f^* \dashv f_*$, the map $\mathrm{Hom}_{\mathcal{E}}(f^*u, F)$ is also an isomorphism for all $F \in \mathcal{E}$; by the Yoneda lemma, $f^*S \rightarrowtail f^*\mathbf{y}_{\mathcal{C}}(c)$ is then an isomorphism. Writing the sieve as a colimit of representable presheaves $S \cong \varinjlim_{d \longrightarrow c \in S} \mathbf{y}_{\mathcal{C}}(d)$ and applying the colimit-preserving functor f^*, we get an isomorphism :

$$\varinjlim_{d \longrightarrow c \in S} f^*\mathbf{y}_{\mathcal{C}}(d) \cong f^*\mathbf{y}_{\mathcal{C}}(c).$$

$(ii) \Rightarrow (iii)$: this is immediate since a colimiting cocone is obviously an epimorphic family.

$(iii) \Rightarrow (i)$: it suffices to show that f^* sends dense monomorphisms $B \subset P$ to isomorphisms. As in the first part of the proof, write P as a colimit of representables $\varinjlim_{i \in I} \mathbf{y}_{\mathcal{C}}(c_i)$. For each i, define B_i as the pullback of $B \rightarrowtail P$ along the cocone leg $\mathbf{y}_{\mathcal{C}}(c_i) \longrightarrow P$. Since pullbacks preserve colimits, $B \cong \varinjlim_{i \in I} B_i$, and $B_i \rightarrowtail \mathbf{y}_{\mathcal{C}}(c_i)$ is a dense monomorphism. This means that B_i is a J-covering sieve on c_i. For each arrow $u : d \longrightarrow c_i$ in this sieve B_i, draw the triangle

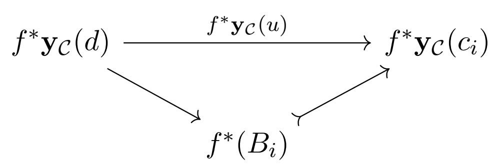

Since $f^* \circ \mathbf{y}_{\mathcal{C}}$ is J-continuous, the arrows $f^*\mathbf{y}_{\mathcal{C}}(u)$ are an epimorphic family, so the lower-right arrows are as well, and they form an isomorphism. Moreover, f^* preserves colimits so $f^*(B) \rightarrowtail f^*(P)$ is an isomorphism. $\square$

Theorem 2.28 (Strong Diaconescu's equivalence)**.** *The equivalence of theorem 2.24 restricts to an equivalence :*

$$\mathfrak{Geom}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}, J)) \simeq \mathbf{Flat}_J(\mathcal{C}, \mathcal{E}).$$

Proof. Geometric morphisms $\mathcal{E} \longrightarrow \mathbf{Sh}(\mathcal{C}, J)$ are exactly the geometric morphisms $f : \mathcal{E} \longrightarrow \mathbf{Psh}(\mathcal{C})$ which factor through $\mathbf{Sh}(\mathcal{C}, J) \rightarrowtail \mathbf{Psh}(\mathcal{C})$, and by lemma 2.27 they are exactly those for which $f^* \circ \mathbf{y}_{\mathcal{C}}$ is J-continuous. $\square$

3 Subobjects lattices in toposes

Toposes are often described as *good mathematical universes* in which one can do usual mathematics, replacing sets with the objects of the topos. Every topos has an *internal language* in which one can study first-order logic intrinsically. In general, this logic is not classical but intuitionistic : the law of excluded middle need not hold. This will appear in this section as toposes are Heyting categories, ant not always boolean. Other classes of categories are introduced, and they serve as natural settings in which one can interpret the fragments of logic which will be introduced in the next section.

Definition 3.1. Let $\mathcal{C}$ be a category with pullbacks. Since the pullback of a monomorphism is itself monomorphic, any arrow $f : c \longrightarrow d$ induces a functor on the subobject posets

$$f^* : \mathrm{Sub}_{\mathcal{C}}(d) \longrightarrow \mathrm{Sub}_{\mathcal{C}}(c),$$

sending a subobject s to the pullback $f^*s \rightarrowtail c$ of $s \rightarrowtail d$ along f.

Definition 3.2. A *cartesian category* is a category which has all finite limits ; a *cartesian functor* is a functor between cartesian categories which preserves finite limits. Write $\mathfrak{Cart}(\mathcal{C}, \mathcal{D})$ for the (full sub)category on cartesian functors between two cartesian categories $\mathcal{C}$ and $\mathcal{D}$.

Definition 3.3. Let $\mathcal{C}$ be a cartesian category. We say that $\mathcal{C}$ has *images* if for any morphism $f : c \longrightarrow d$ there is a subobject $\mathrm{Im}(f)$ of d which is the least (in $\mathrm{Sub}_{\mathcal{C}}(d)$) through which f factors. In the factorization $f : c \longrightarrow \mathrm{Im}(f) \longrightarrow d$, the arrow $c(f) : c \longrightarrow \mathrm{Im}(f)$ is called a *cover*. If $\mathcal{C}$ has images and they are stable under pullback (see definition 3.1), we say that $\mathcal{C}$ is *regular*. A *regular functor* is a cartesian functor between regular categories which preserves covers ; write their category $\mathfrak{Reg}(\mathcal{C}, \mathcal{D})$.

Proposition 3.4. *If $f : c \longrightarrow d$ is an arrow in a regular category $\mathcal{C}$, the pullback functor f^* has a left adjoint $\exists_f$.*

Proof. The left adjoint assigns to $m : b \rightarrowtail c$ the image $\mathrm{Im}(fm) \rightarrowtail d$. The adjunction is immediate : given a subobject $s \rightarrowtail d$, $\mathrm{Im}(fm) \leq s$ if and only if $m \leq f^*s$ by the universal property of the pullback. □

Definition 3.5. A *coherent category* is a regular category in which subobject posets have finite unions which are stable under pullback. Subobject posets in coherent categories have a structure of a distributive lattice, so we call them *subobject lattices*. A *coherent functor* between coherent categories is a regular functor which preserves finite unions. Write $\mathfrak{Coh}(\mathcal{C}, \mathcal{D})$ for the category of coherent functors between $\mathcal{C}$ and $\mathcal{D}$.

Definition 3.6. A *Heyting algebra* is a lattice L (a poset with finite products and finite coproducts) such that every functor $- \wedge a$ has a right adjoint $a \Rightarrow -$; namely there is an operation $\Rightarrow : L^{\mathrm{op}} \times L \longrightarrow L$ satisfying the universal property that $(x \wedge a) \leq b$ if and only if $x \leq (a \Rightarrow b)$. A *Heyting category* is a coherent category such that every pullback functor f^* has a right adjoint $\forall_f$.

Proposition 3.7. *Let $\mathcal{C}$ be a Heyting category and $b, c \rightarrowtail d$ be two subobjects. Then there exists a largest subobject $(b \Rightarrow c) \rightarrowtail d$ such that $(b \Rightarrow c) \cap b \leq c$. This defines an operation $\Rightarrow : \mathrm{Sub}_{\mathcal{C}}(d) \times \mathrm{Sub}_{\mathcal{C}}(d) \longrightarrow \mathrm{Sub}_{\mathcal{C}}(d)$ which is preserved by pullback functors. In particular, the subobject lattices in a Heyting category are Heyting algebras.*

Proof. Write $m : b \rightarrowtail d$ and let $(b \Rightarrow c) = \forall_m(b \cap c)$ (the intersection $b \cap c \rightarrowtail b$ being the pullback of $c \rightarrowtail d$ along m). The property is given by the adjunction $m^* \dashv \forall_m$. □

When $c \rightarrowtail d$ is a subobject in a Heyting category, write $\neg c = c \Rightarrow 0$. Notice that in general $c \cup \neg c$ is different from d, which explains why the law of excluded middle does not hold in the logic of Heyting categories. It however holds in boolean categories which carry classical logic :

Definition 3.8. A *boolean category* is a coherent category in which every subobject $c \rightarrowtail d$ is complemented : there exists a unique subobject $b \rightarrowtail d$ such that $b \cup c = 1$ and $b \cap c = 0$.

Definition 3.9. A category $\mathcal{C}$ is *well-powered* if for every object c, the poset $\mathrm{Sub}_{\mathcal{C}}(c)$ is essentially small. A *geometric category* is a well-powered regular category in which subobject posets have small (hence arbitrary) unions which are stable under pullbacks. A *geometric functor* between geometric categories is a regular functor which preserves unions ; write $\mathfrak{Geo}(\mathcal{C}, \mathcal{D})$ for their category.

Remark 3.10. *The relations between the different classes of categories which have been defined so far* *are summarized in this diagram :*

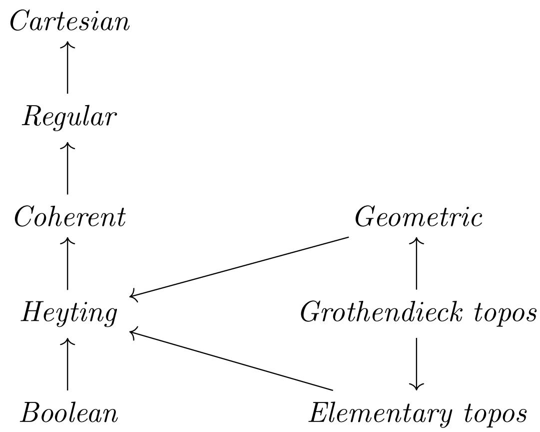

Proof. To see that a boolean category is Heyting, let $c \Rightarrow d = \neg c \cup d$. A geometric category is Heyting by the adjoint functor theorem for posets, which states that if $G : \mathcal{D} \longrightarrow \mathcal{C}$ is a functor between posets such that $\mathcal{D}$ has and G preserves all (small) unions, then G has a right adjoint.

If $\mathcal{E}$ is an elementary topos, Ω is an internal Heyting algebra, by letting :

- $0 : 1 \longrightarrow \Omega$ be the classifying arrow of $0 \rightarrowtail 1$;
- $1 : 1 \longrightarrow \Omega$ be true : $1 \rightarrowtail \Omega$;
- $\wedge : \Omega \times \Omega \longrightarrow \Omega$ be the classifying arrow of $(\text{true}, \text{true}) : 1 \rightarrowtail \Omega \times \Omega$;
- $\vee : \Omega \times \Omega \longrightarrow \Omega$ be the classifying arrow of $\pi_1^*(\text{true}) \cup \pi_2^*(\text{true})$ where π_i is the projection $\Omega \times \Omega \longrightarrow \Omega$ and $\pi_i^*(\text{true})$ is the pullback of true along it ;
- $\Rightarrow : \Omega \times \Omega \longrightarrow \Omega$ be the classifying arrow of the equalizer of $\wedge$ and π_1.

Given a subobject lattice $\text{Sub}_{\mathcal{E}}(c)$, the Heyting algebra structure on it is induced by the internal Heyting algebra structure on Ω. For example, given $a \rightarrowtail c$ and $b \rightarrowtail c$ (with classifying arrows χ_a and χ_b), construct $a \wedge b \rightarrowtail c$ as the pullback of true along the composite

$$c \xrightarrow{(\chi_a, \chi_b)} \Omega \times \Omega \xrightarrow{\wedge} \Omega.$$

□

4 Categorical first-order logic

4.1 Introduction

This section will show how Grothendieck toposes constitute an ideal setting to study first-order theories in logic. Most of the examples will come from algebraic geometry and will prove to be relevant in the next section, but of course many interesting results can be proved outside of algebraic geometry with these tools. We begin with some general first-order logic.

Definition 4.1. A *signature* Σ is a triple (S, F, R) of sets, where :

(i) S is the set of *sorts* of Σ ;

(ii) F is the set of *function symbols* of Σ, each function symbol having a *type*, consisting of a non-empty list of sorts. We write $f : A_1 \cdots A_n \longrightarrow B$ when a symbol f has type $(A_1, \dots, A_n, B)$;

(iii) R is the set of *relation symbols* of Σ, each relation symbol having a *type*, consisting of a (possibly empty) list of sorts. We write $R \rightarrowtail A_1 \cdots A_n$ when a symbol R has type $(A_1, \dots, A_n)$.

If A is a sort, we suppose that there are as many *variables* of sort A as we may wish to use. *Terms* over Σ (and their sorts) are defined to be elements of the smallest set such that the following are terms (if a term t has sort A, we write $t : A$) :

(i) $x : A$ when x is a variable of sort A ;
(ii) $f(t_1, \ldots, t_n) : B$ when $f : A_1 \cdots A_n \longrightarrow B$ is a function symbol and each $t_i : A_i$ is a term of sort A_i.

Definition 4.2. We define the set of *formulae* over a signature Σ recursively (if ϕ is a formula, we also define the set $\mathrm{FV}(\phi)$ of *free variables* of ϕ) :

(i) $R(t_1, \ldots, t_n)$ is a formula when $R \rightarrowtail A_1 \cdots A_n$ is a relation symbol and each $t_i : A_i$ is a term, and the free variables are those occurring in some t_i ;
(ii) $s = t$ is a formula when s and t are terms of the same sort, and the free variables are those occurring in s or t ;
(iii) $\top$ is a formula without any free variable ;
(iv) $\phi \wedge \psi$ is a formula when ϕ and ψ are, and $\mathrm{FV}(\phi \wedge \psi) = \mathrm{FV}(\phi) \cup \mathrm{FV}(\psi)$;
(v) $\perp$ is a formula without any free variable ;
(vi) $\phi \vee \psi$ is a formula when ϕ and ψ are, and $\mathrm{FV}(\phi \vee \psi) = \mathrm{FV}(\phi) \cup \mathrm{FV}(\psi)$;
(vii) $\phi \Rightarrow \psi$ is a formula when ϕ and ψ are, and $\mathrm{FV}(\phi \Rightarrow \psi) = \mathrm{FV}(\phi) \cup \mathrm{FV}(\psi)$;
(viii) $\neg\phi$ is a formula when ϕ is, and $\mathrm{FV}(\neg\phi) = \mathrm{FV}(\phi)$;
(ix) $(\exists x : A)\phi$ is a formula when ϕ is and x is a variable of sort A, and $\mathrm{FV}((\exists x : A)\phi) = \mathrm{FV}(\phi) \setminus \{x\}$;
(x) $(\forall x : A)\phi$ is a formula when ϕ is and x is a variable of sort A, and $\mathrm{FV}((\forall x : A)\phi) = \mathrm{FV}(\phi) \setminus \{x\}$;
(xi) $\bigvee_{i \in I} \phi_i$ is a formula when each ϕ_i is and $\bigcup_{i \in I} \mathrm{FV}(\phi_i)$ is finite, and this set is then $\mathrm{FV}(\bigvee_{i \in I} \phi_i)$;
(xii) $\bigwedge_{i \in I} \phi_i$ is a formula when each ϕ_i is and $\bigcup_{i \in I} \mathrm{FV}(\phi_i)$ is finite, and this set is then $\mathrm{FV}(\bigwedge_{i \in I} \phi_i)$.

A variable occurring in a formula without being free is called a *bound variable.*

Alas, it is too restrictive to consider the twelve types of formulae at once, whence it is convenient to restrict ourselves to certain types of formulae.

Definition 4.3. In this table we define restrictions of logic by only allowing for a given fragment (columns) certain types of formulae (lines).

	Atomic	Horn	Regular	Coherent	First-order	Geometric	∞-First-order
$R(t_1, \ldots, t_n)$	✓	✓	✓	✓	✓	✓	✓
$s = t$	✓	✓	✓	✓	✓	✓	✓
$\top$		✓	✓	✓	✓	✓	✓
$\phi \wedge \psi$		✓	✓	✓	✓	✓	✓
$\perp$				✓	✓	✓	✓
$\phi \vee \psi$				✓	✓	✓	✓
$\phi \Rightarrow \psi$					✓		✓
$\neg\phi$					✓		✓
$(\exists x : A)\phi$			✓	✓	✓	✓	✓
$(\forall x : A)\phi$					✓		✓
$\bigvee_{i \in I} \phi_i$						✓	✓
$\bigwedge_{i \in I} \phi_i$							✓

There is an eighth class of formulae called *cartesian* which will be defined later. It sits between Horn and Regular formulae, only allowing existential quantification in certain cases.

When working in categories other than **Set**, the formulae must come with some additional syntactic data :

Definition 4.4. A *context* is a (possibly empty) finite sequence $\vec{x} = (x_1, \ldots, x_n)$ of distinct variables of a given signature. If A_i is the sort of x_i, the *type* of $\vec{x}$ is the sequence $(A_1, \ldots, A_n)$. If $\vec{x}$ and $\vec{y}$ are contexts and z is a variable, $\vec{x}, z$ is the context $(x_1, \ldots, x_n, z)$ and $\vec{x}, \vec{y}$ is the context $(x_1, \ldots, x_n, y_1, \ldots, y_m)$. If all free variables of a formula ϕ appear in a context $\vec{x}$, we say that $\vec{x}$ is a *suitable* context for ϕ and that the pair $\vec{x}.\phi$ is a *formula-in-context.* The *canonical context* of a formula is the sequence of its free variables, in order of appearance.

Two formulae are *α-equivalent* if one can rename the bound variables of one to get the other. Modulo α-equivalence, we can always assume that formulae have no conflict between their bound and free variables. For example, $R(x) \vee (\forall x : A)S(x, y)$ has one bound and one free x, but is α-equivalent to $R(x) \vee (\forall z : A)S(z, y)$.

If $\vec{x}$ is a suitable context for ϕ and $\vec{s}$ is a sequence of terms of the same length and type as $\vec{x}$, let $\phi[\vec{s}/\vec{x}]$ be the (α-equivalence class) obtained by replacing each free x_i by s_i in ϕ (after some potentially necessary renaming of the bound variables of ϕ).

We are now able to define theories, in a same manner as is done usually in classical first-order logic :

Definition 4.5. A *sequent* is an expression $\phi \vdash_{\vec{x}} \psi$ where $\vec{x}$ is a context suitable for both formulae ϕ and ψ. Intuitively, ψ is meant to be a logical consequence of ϕ in the context $\vec{x}$. Some classical first-order logic lessons use $\phi \vDash \psi$, which we will not do here.

If both ϕ and ψ are atomic/Horn/regular/coherent/first-order/geometric/infinitary first-order, we say that $\phi \vdash_{\vec{x}} \psi$ is an *atomic/.../infinitary first-order sequent.*

A *theory* is a set of sequents (of course over the same signature) called the *axioms* of the theory. If all axioms of a theory $\mathbb{T}$ are atomic/.../infinitary first-order, we say that $\mathbb{T}$ is an *atomic/.../infinitary first-order theory.*

Remark 4.6. *A theory is usually defined to be a set of sequents or formulae which is closed under logical consequence, while a set of axioms generating such a theory is called an axiomatization. This will not make any difference in practice, besides, logical consequence has not been defined yet.*

Example 4.7.
(i) A theory whose signature has no sorts is a *propositional theory*, and its study is basically reduced down to propositional logic ;
(ii) A theory whose signature has no relation symbol and whose axioms are all of the form $\top \vdash_{\vec{x}} s = t$ is an *algebraic theory* (and automatically Horn) ;
(iii) A particularly relevant example is the theory of local rings. Start with the (algebraic) theory of rings, defined over the signature with one sort A, two constants 0 and 1, two binary function symbols $+$ and $\times$ and a unary function symbol $-$, with the following axioms :

$$\begin{aligned}
\top &\vdash_x 0 + x = x \\
\top &\vdash_{x,y} x + y = y + x \\
\top &\vdash_{x,y,z} (x + y) + z = x + (y + z) \\
\top &\vdash_x x + (-x) = 0 \\
\top &\vdash_x 1 \times x = x \\
\top &\vdash_{x,y} x \times y = y \times x \\
\top &\vdash_{x,y,z} (x \times y) \times z = x \times (y \times z) \\
\top &\vdash_{x,y,z} x \times (y + z) = (x \times y) + (x \times z).
\end{aligned}$$

Add to the algebraic theory of rings the coherent axioms of locality :

$$x_1 + \cdots + x_n = 1 \vdash_{x_1,\ldots,x_n} \bigvee_{i=1}^{n} \exists y(x_i \times y = 1) \qquad (n \in \mathbf{N})$$

We can notice two things here. The first is that the symbol $-$ is superfluous, because the fourth axiom can be replaced by $\top \vdash_x \exists y(x + y = 0)$. We do not wish to do this, because then the theory of rings would not be algebraic nor Horn anymore, it would only be regular. The second one is that local rings are usually defined by the property of having a unique maximal ideal ; which is not expressible in (even infinitary) first-order logic. We are saved by the fact that being a local ring is equivalent to asking that for every element x, either x or $1 - x$ is invertible. We cannot express locality in regular logic, so the theory of local rings is at best coherent. Notice that it would be enough to only keep $n = 0$ and $n = 2$, since $n = 1$ is always true and the cases $n > 2$ follow by induction.

Now is time to inject first-order theory into category theory by interpreting formulae and theories inside of nice enough categories.

Definition 4.8. Let $\mathcal{C}$ be a category with finite products and Σ a signature. A Σ-*structure* M in $\mathcal{C}$ consists of :

(i) For each sort A, an object $MA \in \mathcal{C}$ (defining for each finite sequence of sorts $(A_1, \ldots, A_n)$ the object $M(A_1, \ldots, A_n) = MA_1 \times \cdots \times MA_n$) ;
(ii) For each function symbol $f : A_1 \cdots A_n \longrightarrow B$, an arrow $Mf : M(A_1, \ldots, A_n) \longrightarrow MB$;
(iii) For each relation symbol $R \rightarrowtail A_1 \cdots A_n$, a subobject $MR \rightarrowtail M(A_1, \ldots, A_n)$.

Of course we want these structures to form a category, so we define a Σ-*structure homomorphism* $h : M \longrightarrow N$ to be a collection of arrows $h_A : MA \longrightarrow NA$ indexed by the sorts of Σ such that :

– For each function symbol $f : A_1 \cdots A_n \longrightarrow B$, the induced diagram

$$\begin{array}{ccc} M(A_1, \ldots, A_n) & \xrightarrow{Mf} & MB \\ \downarrow{\scriptstyle h_{A_1} \times \cdots \times h_{A_n}} & & \downarrow{\scriptstyle h_B} \\ N(A_1, \ldots, A_n) & \xrightarrow{Nf} & NB \end{array}$$

commutes ;
– For each relation symbol $R \rightarrowtail A_1 \cdots A_n$, there is an arrow $MR \longrightarrow NR$ such that the diagram

$$\begin{array}{ccc} MR & \rightarrowtail & M(A_1, \ldots, A_n) \\ \downarrow & & \downarrow{\scriptstyle h_{A_1} \times \cdots \times h_{A_n}} \\ NR & \rightarrowtail & N(A_1, \ldots, A_n) \end{array}$$

commutes.

Notice that Σ-Str is 2-functorial : if a functor $\mathcal{C} \longrightarrow \mathcal{D}$ preserves finite products and monomorphisms (for example, if it preserves finite limits), it induces a functor $\Sigma\text{-Str}(\mathcal{C}) \longrightarrow \Sigma\text{-Str}(\mathcal{D})$, and any natural transformation $T \longrightarrow T'$ between two such functors induces a natural transformation $\Sigma\text{-Str}(T) \longrightarrow \Sigma\text{-Str}(T')$.

We can now interpret terms and later formulae on a given signature Σ in Σ-structures :

Definition 4.9. Let $\mathcal{C}$ be a category with finite products and $M \in \Sigma\text{-Str}(\mathcal{C})$. We recursively interpret a term-in-context $\vec{x}.t$ (with $x_i : A_i$ and $t : B$) in M as a morphism $[\![\vec{x}.t]\!]_M : M(A_1, \dots, A_n) \longrightarrow MB$:

- If t is a variable x_i, let $[\![\vec{x}.t]\!]_M = \pi_i$ be the i-th product projection ;
- If $t = f(t_1 : C_1, \dots, t_m : C_m)$, let $[\![\vec{x}.t]\!]_M$ be the composition :

$$M(A_1, \dots, A_n) \xrightarrow{([\![\vec{x}.t_1]\!]_M, \dots, [\![\vec{x}.t_m]\!]_M)} M(C_1, \dots, C_m) \xrightarrow{Mf} MB.$$

If this does not introduce confusion, the ambient structure can be omitted ($[\![\vec{x}.t]\!]$) as well as the context if we wish to use the canonical one ($[\![t]\!]$).

Lemma 4.10 (Substitution lemma for terms)**.** *Let $\vec{y}.t$ be a term-in-context, with $y_i : B_i$ and $t : C$. Let $\vec{s}$ be a sequence of terms of same length and type as $\vec{y}$, and $\vec{x}$ a context suitable for all the s_i's. Then $[\![\vec{x}.t[\vec{s}/\vec{y}]]\!]_M$ is the composition :*

$$M(A_1, \dots, A_n) \xrightarrow{([\![\vec{x}.s_1]\!]_M, \dots, [\![\vec{x}.s_m]\!]_M)} M(B_1, \dots, B_m) \xrightarrow{[\![\vec{y}.t]\!]_M} MC.$$

Proof. This is true when t is a variable y_i by the universal property of products, and when $t = f(t_1 : D_1, \dots, t_s : D_s)$ by associativity of composition. ☐

Lemma 4.11 (Naturality lemma for terms)**.** *Let $h : M \longrightarrow N$ be a Σ-structure homomorphism and $\vec{x}.t$ be a term-in-context (with $x_i : A_i$ and $t : B$). Then the diagram*

$$\begin{array}{ccc} M(A_1, \dots, A_n) & \xrightarrow{[\![\vec{x}.t]\!]_M} & MB \\ \Big\downarrow{\scriptstyle h_{A_1} \times \cdots \times h_{A_n}} & & \Big\downarrow{\scriptstyle h_B} \\ N(A_1, \dots, A_n) & \xrightarrow{[\![\vec{x}.t]\!]_N} & NB \end{array}$$

commutes.

Proof. This is similar to the previous lemma : it is obviously true when t is a variable x_i and also true when $t = f(t_1 : C_1, \dots, t_m : C_m)$ by the commutativity of the first diagram defining Σ-structure homomorphisms. ☐

As in tarskian first-order logic in **Set**, terms are easy to interpret, but in general, formulae will require certain properties on the category to be interpretable :

Definition 4.12. Let $\mathcal{C}$ be a cartesian category and $M \in \Sigma\text{-Str}(\mathcal{C})$. A formula-in-context $\vec{x}.\phi$ is recursively interpreted as a subobject of $M(A_1, \dots A_n)$ (if $x_i : A_i$) :

(i) If $\phi = R(t_1 : B_1, \dots, t_m : B_m)$, $[\![\vec{x}.\phi]\!]_M$ is the pullback :

$$\begin{array}{ccc} [\![\vec{x}.\phi]\!] & \longrightarrow & MR \\ \big\downarrow & \lrcorner & \big\downarrow \\ M(A_1, \dots, A_n) & \xrightarrow{([\![\vec{x}.t_1]\!], \dots, [\![\vec{x}.t_m]\!])} & M(B_1, \dots, B_m) \end{array}$$

(ii) If $\phi = (s = t)$ where s and t are of sort B, $[\![\vec{x}.\phi]\!]_M$ is the equalizer :

$$[\![\vec{x}.\phi]\!] \rightarrowtail M(A_1, \dots, A_n) \underset{[\![\vec{x}.t]\!]}{\overset{[\![\vec{x}.s]\!]}{\rightrightarrows}} MB$$

(iii) If $\phi = \top$, $[\![\vec{x}.\phi]\!]_M$ is the top of $\mathrm{Sub}_{\mathcal{C}}(M(A_1,\dots,A_n))$;
(iv) If $\phi = (\psi \wedge \chi)$, $[\![\vec{x}.\phi]\!]_M$ is the pullback

$$\begin{array}{ccc} [\![\vec{x}.\phi]\!] & \rightarrowtail & [\![\vec{x}.\psi]\!] \\ \downarrow & \lrcorner & \downarrow \\ [\![\vec{x}.\chi]\!] & \rightarrowtail & M(A_1,\dots,A_n) \end{array}$$

(v) If $\phi = \bot$ and $\mathcal{C}$ is coherent, $[\![\vec{x}.\phi]\!]_M$ is the bottom of $\mathrm{Sub}_{\mathcal{C}}(M(A_1,\dots,A_n))$;
(vi) If $\phi = (\psi \vee \chi)$ and $\mathcal{C}$ is coherent, $[\![\vec{x}.\phi]\!]_M$ is the union of $[\![\vec{x}.\psi]\!]$ and $[\![\vec{x}.\chi]\!]$;
(vii) If $\phi = (\psi \Rightarrow \chi)$ and $\mathcal{C}$ is Heyting, $[\![\vec{x}.\phi]\!]_M$ is $[\![\vec{x}.\psi]\!] \Rightarrow [\![\vec{x}.\chi]\!]$ in $\mathrm{Sub}_{\mathcal{C}}(M(A_1,\dots,A_n))$;
(viii) If $\phi = \neg\psi$ and $\mathcal{C}$ is Heyting, $[\![\vec{x}.\phi]\!]_M$ is $\neg[\![\vec{x}.\psi]\!]$ in $\mathrm{Sub}_{\mathcal{C}}(M(A_1,\dots,A_n))$;
(ix) If $\phi = (\exists y : B)\psi$ and $\mathcal{C}$ is regular, $[\![\vec{x}.\phi]\!]_M$ is the image of the composition

$$[\![\vec{x}, y.\psi]\!] \rightarrowtail M(A_1,\dots,A_n,B) \xrightarrow{\pi} M(A_1,\dots,A_n)$$

(x) If $\phi = (\forall y : B)\psi$ and $\mathcal{C}$ is Heyting, $[\![\vec{x}.\phi]\!]_M$ is $\forall_\pi([\![\vec{x}, y.\psi]\!])$ with $\pi : M(A_1,\dots,A_n,B) \longrightarrow M(A_1,\dots,A_n)$ the projection ;
(xi) If $\phi = \bigvee_{i\in I} \psi_i$ and $\mathcal{C}$ is geometric, $[\![\vec{x}.\phi]\!]_M$ is the union of the $[\![\vec{x}.\psi_i]\!]$ in $\mathrm{Sub}_{\mathcal{C}}(M(A_1,\dots,A_n))$;
(xii) If $\phi = \bigwedge_{i\in I} \psi_i$ and $\mathrm{Sub}_{\mathcal{C}}(M(A_1,\dots,A_n))$ has arbitrary intersections, $[\![\vec{x}.\phi]\!]_M$ is the intersection of the $[\![\vec{x}.\psi_i]\!]$.

The proofs of the analogues of lemmas 4.10 and 4.11 for formulae are akin and omitted :

Lemma 4.13 (Substitution lemma for formulae)**.** *Let $\vec{y}.\phi$ be a formula-in-context interpretable in $\mathcal{C}$ (with $y_i : B_i$) and $\vec{s}$ be a sequence of terms of same length and type as $\vec{y}$. Let $\vec{x}$ be a context suitable for all the s_i's (with $x_i : A_i$). Then for any $M \in \Sigma\text{-Str}(\mathcal{C})$, there is an arrow $[\![\vec{x}.\phi[\vec{s}/\vec{y}]]\!]_M \longrightarrow [\![\vec{y}.\phi]\!]_M$ such that*

$$\begin{array}{ccc} [\![\vec{x}.\phi[\vec{s}/\vec{y}]]\!] & \dashrightarrow & [\![\vec{y}.\phi]\!] \\ \downarrow & & \downarrow \\ M(A_1,\dots,A_n) & \xrightarrow{([\![\vec{x}.s_1]\!],\dots,[\![\vec{x}.s_m]\!])} & M(B_1,\dots,B_m) \end{array}$$

is a pullback.

Lemma 4.14 (Naturality lemma for formulae)**.** *Let $\vec{x}.\phi$ be a geometric formula-in-context interpretable in $\mathcal{C}$ (with $x_i : A_i$) and $h : M \longrightarrow N$ be a Σ-structure homomorphism. Then there is an arrow $[\![\vec{x}.\phi]\!]_M \longrightarrow [\![\vec{x}.\phi]\!]_N$ such that*

$$\begin{array}{ccc} [\![\vec{x}.\phi]\!]_M & \rightarrowtail & M(A_1,\dots,A_n) \\ \downarrow & & \downarrow{\scriptstyle h_{A_1}\times\cdots\times h_{A_n}} \\ [\![\vec{x}.\phi]\!]_N & \rightarrowtail & N(A_1,\dots,A_n) \end{array}$$

commutes.

Definition 4.15. Let $M \in \Sigma\text{-Str}(\mathcal{C})$. A sequent $\sigma = (\phi \vdash_{\vec{x}} \psi)$ interpretable in $\mathcal{C}$ is *satisfied* in M when $[\![\vec{x}.\phi]\!]_M \leq [\![\vec{x}.\psi]\!]_M$ in $\mathrm{Sub}_{\mathcal{C}}(M(A_1,\dots,A_n))$, namely when the interpretation of $\vec{x}.\phi$ is a subobject of the interpretation of $\vec{x}.\psi$. In this case, we write $M \vDash \sigma$. If all the axioms of a theory $\mathbb{T}$ are satisfied in M, we say that M is a $\mathbb{T}$-*model* and write $M \vDash \mathbb{T}$.

Let $\mathbb{T}\text{-Mod}(\mathcal{C})$ be the full subcategory of $\Sigma\text{-Str}(\mathcal{C})$ on $\mathbb{T}$-models.

Lemma 4.16. *Let $T : \mathcal{C} \longrightarrow \mathcal{D}$ be a cartesian/regular/coherent/Heyting/geometric functor between such categories. For every $M \in \Sigma\text{-Str}(\mathcal{C})$ and every sequent σ interpretable in $\mathcal{C}$ such that $M \vDash \sigma$, $\Sigma\text{-Str}(T)(M) \vDash \sigma$. The converse is true if T is conservative. In particular, if $\mathbb{T}$ is a regular/.../geometric theory and T is a functor of the same kind, $\Sigma\text{-Str}(T)$ restricts to a functor $\mathbb{T}\text{-Mod}(T) : \mathbb{T}\text{-Mod}(\mathcal{C}) \longrightarrow \mathbb{T}\text{-Mod}(\mathcal{D})$.*

Proof. T preserves the interpretations of formulae because it is of the appropriate kind. For example, if T is coherent and $\phi = (\psi \vee \chi)$, $T([\![\phi]\!]_M) = T([\![\psi]\!]_M \cup [\![\chi]\!]_M) = T([\![\psi]\!]_M) \cup T([\![\chi]\!]_M)$ since T preserves finite unions ; whence $T([\![\phi]\!]_M) = [\![\phi]\!]_{\Sigma\text{-Str}(T)(M)}$. T is always at least cartesian, so it always preserves order in subobject lattices.

Notice that $M \vDash (\phi \vdash_{\vec{x}} \psi)$ if and only if $[\![\vec{x}.\phi \wedge \psi]\!]_M \rightarrowtail [\![\vec{x}.\phi]\!]_M$ is an isomorphism. If $\Sigma\text{-Str}(T)(M) \vDash (\phi \vdash_{\vec{x}} \psi)$, then $T([\![\vec{x}.\phi \wedge \psi]\!]_M) = [\![\vec{x}.\phi \wedge \psi]\!]_{\Sigma\text{-Str}(T)(M)} \rightarrowtail [\![\vec{x}.\phi]\!]_{\Sigma\text{-Str}(T)(M)} = T([\![\vec{x}.\phi]\!]_M)$ is an isomorphism. If T is conservative, this in turn implies that $M \vDash (\phi \vdash_{\vec{x}} \psi)$. □

4.2 A soundness theorem

The following definition introduces a *deduction system* providing rules for inferring the validity of sequents from other sequents. The section ends with a soundness theorem, stating that these rules can only prove sequents that are satisfied in all models of a given theory.

Definition 4.17. Rules are written in the form $\frac{\Gamma}{\sigma}$, Γ being a (possibly empty) list of sequents and σ the sequent that can be inferred from the validity of the sequents in Γ. Obviously, the rules involving certain symbols only concern the fragments of logic which include this symbol.

– The structural rules include the identity axiom

$$\frac{}{\phi \vdash_{\vec{x}} \phi}$$

the substitution rule

$$\frac{\phi \vdash_{\vec{x}} \psi}{\phi[\vec{s}/\vec{x}] \vdash_{\vec{y}} \psi[\vec{s}/\vec{x}]}$$

where $\vec{y}$ includes all the variables of $\vec{s}$, and the cut rule

$$\frac{(\phi \vdash_{\vec{x}} \psi) \ (\psi \vdash_{\vec{x}} \chi)}{\phi \vdash_{\vec{x}} \chi}$$

– The equality rules include the equality axiom

$$\frac{}{\top \vdash_x x = x}$$

and the replacement axiom

$$\frac{}{(\vec{x} = \vec{y}) \wedge \phi \vdash_{\vec{z}} \phi[\vec{y}/\vec{x}]}$$

where $\vec{z}$ contains the variables of $\vec{x}$ and $\vec{y}$ and the free variables of ϕ ;

– The finite conjunction rules include the axioms

$$\frac{}{\phi \vdash_{\vec{x}} \top} \quad \frac{}{\phi \wedge \psi \vdash_{\vec{x}} \phi} \quad \frac{}{\phi \wedge \psi \vdash_{\vec{x}} \psi}$$

and the rule

$$\frac{(\phi \vdash_{\vec{x}} \psi) \ (\phi \vdash_{\vec{x}} \chi)}{\phi \vdash_{\vec{x}} \psi \wedge \chi}$$

– The finite disjunction rules are the analogous

$$\frac{}{\bot \vdash_{\vec{x}} \phi} \quad \frac{}{\phi \vdash_{\vec{x}} \phi \vee \psi} \quad \frac{}{\psi \vdash_{\vec{x}} \phi \vee \psi}$$

and

$$\frac{(\phi \vdash_{\vec{x}} \chi) \ (\psi \vdash_{\vec{x}} \chi)}{\phi \vee \psi \vdash_{\vec{x}} \chi}$$

– The implication rules are the double rule

$$\frac{\phi \wedge \psi \vdash_{\vec{x}} \chi}{\psi \vdash_{\vec{x}} \phi \Rightarrow \chi}$$

yielding a rule for negation by setting $\chi = \bot$ and $\neg\phi = (\phi \Rightarrow \bot)$;

– The existential quantification rules are the double rule

$$\frac{\phi \vdash_{\vec{x},y} \psi}{\exists y \phi \vdash_{\vec{x}} \psi}$$

(under the assumption that the sequents are well-formed, namely y is not free in ψ) ;

– The universal quantification rules are the double rule

$$\frac{\phi \vdash_{\vec{x},y} \psi}{\phi \vdash_{\vec{x}} \forall y \psi}$$

– The infinitary conjunction and disjunction rules are the infinitary analogues of the finite conjunction and disjunction rules ;

– The two remaining mixed axioms for coherent logic are the distributive axiom

$$\frac{}{\phi \wedge (\psi \vee \chi) \vdash_{\vec{x}} (\phi \wedge \psi) \vee (\phi \wedge \chi)}$$

(and its infinitary analogue for geometric logic) and the Frobenius axiom

$$\frac{}{\phi \wedge \exists y \psi \vdash_{\vec{x}} \exists y (\phi \wedge \psi)}$$

where y is not in $\vec{x}$.

A sequent σ is *provable* in a theory $\mathbb{T}$ (or $\mathbb{T}$-*provable*) if there is a derivation from the axioms of $\mathbb{T}$ to σ in the appropriate fragment of logic.

Theorem 4.18 (Soundness theorem)**.** *Let $\mathbb{T}$ be a Horn (resp. regular, coherent, first-order, geometric) theory and M be a model of $\mathbb{T}$ in a cartesian (resp. regular, coherent, Heyting, geometric) category. If σ is a sequent which is provable in $\mathbb{T}$ then $M \vDash \sigma$.*

Proof. By induction on the height of the derivation tree leading to σ, it is enough to show that each rule in definition 4.17 is sound, that is if M satisfies the sequents above the line then it satisfies the sequent below it. This is in most cases trivial. For example, the rule

$$\frac{(\phi \vdash_{\vec{x}} \psi) \ (\phi \vdash_{\vec{x}} \chi)}{\phi \vdash_{\vec{x}} \psi \wedge \chi}$$

is sound because if M satisfies the two sequents above the line, then $[\![\vec{x}.\phi]\!]_M \leq [\![\vec{x}.\psi]\!]_M$ and $[\![\vec{x}.\phi]\!]_M \leq [\![\vec{x}.\chi]\!]_M$ in $\mathrm{Sub}_{\mathcal{C}}(M(A_1, \ldots, A_n))$ and of course the universal property of the pullback yields $[\![\vec{x}.\phi]\!]_M \leq [\![\vec{x}.\psi]\!]_M \cap [\![\vec{x}.\chi]\!]_M = [\![\vec{x}.(\psi \wedge \chi)]\!]_M$. ☐

Remark 4.19. *Notice that the law of excluded middle $\top \vdash_{\vec{x}} \phi \vee \neg\phi$ does not appear in definition 4.17, since adding it would make the soundness theorem false. Indeed, as said in the previous section, Grothendieck toposes are Heyting categories and not boolean categories, which means that subobjects are not always complemented. This is why one has to prove theorems intuitionistically in order for them to hold in every topos.*

5 Syntactic sites and classifying toposes

5.1 Syntactic categories and a completeness theorem

According to the soundness theorem, any sequent that is provable in a given theory $\mathbb{T}$ is true in $\mathbb{T}$. A converse of this, known as a completeness theorem, shall be proven using syntactic categories : any sequent which is true in $\mathbb{T}$ is indeed provable.

Definition 5.1. Let $\mathbb{T}$ be a cartesian (resp. regular, coherent, first-order, geometric) theory over Σ. The *syntactic category* $\mathcal{C}_{\mathbb{T}}^{\text{cart}}$ (resp. $\mathcal{C}_{\mathbb{T}}^{\text{reg}}$, $\mathcal{C}_{\mathbb{T}}^{\text{coh}}$, $\mathcal{C}_{\mathbb{T}}^{\text{fo}}$, $\mathcal{C}_{\mathbb{T}}$) has as objects the α-equivalence classes (see definition 4.4) of cartesian (resp. regular, ...) formulae-in-context $\{\vec{x}.\phi\}$ and as arrows $\{\vec{x}.\phi\} \longrightarrow \{\vec{y}.\psi\}$ (where $\vec{x}$ and $\vec{y}$ are chosen disjoint) the $\mathbb{T}$-provable equivalence classes $[\theta]$ of cartesian (resp. regular, ...) formulae $\theta(\vec{x},\vec{y})$ which are $\mathbb{T}$-provably functional, that is such that the sequents

$$\phi \vdash_{\vec{x}} \exists \vec{y}\theta$$
$$\theta \vdash_{\vec{x},\vec{y}} \phi \wedge \psi$$
$$\theta \wedge \theta[\vec{z}/\vec{y}] \vdash_{\vec{x},\vec{y},\vec{z}} \vec{y} = \vec{z}$$

are $\mathbb{T}$-provable. The composite $[\gamma] \circ [\theta]$ of composable arrows $[\theta] : \{\vec{x}.\phi\} \longrightarrow \{\vec{y}.\psi\}$ and $[\gamma] : \{\vec{y}.\psi\} \longrightarrow \{\vec{z}.\chi\}$ is the class $[\exists \vec{y}(\theta \wedge \gamma)]$ and the identity arrow of $\{\vec{x}.\phi\}$ is the class $[\phi \wedge \vec{x}' = \vec{x}] : \{\vec{x}.\phi\} \longrightarrow \{\vec{x}'.\phi[\vec{x}'/\vec{x}]\}$.

Lemma 5.2. *Let $\mathbb{T}$ be a cartesian (resp. regular, ...) theory.*

(i) A morphism $[\theta] : \{\vec{x}.\phi\} \longrightarrow \{\vec{y}.\psi\}$ is an isomorphism if and only if θ is $\mathbb{T}$-provably functional from $\{\vec{y}.\psi\}$ to $\{\vec{x}.\phi\}$.

(ii) $[\theta]$ is a monomorphism if and only if the sequent

$$\theta \wedge \theta[\vec{x}'/\vec{x}] \vdash \vec{x} = \vec{x}'$$

is provable in $\mathbb{T}$.

(iii) Any subobject of $\{\vec{x}.\phi\}$ is (isomorphic to one) of the form

$$\{\vec{x}'.\psi[\vec{x}'/\vec{x}]\} \xrightarrow{[\psi \wedge \vec{x}' = \vec{x}]} \{\vec{x}.\phi\},$$

with $\vec{x}.\psi$ a formula such that $\psi \vdash_{\vec{x}} \phi$ is provable in $\mathbb{T}$. For two such subobjects $[\psi]$ and $[\chi]$, we have $[\psi] \leq [\chi]$ (in $\text{Sub}_{\mathcal{C}_{\mathbb{T}}}(\{\vec{x}.\phi\})$) if and only if $\psi \vdash_{\vec{x}} \chi$ is provable in $\mathbb{T}$.

Proof. (i) If $[\theta]$ is $\mathbb{T}$-provably functional in the other direction, then $[\theta] : \{\vec{y}.\psi\} \longrightarrow \{\vec{x}.\phi\}$ is clearly the required inverse. Conversely, if $[\gamma]$ is the inverse of $[\theta]$ then γ is $\mathbb{T}$-provably equivalent to θ so $[\gamma] = [\theta]$.

(ii) Form the kernel pair $\{\vec{x}, \vec{x}'.\theta \wedge \theta[\vec{x}'/\vec{x}]\}$ of $[\theta]$ and notice that the sequent is $\mathbb{T}$-provable if and only if the diagonal map $\{\vec{x}.\phi\} \longrightarrow \{\vec{x}, \vec{x}'.\theta \wedge \theta[\vec{x}'/\vec{x}]\}$ is an isomorphism, which is the case if and only if $[\theta]$ is a monomorphism.

(iii) Such a morphism is always monic by (ii). If $[\theta] : \{\vec{y}.\psi\} \longrightarrow \{\vec{x}.\phi\}$ is a monomorphism, then again (ii) ensures that $\vec{x}.\exists \vec{y}\theta$ is cartesian relative to $\mathbb{T}$, and (i) shows that $[\theta]$ is an isomorphism over $\{\vec{x}.\phi\}$ from $\{\vec{y}.\psi\}$ to $\{\vec{x}.\exists \vec{y}\theta\}$ (and indeed, $\phi \vdash \exists \vec{y}\theta$ is provable in $\mathbb{T}$ by functionality). Moreover, if $[\psi]$ and $[\chi]$ are two such subobjects, then the only possible morphism over $\{\vec{x}.\phi\}$ from $[\psi]$ to $[\chi]$ is $[\psi \wedge \vec{x} = \vec{x}'] : \{\vec{x}.\psi[\vec{x}'/\vec{x}]\} \longrightarrow \{\vec{x}.\chi\}$, which is a morphism exactly when $\psi \vdash \chi$ is provable in $\mathbb{T}$.

□

This identification between subobjects in a syntactic category and provability of sequents enables us to directly prove :

Theorem 5.3. *(i) If $\mathbb{T}$ is a cartesian theory, then $\mathcal{C}_{\mathbb{T}}^{\text{cart}}$ is a cartesian category ;*

(ii) *If $\mathbb{T}$ is a regular theory, then $\mathcal{C}_{\mathbb{T}}^{\text{reg}}$ is a regular category ;*
(iii) *If $\mathbb{T}$ is a coherent theory, then $\mathcal{C}_{\mathbb{T}}^{\text{coh}}$ is a coherent category ;*
(iv) *If $\mathbb{T}$ is a first-order theory, then $\mathcal{C}_{\mathbb{T}}^{\text{fo}}$ is a Heyting category ;*
(v) *If $\mathbb{T}$ is a geometric theory, then $\mathcal{C}_{\mathbb{T}}$ is a geometric category.*

Now let us prove the completeness theorem which comes from the rich structure of syntactic categories.

Definition 5.4. Let $\mathbb{T}$ be a cartesian (resp. regular, coherent, first-order, geometric) theory over Σ. The *universal model* $M_{\mathbb{T}}$ is the Σ-structure in $\mathcal{C}_{\mathbb{T}}$ where :

– for each sort A, $M_{\mathbb{T}}A$ is the object $\{x : A.\top\}$;
– for each function symbol $f : A_1 \cdots A_n \longrightarrow B$, $M_{\mathbb{T}}f$ is the morphism

$$\{x_1 : A_1, \ldots, x_n : A_n.\top\} \xrightarrow{[f(x_1,\ldots,x_n)=y]} \{y : B.\top\} ;$$

– for each relation symbol $R \rightarrowtail A_1 \cdots A_n$, $M_{\mathbb{T}}R$ is the subobject

$$\{x_1 : A_1, \ldots, x_n : A_n.R(x_1, \ldots, x_n)\} \xrightarrow{[R(x_1,\ldots,x_n)]} \{x_1 : A_1, \ldots, x_n : A_n.\top\}.$$

Theorem 5.5. *Let $\mathbb{T}$ be a cartesian (resp. regular, ...) theory over Σ. Then for any cartesian (resp. regular, ...) formula-in-context $\vec{x}.\phi$, its interpretation in $M_{\mathbb{T}}$ is the subobject $[\![\vec{x}.\phi]\!]_{M_{\mathbb{T}}} = [\phi] : \{\vec{x}.\phi\} \rightarrowtail \{\vec{x}.\top\}$. Moreover, a cartesian (resp. regular, ...) sequent over Σ is satisfied in $M_{\mathbb{T}}$ if and only if it is provable in $\mathbb{T}$.*

Proof. The first part is an easy induction on the shape of ϕ. If σ is a sequent which is provable in $\mathbb{T}$ then by soundness it is satisfied in $M_{\mathbb{T}}$, and if $\phi \vdash \psi$ is satisfied in $M_{\mathbb{T}}$ then $[\phi] \leq [\psi]$ and by lemma 5.2.(iii), the sequent is provable in $\mathbb{T}$. □

Corollary 5.6 (Completeness theorem)**.** *Let $\mathbb{T}$ be a cartesian (resp. regular, ...) theory. Then any sequent which is satisfied in all models of $\mathbb{T}$ in cartesian (resp. regular, ...) categories is provable in $\mathbb{T}$.*

5.2 Classifying geometric theories

Another strong result that can be proven with syntactic categories will be theorem 5.12, for which we shall prove a lemma here.

Lemma 5.7. *Let $\mathbb{T}$ be a cartesian (resp. regular, coherent, geometric) theory. Then, for any cartesian (resp. regular, coherent, geometric) category $\mathcal{D}$, the functor from $\mathfrak{Cart}(\mathcal{C}_{\mathbb{T}}^{\text{cart}}, \mathcal{D})$ (resp. $\mathfrak{Reg}(\mathcal{C}_{\mathbb{T}}^{\text{reg}}, \mathcal{D})$, $\mathfrak{Coh}(\mathcal{C}_{\mathbb{T}}^{\text{coh}}, \mathcal{D})$, $\mathfrak{Geo}(\mathcal{C}_{\mathbb{T}}, \mathcal{D})$) to $\mathbb{T}$-Mod$(\mathcal{D})$ which sends F to $F(M_{\mathbb{T}})$ is an equivalence of categories.*

Proof. By lemma 4.16, this functor is well-defined. Given a $\mathbb{T}$-model M in $\mathcal{D}$, define $F_M : \mathcal{C}_{\mathbb{T}} \longrightarrow \mathcal{D}$ by letting $F_M(\{\vec{x}.\phi\}) = [\![\vec{x}.\phi]\!]_M$ and $F_M([\theta])$ be the morphism whose graph is $[\![\vec{x}, \vec{y}.\theta]\!]_M$. This defines a functor which is cartesian (resp. regular, coherent, geometric). The functoriality of the assignment $M \mapsto F_M$ is a consequence of lemma 4.14. Clearly, $F_M(M_{\mathbb{T}}) \cong M$ naturally in M. Conversely, if we are given a cartesian (resp. regular, coherent, geometric) functor $F : \mathcal{C}_{\mathbb{T}} \longrightarrow \mathcal{D}$ with $F(M_{\mathbb{T}}) \cong M$, then for every $\vec{x}.\phi$ there is an isomorphism $F(\{\vec{x}.\phi\}) \cong [\![\vec{x}.\phi]\!]_M$, yielding a natural isomorphism $F \cong F_M$, naturally in F. □

Let us also define Grothendieck topologies on some of the syntactic categories, since their corresponding toposes will be useful later.

Definition 5.8. In a regular category, a *covering family* is a family of arrows with common codomain c whose union of images is the maximal subobject of c (that is, id_c).

- Let $\mathcal{C}$ be a regular category. The *regular topology* $J_{\mathcal{C}}^{\text{reg}}$ is the topology whose covering sieves are those which contain a cover.
- Let $\mathcal{C}$ be a coherent category. The *coherent topology* $J_{\mathcal{C}}^{\text{coh}}$ is the topology whose covering sieves are those which contain a finite covering family.
- Let $\mathcal{C}$ be a geometric category. The *geometric topology* $J_{\mathcal{C}}$ is the topology whose covering sieves are those which contain a small covering family.

One can notice that all these topologies are subcanonical (see definition 2.4).

Proposition 5.9. *Let $\mathcal{C}$ and $\mathcal{D}$ be regular (resp. coherent, geometric) categories. A cartesian functor is regular (resp. coherent, geometric) if and only if it sends $J_{\mathcal{C}}^{\text{reg}}$- (resp. $J_{\mathcal{C}}^{\text{coh}}$-, $J_{\mathcal{C}}$-) covering sieves to covering families.*

A useful tool in algebraic topology is the notion of classifying spaces. It turns out that a similar construction is possible for (geometric) theories, and it is of great value when studying said theories.

Definition 5.10. Let $\mathbb{T}$ be a geometric theory. A *classifying topos* for $\mathbb{T}$ is a Grothendieck topos $\mathbf{Set}[\mathbb{T}]$ such that for every Grothendieck topos $\mathcal{E}$ we have an equivalence :

$$\mathfrak{Geom}(\mathcal{E}, \mathbf{Set}[\mathbb{T}]) \simeq \mathbb{T}\text{-Mod}(\mathcal{E})$$

natural in $\mathcal{E}$ in the sense that for every geometric morphism $f : \mathcal{F} \longrightarrow \mathcal{E}$, the induced diagram

$$\begin{array}{ccc} \mathfrak{Geom}(\mathcal{E}, \mathbf{Set}[\mathbb{T}]) & \xrightarrow{\ \simeq\ } & \mathbb{T}\text{-Mod}(\mathcal{E}) \\ \Big\downarrow {\scriptstyle -\circ f} & & \Big\downarrow {\scriptstyle f^*} \\ \mathfrak{Geom}(\mathcal{F}, \mathbf{Set}[\mathbb{T}]) & \xrightarrow{\ \simeq\ } & \mathbb{T}\text{-Mod}(\mathcal{F}) \end{array}$$

commutes up to isomorphism. Clearly this implies the uniqueness, up to (canonical) isomorphism, of a classifying topos for a given theory.

Remark 5.11. *In the diagram of definition 5.10, set $\mathcal{E} = \mathbf{Set}[\mathbb{T}]$ to see that there is a universal $\mathbb{T}$-model U in $\mathbf{Set}[\mathbb{T}]$ (the image of $\mathrm{id}_{\mathbf{Set}[\mathbb{T}]}$ under the equivalence) such that every $\mathbb{T}$-model in a Grothendieck topos $\mathcal{F}$ appears as the pullback of U along a (unique up to isomorphism) geometric morphism $f : \mathcal{F} \longrightarrow \mathbf{Set}[\mathbb{T}]$ (the image in $\mathfrak{Geom}(\mathcal{F}, \mathbf{Set}[\mathbb{T}])$ of the model under the equivalence).*

Several examples will be discussed later. One could legitimately ask why this definition is restricted to geometric theories. Nothing keeps us from defining classifying toposes of non-geometric theories, but the two following theorems give a convincing answer.

Theorem 5.12. *(i) For any cartesian theory $\mathbb{T}$, $\mathbf{Psh}(\mathcal{C}_{\mathbb{T}}^{\text{cart}})$ classifies $\mathbb{T}$.*
(ii) For any regular theory $\mathbb{T}$, $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}^{\text{reg}}, J_{\mathcal{C}_{\mathbb{T}}}^{\text{reg}})$ classifies $\mathbb{T}$.
(iii) For any coherent theory $\mathbb{T}$, $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}^{\text{coh}}, J_{\mathcal{C}_{\mathbb{T}}}^{\text{coh}})$ classifies $\mathbb{T}$.
(iv) For any geometric theory $\mathbb{T}$, $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J_{\mathcal{C}_{\mathbb{T}}})$ classifies $\mathbb{T}$.

Proof. Let $\mathcal{E}$ be a Grothendieck topos and $\mathbb{T}$ a geometric theory. By Diaconescu's equivalence, we have $\mathfrak{Geom}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J_{\mathbb{T}})) \simeq \mathbf{Flat}_{J_{\mathbb{T}}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E})$; and by lemma 2.26 this category is in turn equivalent to cartesian $J_{\mathbb{T}}$-continuous functors $\mathcal{C}_{\mathbb{T}} \longrightarrow \mathcal{E}$. Proposition 5.9 shows that this is equivalent to $\mathfrak{Geo}(\mathcal{C}_{\mathbb{T}}, \mathcal{E})$, and lemma 5.7 concludes : $\mathfrak{Geom}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J_{\mathbb{T}})) \simeq \mathbb{T}\text{-Mod}(\mathcal{E})$. All these equivalences are natural in $\mathcal{E}$, and the same proof works when $\mathbb{T}$ is cartesian, regular or coherent. □

Theorem 5.13. *Every Grothendieck topos classifies a geometric theory.*

Proof. Let $(\mathcal{C}, J)$ be a small site and $\mathcal{E}$ any Grothendieck topos. By Diaconescu's theorem, $\mathfrak{Geom}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}, J)) \simeq \mathbf{Flat}_J(\mathcal{C}, \mathcal{E})$ naturally in $\mathcal{E}$. A good candidate for a theory classified by $\mathbf{Sh}(\mathcal{C}, J)$ would then be a "theory $\mathbb{T}_J^{\mathcal{C}}$ of flat and J-continuous functors on $\mathcal{C}$".

The signature consists of one sort $\ulcorner c \urcorner$ for each object $c \in \mathcal{C}$ and one function symbol $\ulcorner f \urcorner : \ulcorner c \urcorner \longrightarrow \ulcorner d \urcorner$ for each arrow $f : c \longrightarrow d$ in $\mathcal{C}$. The axioms are the following :

- Functoriality :

$$\top \vdash_{x:\ulcorner c \urcorner} \ulcorner \mathrm{id}_c \urcorner(x) = x \qquad \text{for every } c \in \mathcal{C}$$

$$\top \vdash_{x:\ulcorner \mathrm{dom}(g) \urcorner} \ulcorner f \circ g \urcorner(x) = \ulcorner f \urcorner(\ulcorner g \urcorner(x)) \qquad \text{for every composable } f \text{ and } g$$

- Filteringness (equivalent to flatness) :

$$\top \vdash_{[]} \bigvee_{c \in \mathcal{C}} (\exists x : \ulcorner c \urcorner)\top$$

$$\top \vdash_{x:\ulcorner c \urcorner, y:\ulcorner d \urcorner} \bigvee_{c \xleftarrow{f} e \xrightarrow{g} d} (\exists z : \ulcorner e \urcorner)(\ulcorner f \urcorner(z) = x \wedge \ulcorner g \urcorner(z) = y) \qquad \text{for every } c, d \in \mathcal{C}$$

$$\ulcorner f \urcorner(x) = \ulcorner g \urcorner(x) \vdash_{x:\ulcorner c \urcorner} \bigvee_{fh = gh} (\exists z : \ulcorner \mathrm{dom}(h) \urcorner)\ \ulcorner h \urcorner(z) = x \qquad \text{for every parallel } f, g : c \longrightarrow d$$

- J-continuity :

$$\top \vdash_{x:\ulcorner c \urcorner} \bigvee_{i \in I} (\exists y : \ulcorner d_i \urcorner)\ \ulcorner f_i \urcorner(y) = x \qquad \text{for every } J\text{-cover } (f_i : d_i \longrightarrow c)_{i \in I}$$

It is straightforward to check that $\mathbb{T}_J^{\mathcal{C}}$-Mod$(\mathcal{E}) \simeq \mathbf{Flat}_J(\mathcal{C}, \mathcal{E})$, whence $\mathbf{Sh}(\mathcal{C}, J)$ is a classyfing topos for $\mathbb{T}_J^{\mathcal{C}}$. □

6 Theories of presheaf type

Theories classified by presheaf toposes are especially nice to work with, as this topos then has a canonical site of definition which is simpler than the syntactic site. Knowing this and a remarkable duality theorem, computations to switch between theories and topologies will become easy. From now on, theories will all be assumed geometric.

Definition 6.1. A (geometric) theory $\mathbb{T}$ is of *presheaf type* if it is classified by a presheaf topos.

Example 6.2. Because of theorem 5.12, every cartesian theory is of presheaf type.

Definition 6.3. Let $\mathbb{T}$ be a theory. A model $M \in \mathbb{T}$-Mod$(\mathbf{Set})$ is *finitely presentable* if $\mathrm{Hom}_{\mathbb{T}\text{-Mod}(\mathbf{Set})}(M, -) : \mathbb{T}\text{-Mod}(\mathbf{Set}) \longrightarrow \mathbf{Set}$ preserves filtered colimits (see definition A.7). In general, such an object is also called *compact*. It is equivalent to ask that every epimorphic family with codomain M contains a finite epimorphic subfamily.

A model M is *presented by* a geometric formula-in-context $\vec{x}.\phi$ if the functor $\mathrm{Hom}_{\mathbb{T}\text{-Mod}(\mathbf{Set})}(M, -)$ is isomorphic to $[\![\vec{x}.\phi]\!]_- : N \mapsto [\![\vec{x}.\phi]\!]_N$. Since ϕ is geometric, this implies that M is finitely presented.

Explicitely, an element of $[\![\vec{x}.\phi]\!]_N$ is given by elements $a_i \in A_{iN}$ such that ϕ holds for $a_1, \ldots, a_n$. Thus, M is presented by ϕ if and only if there are $a_i \in A_{iM}$, called the *generators* of M, such that ϕ holds for $a_1, \ldots, a_n$ and such that for any $N \in \mathbb{T}$-Mod$(\mathbf{Set})$ and any $b_i \in A_{iN}$ for which ϕ holds, there is a unique $f : M \longrightarrow N$ sending a_i to b_i for every $1 \leq i \leq n$.

Theorem 6.4. *Any theory of presheaf type $\mathbb{T}$ is classified by the presheaf topos* $[\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}), \mathbf{Set}]$.

Proof. Let $\mathbb{T}$ be classified by $\mathbf{Set}[\mathbb{T}] \simeq [\mathcal{C}, \mathbf{Set}]$ with $\mathcal{C}$ a small category. By theorem A.19.(iv), we can replace $\mathcal{C}$ by its Cauchy completion $\hat{\mathcal{C}}$. By Diaconescu's equivalence, we have $\mathbf{Flat}(\hat{\mathcal{C}}^{\text{op}}, \mathbf{Set}) \simeq \mathfrak{Geom}(\mathbf{Set}, [\hat{\mathcal{C}}, \mathbf{Set}]) \simeq \mathbb{T}\text{-Mod}(\mathbf{Set})$. One can check that the compact objects of $\mathbf{Flat}(\hat{\mathcal{C}}^{\text{op}}, \mathbf{Set})$ are exactly the retracts of the representable presheaves, and since $\hat{\mathcal{C}}$ is Cauchy-complete this states that $\hat{\mathcal{C}} \simeq (\mathbf{Flat}(\hat{\mathcal{C}}^{\text{op}}, \mathbf{Set}))_{\text{fp}}$, hence $\mathbf{Set}[\mathbb{T}] \simeq [\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}), \mathbf{Set}]$. ☐

Definition 6.5. Let $(\mathcal{C}, J)$ be a site. A *J-irreducible object* is an object whose only J-covering sieve is the maximal one. J is said to be *rigid* if for every object c, the collection of arrows from J-irreducible objects to c generates a J-covering sieve.

Theorem 6.6. *Let $(\mathcal{C}, J)$ be a small subcanonical site such that $\mathbf{y}_{\mathcal{C}}(\mathcal{C})$ is closed (in $\mathbf{Sh}(\mathcal{C}, J)$) under retracts. Then $\mathbf{Sh}(\mathcal{C}, J)$ is equivalent to a presheaf topos if and only if J is rigid.*

Idea of proof. Only the *if* part is really important here. If J is rigid, the comparison lemma (theorem 2.11) applied to the trivial topology shows that $\mathbf{Sh}(\mathcal{C}, J) \simeq \mathbf{Psh}(\mathcal{C}^{\text{irr}})$. The *only if* part can be found in [6], theorem 6.1.7, p.202. ☐

Definition 6.7. If $\mathbb{T}$ is a theory, a formula-in-context $\vec{x}.\phi$ is $\mathbb{T}$-*irreducible* if $\{\vec{x}.\phi\}$ is $J_{\mathbb{T}}$-irreducible in $\mathcal{C}_{\mathbb{T}}$.

Proposition 6.8. *Let $\mathbb{T}$ be a theory of presheaf type. Then a model of $\mathbb{T}$ in* **Set** *is finitely presentable if and only if it is presented by a $\mathbb{T}$-irreducible geometric formula.*

Proof. Since $\mathbb{T}$ is of presheaf type, there is a canonical equivalence

$$\tau : \mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J_{\mathbb{T}}) \simeq [\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}), \mathbf{Set}].$$

By theorems 6.6 and A.19.(iv), there is an equivalence $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J_{\mathbb{T}}) \simeq \mathbf{Psh}(\mathcal{C}_{\mathbb{T}}^{\text{irr}}) \simeq \mathbf{Psh}(\widehat{\mathcal{C}_{\mathbb{T}}^{\text{irr}}})$. The resulting equivalence $\mathbf{Psh}(\widehat{\mathcal{C}_{\mathbb{T}}^{\text{irr}}}) \simeq [\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}), \mathbf{Set}]$ restricts to an equivalence

$$l : \widehat{\mathcal{C}_{\mathbb{T}}^{\text{irr}}}^{\text{op}} \simeq \text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}).$$

Given $\{\vec{x}.\phi\} \in \widehat{\mathcal{C}_{\mathbb{T}}^{\text{irr}}}$, τ sends $\mathbf{y}_{\mathcal{C}_{\mathbb{T}}}(\{\vec{x}.\phi\}) = [\![\vec{x}.\phi]\!]_U$ (see remark 5.11 for the definition of U) to $\mathbf{y}_{\mathcal{C}_{\mathbb{T}}}(l(\{\vec{x}.\phi\})) = [\![\vec{x}.\phi]\!]_N$ (N being the universal model in $[\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}), \mathbf{Set}]$, given by $NA(M) = MA$, $Nf(M) = Mf$ and $NR(M) = MR$, which moreover verifies $[\![\vec{x}.\phi]\!]_N(M) = [\![\vec{x}.\phi]\!]_M$). So $l(\{\vec{x}.\phi\})$ is presented by $\vec{x}.\phi$. Since all representables in presheaf toposes are irreducible, $\vec{x}.\phi$ is $\mathbb{T}$-irreducible. So $\widehat{\mathcal{C}_{\mathbb{T}}^{\text{irr}}} \simeq \mathcal{C}_{\mathbb{T}}^{\text{irr}}$. l is then an equivalence $(\mathcal{C}_{\mathbb{T}}^{\text{irr}})^{\text{op}} \simeq \text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})$, as required : this equivalence sends objects and morphisms of $\mathcal{C}_{\mathbb{T}}$ to the models and homomorphisms that they represent. ☐

Typically, theories of presheaf type are those with few enough axioms for such a big category as a presheaf topos to be classifying it. We can expect that adding axioms amounts to removing sheaves from a classifying topos, and this is indeed the case, known as the duality theorem.

Definition 6.9. Let $\mathbb{T}$ be a theory. A *quotient* of $\mathbb{T}$ is a theory $\mathbb{T}'$ such that each axiom of $\mathbb{T}$ is provable in $\mathbb{T}'$. Two theories $\mathbb{T}$ and $\mathbb{S}$ are *syntactically equivalent* (denoted $\mathbb{T} \equiv_s \mathbb{S}$) if every (geometric) sequent is provable in $\mathbb{T}$ exactly when it is provable in $\mathbb{S}$.

Theorem 6.10 (Duality theorem)**.** *Let $\mathbb{T}$ be a theory. There is a bijection between quotients of $\mathbb{T}$ modulo syntactic equivalence and subtoposes of* $\mathbf{Set}[\mathbb{T}]$*, sending a quotient $\mathbb{T}'$ to its classifying topos.*

Ideas of the proof. Let $\mathbb{T}'$ be a quotient of $\mathbb{T}$ obtained by adding (geometric) axioms of the form $\phi \vdash_{\vec{x}} \psi$. Given a Grothendieck topos $\mathcal{E}$ and $M \in \mathbb{T}\text{-Mod}(\mathcal{E})$, the functor $F_M : \mathcal{C}_{\mathbb{T}} \longrightarrow \mathcal{E}$ sends the monomorphism $\{\vec{x}'.\phi \wedge \psi\} \xrightarrow{[\phi \wedge \psi \wedge \vec{x}' = \vec{x}]} \{\vec{x}.\phi\}$ to an epimorphism if and only if $[\![\vec{x}.\phi]\!]_M \leq [\![\vec{x}.\psi]\!]_M$, which is equivalent to $M \vDash (\phi \vdash_{\vec{x}} \psi)$. Thus, the $J_{\mathbb{T}}$-continuous flat functors $\mathcal{C}_{\mathbb{T}} \longrightarrow \mathcal{E}$ which send each of these monomorphisms to epimorphisms correspond to the $\mathbb{T}'$-models in $\mathcal{E}$. Let $J^{\mathbb{T}}_{\mathbb{T}'}$ be the smallest Grothendieck topology on $\mathcal{C}_{\mathbb{T}}$ containing the $J_{\mathbb{T}}$-covering sieves and the sieves containing such a monomorphism. By Diaconescu's equivalence, $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J^{\mathbb{T}}_{\mathbb{T}'})$ classifies $\mathbb{T}'$, and is also a subtopos of $\mathbf{Set}[\mathbb{T}]$.

To show that this assignment does not depend on the syntactic equivalence class of the quotient, let $\mathbb{T}_1$ and $\mathbb{T}_2$ be two syntactically equivalent quotients of $\mathbb{T}$. Since they have the same models in every Grothendieck topos $\mathcal{E}$, we get an equivalence

$$\begin{aligned}\mathfrak{Geom}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J^{\mathbb{T}}_{\mathbb{T}_1})) \simeq \mathbf{Flat}_{J_{\mathbb{T}_1}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E}) \simeq \mathbb{T}_1\text{-Mod}(\mathcal{E}) \\ = \mathbb{T}_2\text{-Mod}(\mathcal{E}) \simeq \mathbf{Flat}_{J_{\mathbb{T}_2}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E}) \simeq \mathfrak{Geom}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J^{\mathbb{T}}_{\mathbb{T}_2}))\end{aligned}$$

natural in $\mathcal{E}$. It follows from a Yoneda-like lemma that $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J^{\mathbb{T}}_{\mathbb{T}_1})$ and $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J^{\mathbb{T}}_{\mathbb{T}_2})$ are equivalent over $\mathbf{Psh}(\mathcal{C}_{\mathbb{T}})$, hence they represent the same local operators and $J^{\mathbb{T}}_{\mathbb{T}_1} = J^{\mathbb{T}}_{\mathbb{T}_2}$.

Let $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J) \rightarrowtail \mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J_{\mathbb{T}})$ be a subtopos and $\mathbb{T}^J$ the quotient of $\mathbb{T}$ with additional axioms the (geometric) sequents $\psi \vdash_{\vec{y}} \exists \vec{x}\theta$, where $[\theta] : \{\vec{x}.\phi\} \longrightarrow \{\vec{y}.\psi\}$ is a monomorphism generating a J-covering sieve. One can show that the natural equivalence $\mathbb{T}\text{-Mod}(\mathcal{E}) \simeq \mathbf{Flat}_{J_{\mathbb{T}}}(\mathcal{C}_{\mathbb{T}}, \mathcal{E})$ restricts to an equivalence $\mathbb{T}^J\text{-Mod}(\mathcal{E}) \simeq \mathbf{Flat}_J(\mathcal{C}_{\mathbb{T}}, \mathcal{E})$, thus $\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J)$ classifies $\mathbb{T}^J$.

It is clear that for any topology J on $\mathcal{C}_{\mathbb{T}}$ we have $J = J^{\mathbb{T}}_{\mathbb{T}^J}$, so we are left to show that $\mathbb{T}' \equiv_s \mathbb{T}^{J^{\mathbb{T}}_{\mathbb{T}'}}$ for any quotient $\mathbb{T}'$ of $\mathbb{T}$. Notice that given a Grothendieck topos $\mathcal{E}$, the $\mathbb{T}'$-models in $\mathcal{E}$ and the $\mathbb{T}^{J^{\mathbb{T}}_{\mathbb{T}'}}$-models in $\mathcal{E}$ both correspond to the $J^{\mathbb{T}}_{\mathbb{T}'}$-continuous flat functors $\mathcal{C}_{\mathbb{T}'} \longrightarrow \mathcal{E}$. Now let $U^{\mathbb{T}}_{\mathbb{T}'}$ be the image of $a_{J^{\mathbb{T}}_{\mathbb{T}'}} \circ \mathbf{y}_{\mathcal{C}_{\mathbb{T}}} \in \mathbf{Flat}_{J^{\mathbb{T}}_{\mathbb{T}'}}(\mathcal{C}_{\mathbb{T}}, \mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J^{\mathbb{T}}_{\mathbb{T}'}))$ under the equivalence with $\mathbb{T}'\text{-Mod}(\mathbf{Sh}(\mathcal{C}_{\mathbb{T}}, J^{\mathbb{T}}_{\mathbb{T}'}))$. This model is both a $\mathbb{T}'$- and a $\mathbb{T}^{J^{\mathbb{T}}_{\mathbb{T}'}}$-universal model, and one can deduce that $\mathbb{T}' \equiv_s \mathbb{T}^{J^{\mathbb{T}}_{\mathbb{T}'}}$ as required. □

Definition 6.11. According to this theorem, the (syntactic equivalence classes of) quotients of a theory $\mathbb{T}$ correspond to the topologies on $\mathbf{Psh}(\mathcal{C}_{\mathbb{T}})$ which contain $J_{\mathbb{T}}$. If $\mathbb{T}'$ is a quotient of a theory $\mathbb{T}$, the topology $J^{\mathbb{T}}_{\mathbb{T}'}$ on $\mathbf{Psh}(\mathcal{C}_{\mathbb{T}})$ as in the theorem is called the *associated* $\mathbb{T}$*-topology of* $\mathbb{T}'$. In the light of theorem 6.4, when $\mathbb{T}$ is of presheaf type this is equivalently a topology on $\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}$.

To make the computation $\mathbb{T}' \mapsto J^{\mathbb{T}}_{\mathbb{T}'}$ convenient in the setting where $\mathbb{T}$ is of presheaf type, we shall state a theorem making this topology more explicit when $\mathbb{T}'$ is presented in the right form.

Definition 6.12. Let $\mathbb{T}$ be a theory. Let $\vec{x}.\phi$ and $\vec{y}.\psi$ be formulae representing models $M_\phi, M_\psi \in \text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})$. Let θ be provably functional $\{\vec{y}.\psi\} \longrightarrow \{\vec{x}.\phi\}$. Then for any $N \in \mathbb{T}\text{-Mod}(\mathbf{Set})$, $[\![\theta]\!]_N \subset [\![\psi]\!]_N \times [\![\psi]\!]_N$ is the graph of a map $[\![\psi]\!]_N \longrightarrow [\![\phi]\!]_N$. If $\vec{a} \in [\![\psi]\!]_{M_\psi}$ is the tuple of generators of M_ψ, then there is a unique $\vec{b} \in [\![\phi]\!]_{M_\psi}$ such that $(\vec{a}, \vec{b}) \in [\![\theta]\!]_{M_\psi}$. Since $\vec{x}.\phi$ presents M_ϕ, there is a unique homomorphism $s_\theta : M_\phi \longrightarrow M_\psi$ sending the generators of M_ϕ to $\vec{b}$. Call this s_θ the arrow *presented by* θ.

Theorem 6.13. *Let* $\mathbb{T}$ *be a theory of presheaf type. Let* $\vec{x}_i.\phi_i$ $(i \in I)$ *be formulae presenting models* $M_i \in \text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})$*. For each* $i \in I$*, let* $\vec{y}^j_i.\psi^j_i$ *and* θ^j_i $(j \in J_i)$ *be formulae where each* ψ^j_i *presents a model* $M^j_i \in \text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})$ *and each* θ^j_i *is* $\mathbb{T}$*-provably functional from* ψ^j_i *to* ϕ_i*. Then the quotient* $\mathbb{T}'$ *obtained from* $\mathbb{T}$ *by adding the axioms*

$$\phi_i \vdash_{\vec{x}_i} \bigvee_{j \in J_i} \exists \vec{y}^j_i \theta^j_i$$

for $i \in I$ *has as its associated* $\mathbb{T}$*-topology the topology on* $\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}$ *generated by the sieves* S_i*, where each* S_i *is the dual of the cosieve on* M_i *generated by the arrows* $s^j_i : M_i \longrightarrow M^j_i$ *presented by* θ^j_i*.*

Ideas of the proof. Let J be the induced $\mathbb{T}$-topology of $\mathbb{T}'$ on $\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}$ and J' the topology generated by the sieves S_i. Also write σ_i for the axiom displayed in the theorem. We have to show that $J = J'$.

Let $F : \text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}} \longrightarrow \mathcal{E}$ be a flat functor into a Grothendieck topos $\mathcal{E}$, and M_F the $\mathbb{T}$-model in $\mathcal{E}$ naturally corresponding to F via Diaconescu's equivalence and the fact that $[\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}), \mathbf{Set}]$

classifies $\mathbb{T}$. One easily shows that $F(M_i) = [\![\vec{x}_i.\phi_i]\!]_{M_F}$ and that the graph of $F(s_i^j)$ is $[\![\vec{y}_i^j, \vec{x}_i.\theta_i^j]\!]_{M_F}$. From this and the equivalence

$$\mathbb{T}'\text{-Mod}(\mathcal{E}) \simeq \mathbf{Flat}_J(\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}, \mathcal{E})$$

we can show that any functor $F \in \mathbf{Flat}_J(\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}, \mathcal{E})$ sends each S_i to an epimorphic family. Apply this and lemma 2.27 to the geometric morphism

$$\mathbf{Sh}(\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}, J) \rightarrowtail [\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}), \mathbf{Set}]$$

to see that the sheafification $a_J : [\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set}), \mathbf{Set}] \longrightarrow \mathbf{Sh}(\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}, J)$ sends each monomorphism $S_i \rightarrowtail \text{Hom}_{\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})}(M_i, -)$ to an isomorphism, hence $J' \subset J$.

To see that $J = J'$ it is in fact enough now to show that for any Grothendieck topos $\mathcal{E}$,

$$\mathbf{Flat}_{J'}(\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}, \mathcal{E}) \subset \mathbf{Flat}_J(\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}, \mathcal{E}).$$

Since J' contains the S_i's, the $\mathbb{T}$-model corresponding to $F \in \mathbf{Flat}_{J'}(\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}, \mathcal{E})$ is already a $\mathbb{T}'$-model. Since $\mathbb{T}'\text{-Mod}(\mathcal{E}) \simeq \mathbf{Flat}_J(\text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})^{\text{op}}, \mathcal{E})$, F is J-continuous. □

Because it is noteworthy, we state a last theorem to underline the simplicity of presheaf type theories :

Theorem 6.14. *Let $\mathbb{T}$ be a theory of presheaf type and $\mathcal{A}$ a full subcategory of* f.p. $\mathbb{T}$-Mod(**Set**). *Then the $\mathcal{A}$-completion $\mathbb{T}_{\mathcal{A}}$ of $\mathbb{T}$ (the theory of geometric sequents valid in all models in $\mathcal{A}$) is of presheaf type, classified by* $[\mathcal{A}, \mathbf{Set}]$.

Proof. The inclusion $\mathcal{A} \rightarrowtail \text{f.p.}\,\mathbb{T}\text{-Mod}(\mathbf{Set})$ induces a canonical geometric inclusion

$$i : [\mathcal{A}, \mathbf{Set}] \rightarrowtail \mathbf{Set}[\mathbb{T}].$$

The quotient of $\mathbb{T}$ corresponding to this subtopos is exactly the collection of sequents which hold in every model in $\mathcal{A}$. □

7 The Zariski topos

From this section onwards, we take a look at examples of classifying toposes in algebraic geometry. The first example is the Zariski topos, which is one of the most natural toposes to consider when dealing with schemes.

Definition 7.1. Let k be a (commutative, unitary) ring. A *finitely presented k-algebra* is a k-algebra of the form $k[X_1, \dots, X_n]/(f_1, \dots, f_m)$. A morphism of schemes $f : T \longrightarrow S$ is *locally of finite presentation* if there is an affine open cover $T = \bigcup_{i \in I} U_i$ with $U_i = \text{Spec}\, A_i$ and affine open subsets $V_i = \text{Spec}\, K_i \subset S$ for $i \in I$ such that for all i, $f(U_i) \subset V_i$ and the induced ring homomorphisms $K_i \longrightarrow A_i$ make every A_i into a finitely presented K_i-algebra.

Definition 7.2. Let S be a scheme. The *(big) Zariski site* is the category $\text{ZAR}^*(S)$ with objects the affine schemes $\text{Spec}\, A \longrightarrow S$ locally of finite presentation over S and scheme morphisms over S, and with a sieve R on $T = \text{Spec}\, A$ being in $J_{\text{ZAR}^*}(T)$ if there are $a_1, \dots, a_k \in A$ with sum 1 and such that all the duals of localizations $(A \longrightarrow A[a_i^{-1}])^{\text{op}}$ are in R.

The resulting topos $S_{\text{ZAR}} = \mathbf{Sh}(\text{ZAR}^*(S), J_{\text{ZAR}^*})$ is the *big Zariski topos* of S.

Remark 7.3. *In general, the Zariski site* $\text{ZAR}(S)$ *has* all schemes *locally of finite presentation over S, and a sieve is J_{ZAR}-covering if it contains immersions of open subsets $U_i \rightarrowtail T$ such that $T = \bigcup_{i \in I} U_i$. The comparison lemma (theorem 2.11) shows that the resulting topos is equivalent when only taking affine schemes (*$\mathbf{Sh}(\text{ZAR}(S), J_{\text{ZAR}}) \simeq \mathbf{Sh}(\text{ZAR}^*(S), J_{\text{ZAR}}|_{\text{ZAR}^*(S)})$*), and the Grothendieck topology J_{ZAR} restricted to* $\text{ZAR}^*(S)$ *is in fact the same as J_{ZAR^*} defined above.*

When S is an affine scheme, there is a simple syntactic presentation for S_{ZAR} (that is, a theory classified by S_{ZAR}) which shall be given here. In the general case, it appears that the simplest theory we know to be classified by S_{ZAR} is quite intricate and can be found in [9], theorem 3.7.6, p.33.

Until the end of this section, let K be a ring and $S = \operatorname{Spec} K$.

Lemma 7.4. *There is an equivalence of categories* $\mathrm{ZAR}^*(S) \simeq (K\mathrm{-Alg})^{\mathrm{op}}_{\mathrm{fp}}$.

Proof. Since Spec gives an equivalence between the affine schemes and the opposite of the category of rings, the claim follows from the observation that $\operatorname{Spec} A \longrightarrow \operatorname{Spec} K$ is locally of finite presentation if and only if A is a a finitely presented K-algebra. □

Lemma 7.5. J_{ZAR^*} *is generated (on* $(K\mathrm{-Alg})^{\mathrm{op}}_{\mathrm{fp}}$*) by the families of duals of localizations :*

$$\big(K[X_1,\ldots,X_n]/(X_1+\cdots+X_n-1) \longrightarrow K[X_1,\ldots,X_n,X_i^{-1}]/(X_1+\cdots+X_n-1)\big)^{\mathrm{op}}$$

for $1 \leq i \leq n$.

Proof. Any family of duals of localizations $(A \longrightarrow A[a_i^{-1}])^{\mathrm{op}}$ with $\sum_{i=1}^n a_i = 1$ as in definition 7.2 is the pullback in $\mathrm{ZAR}^*(S)$ of this universal family along the suitable $K[X_1,\ldots,X_n]/(X_1+\cdots+X_n-1) \longrightarrow A$ sending X_i to a_i. □

Definition 7.6. The *theory* $\mathbb{K}$ *of* K*-algebras* is the theory of rings (see example 4.7.(iii)) with one new constant symbol $c_\lambda : A$ for every $\lambda \in K$, with the axioms

$$\top \vdash_{[]} c_0 = 0 \qquad\qquad \top \vdash_{[]} c_1 = 1$$
$$\top \vdash_{[]} c_\lambda + c_\mu = c_{\lambda+\mu} \qquad\qquad \top \vdash_{[]} c_\lambda \times c_\mu = c_{\lambda\mu}$$

for every $\lambda, \mu \in K$. The quotient of this theory obtained by adding the locality axiom (see example 4.7.(iii)) is the *theory* $\mathrm{loc}-\mathbb{K}$ *of local* K*-algebras.*

Lemma 7.7. *The finitely presentable models in* $\mathbb{K}$-$\mathrm{Mod}(\mathbf{Set})$ *are exactly the finitely presented* K*-algebras.*

Proof. It is easy to see that $K[X]$ is finitely presentable. Since any $A = K[X_1,\ldots,X_n]/(f_1,\ldots,f_m)$ can be constructed from $K[X]$ by finite colimits, the finitely presented algebras are all finitely presentable $\mathbb{K}$-models. This comes from the fact that $K[X_1,\ldots,X_n] = K[X]^{\otimes n}$ and A is then the coequalizer of $X_i \mapsto f_i$ and $X_i \mapsto 0$.

Let $A \in \mathrm{f.p.}\mathbb{K}\text{-}\mathrm{Mod}(\mathbf{Set})$. The set I of finitely generated sub-K-algebras (that is, quotients of polynomial rings $K[X_1,\ldots,X_k]$) of A if clearly a filtered category, and the inclusions to A form a cocone in $\mathbb{K}$-$\mathrm{Mod}(\mathbf{Set})$. We get a canonical homomorphism $\mathrm{colim}_{B\in I} B \longrightarrow A$. It is surjective since every $a \in A$ is in the image of $K[X] \longrightarrow A$ mapping X to a, and injective since two elements $a_1 \in A_1$ and $a_2 \in A_2$ can be mapped into a common finitely generated subalgebra of A. Since A is finitely presentable, there is an image of id_A under the isomorphism

$$\mathrm{Hom}_{\mathbb{K}\text{-}\mathrm{Mod}(\mathbf{Set})}(A,A) \cong \mathrm{Hom}_{\mathbb{K}\text{-}\mathrm{Mod}(\mathbf{Set})}(A, \mathrm{colim}_{B\in I} B) \cong \mathrm{colim}_{B\in I} \mathrm{Hom}_{\mathbb{K}\text{-}\mathrm{Mod}(\mathbf{Set})}(A,B)$$

which has to be a section of some inclusion $B \subset A$ with $B \in I$, hence $A \in I$ and A is finitely generated.

Let $q : K[X_1,\ldots,X_n] \longrightarrow A$ be a surjective K-algebra homomorphism, which exists since we just showed that A is finitely generated. Again, the set J of finitely generated ideals of $K[X_1,\ldots,X_n]$ contained in the kernel of q is a filtered category. In the same manner, we get a homomorphism

$$\mathrm{colim}_{\mathfrak{j}\in J} K[X_1,\ldots,X_n]/\mathfrak{j} \longrightarrow A$$

which is an isomorphism as well. Again, since A is finitely presentable, take the image of id_A under the isomorphism

$$\mathrm{Hom}_{\mathbb{K}\text{-}\mathrm{Mod}(\mathbf{Set})}(A,A) \cong \mathrm{colim}_{\mathfrak{j}\in J} \mathrm{Hom}_{\mathbb{K}\text{-}\mathrm{Mod}(\mathbf{Set})}(A, K[X_1,\ldots,X_n]/\mathfrak{j})$$

to see that there is a finitely presented K-algebra $B = K[X_1, \dots, X_n]/\mathfrak{j}$ and a section $s : A \longrightarrow B$ of the quotient $q_B : B \longrightarrow A$. Notice that for any $B' = K[X_1, \dots, X_n]/\mathfrak{l}$ with $\mathfrak{l} \supset \mathfrak{j}$ in J we have $q_{B'}rs = \mathrm{id}_A$. To see that A is finitely presented, we just need to see that there is a B' isomorphic to A, that is such that $rsq_{B'} = \mathrm{id}_{B'}$. Since B is finitely presented, it is in f.p.$\mathbb{K}$-Mod(**Set**) and as before, there is an isomorphism

$$\mathrm{colim}_{\mathfrak{l}\in J} \mathrm{Hom}_{\mathbb{K}\text{-Mod}(\mathbf{Set})}(B, K[X_1, \dots, X_n]/\mathfrak{l}) \cong \mathrm{Hom}_{\mathbb{K}\text{-Mod}(\mathbf{Set})}(B, A).$$

Since sq_B and id_B are equal in the right hand side, they are in the left hand side. So there is B' with $rsq_B = r$. Thus, $rsq_{B'}r = rsq_B = r$ and since r is surjective, $rsq_{B'} = \mathrm{id}_{B'}$, so $B' \cong A$ and A is indeed finitely presented. □

Remark 7.8. *Since it is algebraic, the theory of K-algebras is of presheaf type (see theorem 5.12). Also, the finitely presentable models of $\mathbb{K}$ are exactly the finitely presented K-algebras. Hence,* $[\text{f.p.}\mathbb{K}\text{-Mod}(\mathbf{Set}), \mathbf{Set}] \simeq \mathbf{Psh}(\mathrm{ZAR}^*(\mathrm{Spec}\, K))$.

Theorem 7.9. *The Zariski topos S_{ZAR} of an affine scheme $S = \mathrm{Spec}\, K$ classifies the theory of local K-algebras.*

Proof. According to the above remark, it is enough to prove that J_{ZAR^*} is the induced $\mathbb{K}$-topology on f.p.$\mathbb{K}$-Mod(**Set**) of the theory loc$-\mathbb{K}$. Fortunately, the locality axioms are of the required form

$$\phi_n \vdash_{\vec{x}_n} \bigvee_{1\le i\le n} \exists \vec{y}_n^i \theta_n^i$$

as in theorem 6.13, with :

$$\begin{aligned}
\phi_n &= (x_1 + \cdots + x_n = 1) \\
\psi_n^i &= (x'_1 + \cdots + x'_n = 1) \wedge (x'_i \times y = 1) \\
\theta_n^i &= \psi_n^i \wedge (x_1 = x'_1) \wedge \ldots \wedge (x_n = x'_n).
\end{aligned}$$

Since the model $M_{\phi_n} = K[X_1, \dots, X_n]/(X_1 + \cdots + X_n - 1)$ is presented by ϕ_n and the model $M_{\psi_n^i} = K[X_1, \dots, X_n, X_i^{-1}]/(X_1 + \cdots + X_n - 1)$ is presented by ψ_n^i, it is enough to show that θ_n^i presents the corresponding localization homomorphism $M_{\phi_n} \longrightarrow M_{\psi_n^i}$. Set x'_i to be X_i and y to be X_i^{-1} in the presentation ψ_n^i of $M_{\psi_n^i}$ to see that θ_n^i is the unique morphism $M_{\phi_n} \longrightarrow M_{\psi_n^i}$ sending X_i to X_i, that is the localization. □

Remark 7.10. *In particular, let $K = \mathbf{Z}$ to notice that the (usual) local rings are exactly the points of the Zariski topos* $(\mathrm{Spec}\, \mathbf{Z})_{\mathrm{ZAR}}$.

8 The crystalline topos

In this section we present another topos used in algebraic geometry, playing a key role in crystalline cohomology. The interest is that the crystalline topos involves more data, which makes a syntactic presentation harder to find than for the Zariski topos. In rings of characteristic $p > 0$, it is not possible to divide by p hence the polynomial derivative is never surjective. To overcome this, there is a notion of divided powers structure which will allow us to divide by p. Once again, only the affine case is treated, the general result can be found in [9], theorem 5.8.3, p.75.

Definition 8.1. Let A be a ring and $I \subset A$ an ideal. A *divided powers structure* (or *PD-structure*) γ on I is a family $(\gamma_n)_{n\in\mathbf{N}}$ of maps $I \longrightarrow I$ such that for all $n, m \geq 0$, $a \in A$ and $x, y \in I$:

(i) $\gamma_0(x) = 1$ and $\gamma_1(x) = x$;
(ii) $\gamma_n(x)\gamma_m(x) = \frac{(n+m)!}{n!m!}\gamma_{n+m}(x)$;

(iii) $\gamma_n(ax) = a^n\gamma_n(x)$;
(iv) $\gamma_n(x+y) = \sum_{i=0}^{n}\gamma_i(x)\gamma_{n-i}(y)$;
(v) $\gamma_n(\gamma_m(x)) = \frac{(nm)!}{n!(m!)^n}\gamma_{nm}(x)$.

The numbers $\frac{(n+m)!}{n!m!}$ and $\frac{(nm)!}{n!(m!)^n}$ are always integers. Typically, $\gamma_n(x)$ should be thought of as a legitimate way to write $x^n/n!$. The data (A, I, γ) is called a *divided powers ring* (or *PD-ring*). A *divided powers homomorphism* $(A, I, \gamma) \longrightarrow (B, J, \delta)$ is a ring homomorphism $A \longrightarrow B$ sending I into J and commuting with the PD-structures. This defines the category **PDRing** of PD-rings. (A, I, γ)-algebras are defined as the coslice category $(A, I, \gamma)/\mathbf{PDRing}$.

The ideal I is *PD-generated* by a subset $S \subset I$ if it is generated as an ideal by the $\gamma_n(s)$ for $s \in S$ and $n \geq 1$.

The polynomial ring $A[X]$ has a natural PD-structure $(I[X], \gamma')$ such that (A, I, γ)-algebra homomorphisms $(A[X], I[X], \gamma') \longrightarrow (B, J, \delta)$ correspond to the elements of B. Similarly, define the PD-polynomial algebra $A\langle X\rangle$ as the A-module freely generated by the $\gamma_n(X)$ for $n \geq 0$, with the PD-ideal generated by I and the $\gamma_n(X)$ for $n \geq 1$. The (A, I, γ)-algebra homomorphisms $A\langle X\rangle \longrightarrow (B, J, \delta)$ correspond to the elements of J.

Definition 8.2. We turn definition 8.1 into actual geometric theories.

(i) Let $\mathbb{R}$ be the theory of rings (see example 4.7) ;
(ii) Let $\mathbb{I}$ be the theory of rings with a new relation symbol $I \rightarrowtail A$ and the axioms

$$\top \vdash I(0)$$
$$I(x) \wedge I(y) \vdash_{x,y} I(x+y)$$
$$I(x) \vdash_{\lambda,x} I(\lambda \times x)$$

(iii) Let **nil** be the additional axiom $I(x) \vdash_x \bigvee_{n\in\mathbf{N}} x^n = 0$;
(iv) Let $\mathbb{PDI}$ be the theory of rings with a new sort S_I and new functions symbols $i : S_I \longrightarrow A$, $0 : S_I$, $+ : S_I \times S_I \longrightarrow S_I$, $\times : A \times S_I \longrightarrow S_I$ and $\gamma_n : S_I \longrightarrow S_I$ for $n \geq 0$, the axioms stating that i is an A-module homomorphism, and the axioms

$$i(x) = i(y) \vdash_{x,y:S_I} x = y$$
$$\top \vdash_x \gamma_0(x) = 1$$
$$\top \vdash_x \gamma_1(x) = x$$
$$\top \vdash_{x,y} i(\gamma_n(x))\gamma_m(x) = \frac{(n+m)!}{n!m!}\gamma_{n+m}(x)$$
$$\top \vdash_{\lambda,x} \gamma_n(\lambda \times x) = \lambda^n \times \gamma_n(x)$$
$$\top \vdash_{x,y} \gamma_n(x+y) = \sum_{i=0}^{n} i(\gamma_i(x)) \times \gamma_{n-i}(y)$$
$$\top \vdash_{x,y} \gamma_n(\gamma_m(x)) = \frac{(nm)!}{n!(m!)^n}\gamma_{nm}(x)$$

(v) Define $\mathbb{PD}$ by extending the theory $\mathbb{R} + \mathbb{I}$ with the extension $\mathbb{PDI}$ and the axiom

$$\exists x : S_I,\ i(x) = y \dashv\vdash_y I(y)$$

to state that the inclusion i is the inclusion of the ideal I (we get two theories $\mathbb{R} + \mathbb{I} + \mathbb{PD}$ and $\mathbb{R} + \mathbb{PDI}$ which are syntactically equivalent, so they are in particular Morita-equivalent) ;
(vi) Let K be a ring a R a K-algebra. Let $(K, R)-$Alg be the disjoint union of the theories of K-algebras and R-algebras, with a new function symbol $f : A \longrightarrow B$ (A being the sort of the K-algebras and B the sort of the R-algebras) and axioms stating that f is a K-algebra homomorphism ;

(vii) Let **surj** be the additional axiom $\top \vdash_{y:B} \exists x : A,\ f(x) = y$ and $(K,R)-\mathrm{Quot} = (K,R)-\mathrm{Alg} + \mathbf{surj}$;

(viii) If $I \subset K$ is an ideal, let $\mathbb{I}_I$ be the extension of the theory of K-algebras by $\mathbb{I}$ (see (ii)) and the axiom $\top \vdash_\lambda I(c_\lambda)$ (see definition 7.6) for every $\lambda \in I$;

(ix) If (K, I, δ) is a PD-ring, let $\mathbb{PD}_\delta$ be the extension of $\mathbb{I}_I$ by $\mathbb{PD}$ (see (v)) and the axioms $i(x) = c_\lambda \vdash_{x:S_I} i(\gamma_n(x)) = c_{\delta_n(\lambda)}$ for $\lambda \in I$.

Lemma 8.3. *Let (A, I, γ) be a PD-ring. If I is PD-generated by nilpotent elements, then it is a nilpotent ideal.*

Proof. It is enough to show that if $a \in I$ is nilpotent, then every $\gamma_n(a)$ is nilpotent. Let $m \geq 1$ be such that $a^m = 0$.

Compute on the canonical PD-structure on $\mathbf{Q}[X]$:

$$\gamma_n(X)^k = \frac{1}{(n!)^k} X^{kn} = \frac{(kn-m)!}{(n!)^k} \gamma_{kn-m}(X) X^m.$$

Notice that $\frac{(kn-m)!}{(n!)^k}$ is an integer when k is large enough, hence this stays true in $\mathbf{Z}\langle X\rangle \subset \mathbf{Q}[X]$. After applying the unique PD-morphism $\mathbf{Z}\langle X\rangle \longrightarrow A$ sending X to a (see definition 8.1), we get $\gamma_n(a)^k = \frac{(kn-m)!}{(n!)^k} \gamma_{kn-m}(a) a^m = 0$. □

Definition 8.4. A *PD-scheme* is a triple (S, I, γ) where S is a scheme, $I \subset \mathcal{O}_S$ is a quasi-coherent sheaf of ideals (every $I(U)$ is a sub-$\mathcal{O}_S(U)$-module of $\mathcal{O}_S(U)$, the restriction maps of I are compatible with those of $\mathcal{O}_S$ and for each $x \in S$ there is an open neighborhood U in which there is an exact sequence $\mathcal{O}_S^{\oplus I}|_U \longrightarrow \mathcal{O}_S^{\oplus J}|_U \longrightarrow I|_U \longrightarrow 0$) and γ is a PD-structure on I, that is such that every $(\mathcal{O}_S(U), I(U), \gamma)$ is a PD-ring. Equivalently, this is a model of $\mathbb{PD}$ in $\mathbf{Sh}(S)$.

It is a theorem of basic algebraic geometry that the quasi-coherent sheaves of ideals $I \subset \mathcal{O}_S$ correspond bijectively with closed immersions $i : U \longrightarrow S$, by $I = \mathrm{Ker}(i^\sharp)$. Thus, given a PD-scheme (S, I, γ) there is a natural closed immersion $V(I) \longrightarrow S$ whose comorphism has kernel I.

A *morphism of PD-schemes* $(S, I, \gamma) \longrightarrow (S', I', \gamma')$ is a morphism of schemes $f : S \longrightarrow S'$ such that $f^\sharp : f^*\mathcal{O}_{S'} \longrightarrow \mathcal{O}_S$ is a $\mathbb{PD}$-model homomorphism. Explicitely, it means that $(f^{-1}I')\mathcal{O}_S \subset I$ and that every $(\mathcal{O}_{S'}(U'), I'(U'), \gamma') \longrightarrow (\mathcal{O}_S(f^{-1}U'), I(f^{-1}U'), \gamma)$ with U' open in S' is a PD-homomorphism. In particular, such a morphism induces a morphism between the closed immersions $V(I) \longrightarrow V(I')$.

A morphism of PD-schemes $S \longrightarrow S'$ is *locally of finite PD-presentation* if there are affine open covers $S = \bigcup \operatorname{Spec} A_i$ and $S' = \bigcup \operatorname{Spec} K_j$ such that for every i there is a map $\operatorname{Spec} A_i \longrightarrow \operatorname{Spec} K_{j_i}$ whose induced $K_{j_i} \longrightarrow A_i$ is of the form $(K, I, \gamma) \longrightarrow K\langle X_1, \ldots, X_n\rangle[Y_1, \ldots, Y_m]/(r_1, \ldots, r_k)$.

A *PD-thickening* is a PD-scheme (S, I, γ) such that the $\mathbb{PD}$-model $(\mathcal{O}_S, I, \gamma)$ in $\mathbf{Sh}(S)$ satisfies the axiom **nil**.

If (S, I, γ) is a PD-scheme and $X \longrightarrow V(I)$ is a morphism of schemes, a *PD-thickening over S and X* is a PD-thickening (T, J, δ) together with a PD-morphism $(T, J, \delta) \longrightarrow (S, I, \gamma)$ and a morphism of schemes $V(J) \longrightarrow X$ over S. A *morphism of PD-thickenings over S and X* is a PD-morphism $(T, J, \delta) \longrightarrow (T', J', \delta')$ over (S, I, γ) such that $V(J) \longrightarrow V(J')$ is a morphism over X :

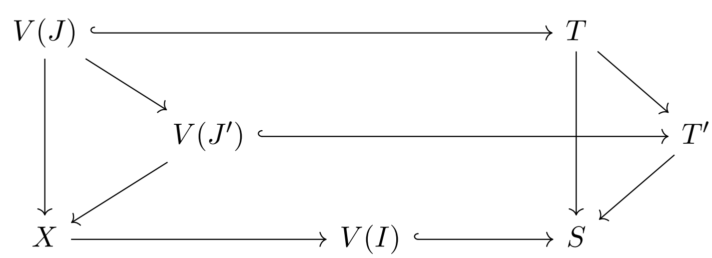

The *(big) crystalline site* $\mathbf{Cris}(X/S)$ is the category of PD-thickenings over X and S for which $T \longrightarrow S$ is locally of finite PD-presentation, and their morphisms, with the same topology as in definition 7.2 : a sieve on (T, J, δ) is in the topology if it contains the inclusions $(T_i, J|_{T_i}, \delta|_{T_i}) \rightarrowtail (T, J, \delta)$ of an open cover $T = \bigcup T_i$. As in definition 7.2 and remark 7.3, thanks to the comparison lemma (theorem 2.11) we can restrict to the affine schemes without changing the resulting topos, which we call $(X/S)_{\mathbf{Cris}}$.

Proposition 8.5. *Let $\mathbb{T}$ be a theory of presheaf type and A a sort of the signature of $\mathbb{T}$. Then the theory $\mathbb{T}'$ obtained by adding a constant symbol $c : A$ is also of presheaf type, and the finitely presentable models of $\mathbb{T}'$ are exactly the finitely presentable models of $\mathbb{T}$.*

Ideas of proof. Let $M_{\mathbb{T}}$ be the universal model in the classifying topos $\mathbf{Set}[\mathbb{T}]$. Then one can show that $\mathbb{T}'$ is classified by $\mathbf{Set}[\mathbb{T}]/[\![A]\!]_{M_{\mathbb{T}}}$. If $\mathbf{Set}[\mathbb{T}] = \mathbf{Psh}(\mathcal{C})$, then the equivalence

$$\mathbf{Psh}(\mathcal{C})/X \simeq \mathbf{Psh}(\int_{\mathcal{C}} X)$$

taken at $X = [\![A]\!]_{M_{\mathbb{T}}}$ shows that $\mathbb{T}'$ is of presheaf type. Notice that there is an equivalence $\mathbb{T}'\text{-Mod}(\mathbf{Set}) \simeq \int^{\mathbb{T}\text{-Mod}(\mathbf{Set})} [\![A]\!]_{-}$ over $\mathbb{T}\text{-Mod}(\mathbf{Set})$. The functor $[\![A]\!]_{-}$ preserves filtered colimits, and one can deduce that the projection functor $\pi : \int^{\mathbb{T}\text{-Mod}(\mathbf{Set})} [\![A]\!]_{-} \longrightarrow \mathbb{T}\text{-Mod}(\mathbf{Set})$ preserves and reflects compact objects. ☐

Lemma 8.6. *Let (K, I, γ) be a PD-ring. The finitely presentable models of the theory $\mathbb{T}_0$ of K-algebras $+\mathbb{I}_I + \mathbb{PD}_\gamma$ are exactly the PD-rings over K of the finite PD-presentation form :*

$$K\langle X_1, \dots, X_n\rangle[Y_1, \dots, Y_m]/(r_1, \dots, r_k).$$

Proof. The theory $\mathbb{T}_0$ is syntactically equivalent to a Horn theory, up to adding constants. But the finitely presentable models of a Horn theory are all presented by a Horn formula (see proposition 6.8). Here, $\vec{x} : S_I^n, \vec{y} : A^m.\top$ presents $K\langle X_1, \dots, X_n\rangle[Y_1, \dots, Y_m]$. Since $\mathbb{PD}_\gamma$ asks the interpretation of the symbol i to be injective, any atomic formula is provably equivalent to one of the form $r = 0$ where r is a term representing an element of $K\langle X_1, \dots, X_n\rangle[Y_1, \dots, Y_m]$. A Horn formula is just a conjunction of atomic formulae, which concludes. ☐

Lemma 8.7. *Let (K, I, γ) be a PD-ring with I finitely PD-generated, and R be a finitely presented K/I-algebra. Let $\mathbb{T} = (K, R)\text{−Quot} + \mathbb{PD}_\gamma + \mathbf{nil}$. Then $\mathbb{T}$ is of presheaf type and its finitely presentable models are the $(A, J, \delta, B = A/J)$ where (A, J, δ) is of finite PD-presentation over (K, I, γ).*

Sketch of proof. Let $\mathbb{T}_0$ be the cartesian theory of K-algebras $+ \mathbb{I}_I + \mathbb{PD}_\gamma$. By lemma 8.6, the objects of f.p.$\mathbb{T}_0$-Mod($\mathbf{Set}$) are the PD-rings of finite presentation over (K, I, γ). Since every $\mathbb{T}_0$-model in $\mathbf{Set}$ gives exactly one $\mathbb{T}'$-model (with $\mathbb{T}' = (K, K/I)\text{−Quot} + \mathbb{PD}_\gamma + \mathbf{nil}$), the forgetful functor f.p.$\mathbb{T}'$-Mod($\mathbf{Set}$) $\longrightarrow$ f.p.$\mathbb{T}_0$-Mod($\mathbf{Set}$) is an equivalence of categories. Since R is a finitely presented K/I-algebra, proposition 8.5 ensures that $(K, R)\text{−Quot} + \mathbb{PD}_\gamma$ is of presheaf type, with finitely presentable models those of finite PD-presentation over K.

Now let $M = (A, J, \delta, B = A/J)$ be a finitely presentable $((K, R)\text{−Quot} + \mathbb{PD}_\gamma)$-model in $\mathbf{Set}$. Since I is finitely PD-generated and A is of finite PD-presentation over K, J is also finitely PD-generated. Let $(a_1, \dots, a_n)$ be a family of generators of J. Let $J_{\mathbf{nil}}$ be the induced $((K, R)\text{−Quot} + \mathbb{PD}_\gamma)$-topology of $\mathbb{T}$. Then the covering sieve S_i given by theorem 6.13 is the cosieve of all $M \longrightarrow M'$ sending a_i to a nilpotent element. It is generated by the family $M \longrightarrow (A/J_n, J/J_n, \gamma, B)$ where J_n is the ideal PD-generated by a_i^n. Each of these models are still finitely generated, and their ideal J/J_n are each PD-generated by the images of the PD-generators of J. This gives a cover of M by models where the PD-ideal is PD-generated by nilpotent elements, and thanks to lemma 8.3, these ideals are nilpotent. This shows that $J_{\mathbf{nil}}$ is a rigid topology, hence by theorem 6.6 $\mathbb{T}$ is of presheaf type. ☐

Lemma 8.8. *Let (K, I, γ) be a PD-ring and R a K/I-algebra. Let $\mathbb{T} = (K, R)\text{−Quot} + \mathbb{PD}_\gamma + \mathbf{nil}$. The topology on $\mathbb{T}\text{-Mod}(\mathbf{Set})^{\mathrm{op}}$ seen as a full subcategory of $\mathbf{Cris}(\operatorname{Spec} R/\operatorname{Spec} K)$ has cosieves on objects $(A, J, \delta, B = A/J)$ those which contain the canonical localization arrows $(A, J, \delta, B = A/J) \longrightarrow (A[a_i^{-1}], J_{a_i}, \delta_{a_i}, B[a_i^{-1}])$ with $(a_1, \dots, a_n) = A$.*

Proof. A cosieve on $(A, J, \delta, B = A/J)$ is covering if it generates a covering sieve in $\mathbf{Cris}(\operatorname{Spec} R/\operatorname{Spec} K)$. Such a sieve contains an open cover of $\operatorname{Spec} A$ if and only if it contains a cover by principal open subsets $D(a_i)$, if and only if $(a_1, \ldots, a_n) = A$. □

Theorem 8.9. *Let (K, I, γ) be a PD-ring with I finitely PD-generated, and R a finitely presented K/I-algebra. Then the crystalline topos*

$$(X/S)_{\mathbf{Cris}} = \mathbf{Sh}(\mathbf{Cris}(\operatorname{Spec} R/\operatorname{Spec} K))$$

classifies the theory

$$(K, R)\text{−Quot} + \mathbb{PD}_\gamma + \mathbf{nil} + \mathbf{loc}$$

where **loc** *is the axiom of locality (see example 4.7.(iii)).*

Proof. Let $\mathbb{T} = (K, R)\text{−Quot}+\mathbb{PD}_\gamma+\mathbf{nil}$. Then $\mathbf{Cris}(\operatorname{Spec} R/\operatorname{Spec} K)^{\mathrm{op}}$ is the subcategory of $\mathbb{T}\text{-Mod}(\mathbf{Set})$ where A is of finite PD-presentation over K. By lemma 8.7, this is f.p. $\mathbb{T}\text{-Mod}(\mathbf{Set})$. Hence, the presheaf topos

$$\mathbf{Psh}(\mathbf{Cris}(\operatorname{Spec} R/\operatorname{Spec} K))$$

classifies $\mathbb{T}$. Thanks to lemma 8.8, it is for the exact same reason as in theorem 7.9 that the induced topology of $\mathbb{T} + \mathbf{loc}$ on $\mathbf{Cris}(\operatorname{Spec} R/\operatorname{Spec} K)$ is the topology defining $(X/S)_{\mathbf{Cris}}$. □

A Categorical prerequisites

A.1 Glossary

Definition A.1. A category is *small* if its objects and morphisms form sets. A category is *essentially small* if it is equivalent to a small category. A category is *locally small* if between any two objects there is a set of morphisms.

Definition A.2. A functor $F : \mathcal{C} \longrightarrow \mathcal{D}$ is full (resp. faithful) if given two objects c and d, the induced map $\mathrm{Hom}_{\mathcal{C}}(c, d) \longrightarrow \mathrm{Hom}_{\mathcal{C}}(Fc, Fd)$ is surjective (resp. injective). An *embedding* is a faithful functor that is injective on objects ; every fully faithful functor is equivalent to a fully faithful functor that is injective on objects. Given a category $\mathcal{C}$, the *full subcategory of $\mathcal{C}$ on the objects $\mathcal{O} \subset$ Objects$(\mathcal{C})$* is the category with objects $\mathcal{O}$ and all the morphisms from $\mathcal{C}$ between them.

Definition A.3. In a category, a family of morphisms with same codomain $(f_i : X_i \longrightarrow X)_i$ is *epimorphic* if for any two morphisms $g, h : X \longrightarrow Y$ such that $gf_i = hf_i$ for all i, we have $g = h$.

Definition A.4. A functor F is *conservative* if for any morphism f, $F(f)$ being an isomorphism implies that f was already an isomorphism.

Definition A.5. A category is *complete* (resp. *cocomplete*) if it has all small limits (resp. all small colimits).

Definition A.6. Two functors $L : \mathcal{C} \longrightarrow \mathcal{D}$ and $R : \mathcal{D} \longrightarrow \mathcal{C}$ are *adjoint* (denoted $L \dashv R$, L is the *left adjoint* and R is the *right adjoint*) if the functors $\mathrm{Hom}_{\mathcal{D}}(L(-), -)$ and $\mathrm{Hom}_{\mathcal{C}}(-, R(-))$ are naturally isomorphic. Under the isomorphism $\mathrm{Hom}_{\mathcal{D}}(L(c), L(c)) \cong \mathrm{Hom}_{\mathcal{C}}(c, R(L(c)))$, the image of $\mathrm{id}_{L(c)}$ is called the *unit* of the adjunction and is denoted $\eta_c : c \longrightarrow RLc$. Dually, under the isomorphism $\mathrm{Hom}_{\mathcal{C}}(R(d), R(d)) \cong \mathrm{Hom}_{\mathcal{D}}(L(R(d)), d)$, the image of $\mathrm{id}_{R(d)}$ is called the *counit* and is denoted $\epsilon_d : LRd \longrightarrow d$. These constructions give natural transformations $\eta : \mathrm{id}_{\mathcal{C}} \longrightarrow RL$ and $\epsilon : LR \longrightarrow \mathrm{id}_{\mathcal{D}}$.

Definition A.7. A category $\mathcal{C}$ is *filtered* if every finite diagram has a cocone. Equivalently, if :

- there is an object of $\mathcal{C}$;
- for every objects c_1 and c_2 there are morphisms to a common object $c_1 \longrightarrow c_3$ and $c_2 \longrightarrow c_3$;
- for every pair of parallel arrows $f, g : c_1 \longrightarrow c_2$ there is an arrow $h : c_2 \longrightarrow c_3$ coequalizing f and g ($hf = hg$).

Definition A.8. A *cartesian closed category* is a category with finite products and exponentials (that is, each functor $- \times X$ has a right adjoint $(-)^X$).

Definition A.9. In a category $\mathcal{C}$, a *subobject* of an object X is an equivalence class of monomorphisms $A \rightarrowtail X$, where A and B are equivalent when they are isomorphic in $\mathcal{C}/X$. Since there is at most one morphism in $\mathcal{C}/X$ between two such monomorphisms, the category of subobjects $\mathrm{Sub}_{\mathcal{C}}(X)$ is a (*a priori* large) poset.

Definition A.10. In a category $\mathcal{C}$ with finite limits, a *subobject classifier* is a monomorphism $\mathrm{true} : 1 \rightarrowtail \Omega$ from the terminal object such that for every monomorphism $A \rightarrowtail X$ in $\mathcal{C}$ there is a unique characteristic morphism $\chi_A : X \longrightarrow \Omega$ such that the corresponding square is a pullback :

$$\begin{array}{ccc} A & \xrightarrow{\ !\ } & 1 \\ \downarrow & \lrcorner & \downarrow{\scriptstyle \mathrm{true}} \\ X & \xrightarrow[\chi_A]{} & \Omega \end{array}$$

Definition A.11. The *kernel pair* of an arrow $f : X \longrightarrow Y$ in a category with pullbacks is the pullback of f along itself :

$$\begin{array}{ccc} X \times_Y X & \longrightarrow & X \\ \big\downarrow & \lrcorner & \big\downarrow f \\ Y & \xrightarrow[f]{} & Y \end{array}$$

The corresponding *diagonal* arrow is the arrow $\Delta : X \longrightarrow X \times_Y X$ given by the universal property of pullbacks applied to the arrows id_X and id_X. It is easy to check that Δ is an isomorphism if and only if f is a monomorphism.

Definition A.12. In a category, an object A is a *retract* of B if there are morphisms $i : A \longrightarrow B$ and $r : B \longrightarrow A$ such that $ri = \mathrm{id}_A$. A morphism e is *idempotent* if $e \circ e = e$. Clearly, if r is a retraction of a morphism i, then ir is idempotent. A *split idempotent* is an idempotent of this form.

Let $\mathcal{C}$ be a small category. Its *Cauchy completion* $\hat{\mathcal{C}}$ is the full subcategory of $\mathbf{Psh}(\mathcal{C})$ on the retracts of representable presheaves. A category is *Cauchy-complete* if every idempotent splits.

A.2 Results

Lemma A.13 (Yoneda's lemma)**.** *Let $\mathcal{C}$ be a locally small category and $\mathbf{y}_{\mathcal{C}} : \mathcal{C} \longrightarrow \mathbf{Psh}(\mathcal{C}) = [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}]$ the functor sending an object c to $\mathrm{Hom}_{\mathcal{C}}(-, c)$. This is the image under the adjunction $\mathrm{Hom}_{\mathfrak{CAT}}(\mathcal{C}^{\mathrm{op}} \times \mathcal{C}, \mathbf{Set}) \simeq \mathrm{Hom}_{\mathfrak{CAT}}(\mathcal{C}, [\mathcal{C}^{\mathrm{op}}, \mathbf{Set}])$ of the functor $\mathrm{Hom}_{\mathcal{C}}$. Then for a presheaf $X \in \mathbf{Psh}(\mathcal{C})$, there is a canonical isomorphism :*

$$\mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(\mathbf{y}_{\mathcal{C}}(c), X) \cong X(c).$$

Proof. A natural transformation $\eta \in \mathrm{Hom}_{\mathcal{C}}(-, c) \longrightarrow X$ is uniquely determined by the value $\xi = \eta_c(\mathrm{id}_c) \in X(c)$. Indeed, for any object b in $\mathcal{C}$, the naturality of η implies that η_b must send an element $f \in \mathrm{Hom}_{\mathcal{C}}(b, c)$ to $X(f)(\xi) \in X(b)$. $\square$

Corollary A.14. *The* Yoneda embedding $\mathbf{y}_{\mathcal{C}}$ *is full and faithful, since it induces the following isomorphism for objects c and d :*

$$\mathrm{Hom}_{\mathbf{Psh}(\mathcal{C})}(\mathrm{Hom}_{\mathcal{C}}(-, c), \mathrm{Hom}_{\mathcal{C}}(-, d)) \cong (\mathrm{Hom}_{\mathcal{C}}(-, d))(c) = \mathrm{Hom}_{\mathcal{C}}(c, d).$$

In particular, if $\mathbf{y}_{\mathcal{C}}(c) \cong \mathbf{y}_{\mathcal{C}}(d)$ then this natural isomorphism comes from an isomorphism $c \cong d$.

Proposition A.15. *Two functors $L : \mathcal{C} \longrightarrow \mathcal{D}$ and and $R : \mathcal{D} \longrightarrow \mathcal{C}$ are adjoint if and only if there exist natural transformations $\eta : \mathrm{id}_{\mathcal{C}} \longrightarrow RL$ and $\epsilon : LR \longrightarrow \mathrm{id}_{\mathcal{D}}$ such that the following two triangles commute :*

$$\begin{array}{ccc} L & \xrightarrow{L\eta} & LRL \\ & \searrow^{\mathrm{id}_L} & \big\downarrow \epsilon L \\ & & L \end{array} \qquad \begin{array}{ccc} R & & \\ \big\downarrow \eta R & \searrow^{\mathrm{id}_R} & \\ RLR & \xrightarrow{R\epsilon} & R \end{array}$$

Proof. Given $f : Lc \longrightarrow d$, the corresponding morphism $f^\sharp : c \longrightarrow Rd$ is the composite $R(f) \circ \eta_c$. In the other way, given $g : c \longrightarrow Rd$ the corresponding morphism $g^\flat : Lc \longrightarrow d$ is given by $\epsilon_d \circ L(g)$. The isomorphism $f \mapsto f^\sharp$ follows from :

$$\begin{aligned} f^{\sharp\flat} &= (c \xrightarrow{\eta_c} RLc \xrightarrow{R(f)} Rd)^\flat \\ &= Lc \xrightarrow{L(\eta_c)} LRLc \xrightarrow{LR(f)} LRd \xrightarrow{\epsilon_d} d \\ &= Lc \xrightarrow{L(\eta_c)} LRLc \xrightarrow{\epsilon_{Lc}} Lc \xrightarrow{f} d \\ &= Lc \xrightarrow{f} d. \end{aligned}$$

The naturality of this isomorphism comes from the naturality of η and ϵ. □

Theorem A.16. *Let $L : \mathcal{C} \longrightarrow \mathcal{D} : R$ be functors with $L \dashv R$. Then L preserves all colimits and R preserves all limits.*

Proof. We only show this for R, as the fact for L is dual. Let Y be an object of $\mathcal{D}$ and $X : J \longrightarrow \mathcal{C}$ be a diagram whose limit exists in $\mathcal{C}$. Since the Hom functors preserve limits in the second argument, we have :

$$\begin{aligned} \mathrm{Hom}_{\mathcal{D}}(Y, R(\lim_{j\in J} X_j)) &\cong \mathrm{Hom}_{\mathcal{C}}(LY, \lim_{j\in J} X_j) \\ &\cong \lim_{j\in J} \mathrm{Hom}_{\mathcal{C}}(LY, X_j) \\ &\cong \lim_{j\in J} \mathrm{Hom}_{\mathcal{D}}(Y, RX_j) \\ &\cong \mathrm{Hom}_{\mathcal{D}}(Y, \lim_{j\in J} RX_j), \end{aligned}$$

naturally in Y. The Yoneda lemma concludes that $R(\lim_{j\in J} X_j) \cong \lim_{j\in J} RX_j$. □

Converses to this theorem are called adjoint functor theorems, stating that under assumptions on the categories or the functors involved, them preserving (co)limits suffices to ensure that they have adjoints. A proof of the adjoint functor theorem used in remark 3.10 would be too long for its relevance.

Theorem A.17 (Finite limits commute with filtered colimits in **Set**)**.** *Let I be a finite category and J a small filtered category. Then for any functor $F : I \times J \longrightarrow$ **Set**, the canonical arrow*

$$\kappa : \mathrm{colim}_{j\in J} \lim_{i\in I} F(i,j) \longrightarrow \lim_{i\in I} \mathrm{colim}_{j\in J} F(i,j)$$

is an isomorphism.

Proof. Notice that

$$\mathrm{colim}_{j\in J} F(i,j) = \coprod_{j\in J} F(i,j)/R,$$

R being the equivalence relation where $x \in F(i,j)$ and $x' \in F(i,j')$ are equivalent when there are $u : j \longrightarrow k$ and $u' : j' \longrightarrow k$ with $F(i,u)(x) = F(i,u')(x')$. Any finite family of such equivalence classes $((x_m, j_m))$ (with $x_m \in F(i, j_m)$) can be rewritten as $((y_m, k))$ with a common k, since J is filtered. For the same reason, if two items (y,k) and (y',k) of this list are equivalent then there is $w : k \longrightarrow k'$ such that $F(i,w)(y) = F(i,w)(y')$.

For any functor $G : I \longrightarrow$ **Set**, $\lim_{i\in I} G(i) = \mathrm{Hom}_{[I,\mathbf{Set}]}(*, G)$ is the set of cones on G. Letting $G(i) = \mathrm{colim}_{j\in J} F(i,j)$, since I is finite every cone is a finite family of elements of $\mathrm{colim}_{j\in J} F(i,j)$ satisfying a finite number of equations. According to the remarks above, each cone τ is thus a family $(\tau_i = (y_i, k'))_{i\in I}$ for some common k' where the y_i constitute a cone $y : * \longrightarrow F(-,k')$. Since the equivalence class of y is an element of $\mathrm{colim}_{j\in J} \lim_{i\in I} F(i,j)$, we get a map

$$\begin{aligned} \lim_{i\in I} \mathrm{colim}_{j\in J} F(i,j) &\longrightarrow \mathrm{colim}_{j\in J} \lim_{i\in I} F(i,j) \\ \tau &\longmapsto (y, k') \end{aligned}$$

which can be checked to be the inverse of κ. ☐

Proposition A.18. *The following conditions are equivalent for an idempotent* $e : c \longrightarrow c$ *:*

(i) e *splits as* $i \circ r$ *;*
(ii) the equalizer i *of* e *and* id_c *exists ;*
(iii) the coequalizer r *of* e *and* id_c *exists.*

Theorem A.19. *Let* $\mathcal{C}$ *be a small category and* $\hat{\mathcal{C}}$ *its Cauchy completion. Then*

(i) $\hat{\mathcal{C}}$ *is small and* $\mathcal{C}$ *is a full subcategory of* $\hat{\mathcal{C}}$ *;*
(ii) $\hat{\mathcal{C}}$ *is Cauchy-complete ;*
(iii) the inclusion $\mathcal{C} \rightarrowtail \hat{\mathcal{C}}$ *is an equivalence if and only if* $\mathcal{C}$ *is Cauchy-complete ;*
(iv) there is an equivalence $\mathbf{Psh}(\mathcal{C}) \simeq \mathbf{Psh}(\hat{\mathcal{C}})$*.*

Proof. $\hat{C}$ is small because $\mathbf{Psh}(\mathcal{C})$ is well-powered, and $\mathcal{C}$ is a full subcategory because $\mathbf{y}_{\mathcal{C}}$ is full and faithful. Every idempotent in $\hat{\mathcal{C}}$ splits because it splits in $\mathbf{Psh}(\mathcal{C})$ and the composite of two retractions is a retraction.

A retract of a representable presheaf $\mathbf{y}_{\mathcal{C}}(c)$ induces an idempotent on $\mathbf{y}_{\mathcal{C}}(c)$ and by the Yoneda lemma an idempotent on c. If $\mathcal{C}$ is Cauchy-complete, this idempotent splits and already produces a retraction of $\mathbf{y}_{\mathcal{C}}(c)$. So if $\mathcal{C}$ is Cauchy-complete then $\mathcal{C} \simeq \hat{C}$.

To prove the equivalence between the presheaf categories, it is enough to show that a presheaf $F \in \mathbf{Psh}(\mathcal{C})$ extends uniquely to a presheaf $\hat{F} \in \mathbf{Psh}(\hat{\mathcal{C}})$. Since by proposition A.18 the category **Set** is already Cauchy-complete, $\hat{F}$ has to map the splitting of an idempotent e to the splitting of Fe. If $i : R \rightleftarrows \mathbf{y}_{\mathcal{C}}(c) : r$ and $j : S \rightleftarrows \mathbf{y}_{\mathcal{C}}(d)$ are retractions, every morphism $f : R \longrightarrow S$ is equal to $sjfri$. Since $\hat{F}(i)$ and $\hat{F}(s)$ are already defined and since $\hat{F}(jfr)$ has to be $F(jfr)$, this determines uniquely the presheaf $\hat{F}$. ☐

Bibliography

[1] Michael ARTIN, Alexandre GROTHENDIECK, Jean-Louis VERDIER, *Théorie des topos et cohomologie étale des schémas (SGA4)*, Springer-Verlag, 1972.

[2] Michael BARR, *Toposes without points*, Springer-Verlag, 1971.

[3] Ingo BLECHSCHMIDT, *Using the internal language of toposes in algebraic geometry*, `https://arxiv.org/abs/2111.03685`, 2017.

[4] Francis BORCEUX, Dominique DEJEAN, *Cauchy completion in category theory*, Cahiers de topologie et géométrie différentielle catégoriques (tome 27), 1986.

[5] Olivia CARAMELLO, *Site characterizations for geometric invariants of toposes*, `https://arxiv.org/abs/1112.2542`, 2011.

[6] Olivia CARAMELLO, *Theories, Sites, Toposes*, Oxford University Press, 2018.

[7] Monique HAKIM, *Topos annelés et schémas relatifs*, Springer-Verlag, 1972.

[8] Matthias HUTZLER, *Internal language and classified theories of toposes in algebraic geometry*, 2018.

[9] Matthias HUTZLER *Syntactic presentations for glued toposes and for crystalline toposes*, `https://arxiv.org/abs/2206.11244`, 2022.

[10] Peter T. JOHNSTONE, *Sketches of an Elephant : A Topos Theory Compendium, Volume 2*, Oxford Logic Guides, 2002.

[11] Saunders MAC LANE, *Categories for the Working Mathematician*, Springer-Verlag, 1971.

[12] Saunders MAC LANE, Ieke MOERDIJK, *Sheaves in Geometry and Logic*, Springer-Verlag, 1992.

[13] Conor MAHANY, *Elementary Topos Theory and Intuitionistic Logic*, `https://math.uchicago.edu/~may/REU2012/REUPapers/Mahany.pdf`, 2012.

[14] The Stacks Project authors, *The Stacks Project*, `https://stacks.math.columbia.edu/`.

[15] Ravi VAKIL, *The Rising Sea*, `http://math.stanford.edu/~vakil/216blog/`, 2017.

[16] Gavin C. WRAITH, *Generic Galois theory of local rings* (from the symposium *Applications of Sheaves*), Springer-Verlag 1977.