THE LARGEST GROTHENDIECK TOPOLOGY FOR WHICH A PRESHEAF IS A SHEAF

YU MURAKAMI
NEW YORK GENERAL GROUP
JULY 3, 2025

Abstract. Let A be a small category and let $F : A^{\mathrm{op}} \to$ **Set** be a presheaf. We give a direct construction of the largest Grothendieck topology J_F on A for which F satisfies the sheaf condition. A sieve S on a is declared J_F-covering precisely when every pullback of S is F-effective, i.e. when the canonical Yoneda restriction map

$$F(b) \longrightarrow \mathrm{Nat}(h^*S, F)$$

is bijective for every arrow $h : b \to a$. We prove directly that this assignment is a Grothendieck topology and that it is maximal among all topologies for which F is a sheaf. The construction is a single-presheaf instance of a standard principle implicit in the theory of sites and sheaves; the contribution of the paper is to isolate the stable effective-sieve construction and give a self-contained proof.

1. Introduction

Let A be a small category. A Grothendieck topology on A specifies which sieves are to be regarded as covering. Given such a topology J, a presheaf $F : A^{\mathrm{op}} \to$ **Set** is a J-sheaf when every J-covering sieve admits unique amalgamation of matching families.

This paper studies the inverse question: given a presheaf $F : A^{\mathrm{op}} \to$ **Set**, what is the largest Grothendieck topology on A for which F is a sheaf?

The answer is obtained by considering not merely those sieves S for which the sheaf comparison map

$$F(a) \longrightarrow \mathrm{Nat}(S, F)$$

is bijective, but those sieves for which this property is stable under every pullback. This stability condition is essential. The class of sieves for which the above map is bijective need not itself be pullback-stable, and hence need not form a Grothendieck topology. The correct construction is therefore

$$J_F(a) = \left\{ S \subseteq ya \;\middle|\; \text{for every } h : b \to a,\ F(b) \xrightarrow{\sim} \mathrm{Nat}(h^*S, F) \right\}.$$

The Yoneda lemma enters through the identification

$$F(a) \cong \mathrm{Nat}(ya, F),$$

so that the sheaf condition becomes the assertion that restriction along $S \hookrightarrow ya$ is bijective.

The result belongs to the standard circle of ideas around sites, canonical topologies, and sheafification; see Mac Lane–Moerdijk [1], SGA 4 [2], and Johnstone [3]. The purpose here is not to claim a new sheafification theorem, but to give a precise, self-contained construction for a single presheaf and prove its universal property directly.

2. Sieves, Matching Families, and Yoneda Restriction

Let A be a small category. Write

$$\widehat{A} = [A^{\mathrm{op}}, \mathbf{Set}]$$

for the category of presheaves on A.

For $a \in A$, write

$$ya = A(-, a)$$

for the representable presheaf.

A *sieve* S on a is a subpresheaf

$$S \hookrightarrow ya.$$

Equivalently, S is a class of arrows $u : b \to a$, closed under precomposition: if $u : b \to a$ belongs to S and $v : c \to b$ is any arrow, then $uv : c \to a$ belongs to S.

If $f : b \to a$ is an arrow and S is a sieve on a, the pullback sieve f^*S on b is

$$f^*S = \{g : c \to b \mid fg \in S\}.$$

For a presheaf $F : A^{\mathrm{op}} \to \mathbf{Set}$, an F-matching family on a sieve $S \hookrightarrow ya$ is a natural transformation

$$S \to F.$$

Thus the set of matching families is

$$\mathrm{Nat}(S, F).$$

By the Yoneda lemma [4], there is a natural bijection

$$F(a) \cong \mathrm{Nat}(ya, F).$$

Restriction along the inclusion $S \hookrightarrow ya$ gives the canonical map

$$\rho^F_S : F(a) \longrightarrow \mathrm{Nat}(S, F).$$

Explicitly, for $x \in F(a)$, $\rho^F_S(x)$ is the matching family sending an arrow $u : b \to a$ in S to

$$F(u)(x) \in F(b).$$

A presheaf F is a sheaf for a Grothendieck topology J if and only if ρ^F_S is bijective for every J-covering sieve S.

3. Stable F-Effective Sieves

Definition 3.1. Let $F : A^{\mathrm{op}} \to \mathbf{Set}$ be a presheaf. A sieve S on a is called *stably F-effective* if, for every arrow $h : b \to a$, the restriction map

$$\rho^F_{h^*S} : F(b) \longrightarrow \mathrm{Nat}(h^*S, F)$$

is bijective.

Define

$$J_F(a) = \{S \text{ sieve on } a \mid S \text{ is stably } F\text{-effective}\}.$$

The aim is to prove that J_F is a Grothendieck topology. The stability in the definition is crucial. A sieve S may satisfy that ρ^F_S is bijective without its pullbacks satisfying the same property. Therefore one cannot define J_F merely by asking that ρ^F_S itself be bijective.

4. A Local Transitivity Lemma

We need one elementary lemma. It is the precise form of the usual transitivity argument for gluing matching families.

Lemma 4.1. *Let $F : A^{\mathrm{op}} \to \mathbf{Set}$ be a presheaf. Let R and S be sieves on an object a. Suppose:*

(1) $\rho^F_R : F(a) \to \mathrm{Nat}(R, F)$ *is bijective;*
(2) *for every arrow $f : b \to a$ in R, the map*

$$\rho^F_{f^*S} : F(b) \to \mathrm{Nat}(f^*S, F)$$

is bijective;

(3) *for every arrow $u : c \to a$ in S, the map*

$$\rho^F_{u^*R} : F(c) \to \mathrm{Nat}(u^*R, F)$$

is injective.

Then

$$\rho^F_S : F(a) \longrightarrow \mathrm{Nat}(S, F)$$

is bijective.

Proof. We prove injectivity and surjectivity separately.

First suppose $x, x' \in F(a)$ and

$$\rho^F_S(x) = \rho^F_S(x').$$

We must show $x = x'$.

Because ρ^F_R is injective, it suffices to prove that the restrictions of x and x' to R agree. Let $f : b \to a$ be an arrow in R. We need to show

$$F(f)(x) = F(f)(x').$$

By hypothesis, $\rho^F_{f^*S}$ is injective. Hence it is enough to show that $F(f)(x)$ and $F(f)(x')$ have the same restriction to f^*S.

Let $g : c \to b$ belong to f^*S. Then $fg \in S$. Since x and x' agree on S,

$$F(fg)(x) = F(fg)(x').$$

Equivalently,

$$F(g)F(f)(x) = F(g)F(f)(x').$$

Thus $F(f)(x)$ and $F(f)(x')$ agree on f^*S, so

$$F(f)(x) = F(f)(x').$$

This holds for every $f \in R$, and therefore $x = x'$. Thus ρ^F_S is injective.

Now let

$$\alpha : S \to F$$

be a matching family on S. We must construct a unique $x \in F(a)$ whose restriction to S is α.

For every arrow $f : b \to a$ in R, define a matching family

$$\alpha_f : f^*S \to F$$

by

$$\alpha_f(g) = \alpha(fg)$$

for every $g : c \to b$ with $fg \in S$. This is natural because α is natural.

By hypothesis, $\rho^F_{f^*S}$ is bijective. Hence there is a unique element

$$x_f \in F(b)$$

whose restriction to f^*S is α_f.

We claim that the assignment

$$f \longmapsto x_f$$

is a matching family on R. Let $f : b \to a$ be in R, and let $h : c \to b$ be any arrow. Since R is a sieve, $fh : c \to a$ also lies in R. We need to prove

$$F(h)(x_f) = x_{fh}.$$

Because $\rho^F_{(fh)^*S}$ is injective, it is enough to compare the restrictions of these two elements to $(fh)^*S$.

Let $k : d \to c$ belong to $(fh)^*S$. Then $fhk \in S$. We compute:

$$F(k)F(h)(x_f) = F(hk)(x_f) = \alpha_f(hk) = \alpha(fhk),$$

while

$$F(k)(x_{fh}) = \alpha_{fh}(k) = \alpha(fhk).$$

Therefore

$$F(k)F(h)(x_f) = F(k)(x_{fh})$$

for every $k \in (fh)^*S$. Hence

$$F(h)(x_f) = x_{fh}.$$

Thus $f \mapsto x_f$ is a natural transformation

$$R \to F.$$

Since ρ_R^F is bijective, there exists a unique element

$$x \in F(a)$$

such that

$$F(f)(x) = x_f$$

for every $f : b \to a$ in R.

It remains to prove that x restricts to α on S. Let $u : c \to a$ be an arrow in S. We need to show

$$F(u)(x) = \alpha(u).$$

By hypothesis, $\rho_{u^*R}^F$ is injective. Hence it is enough to prove that the two elements $F(u)(x)$ and $\alpha(u)$ of $F(c)$ have the same restriction to u^*R.

Let $v : d \to c$ belong to u^*R. Then $uv \in R$. Since $u \in S$ and S is a sieve, $uv \in S$ as well.

We compute:

$$F(v)F(u)(x) = F(uv)(x) = x_{uv}.$$

We claim that

$$x_{uv} = \alpha(uv).$$

Indeed, x_{uv} is defined as the unique element of $F(d)$ whose restriction to $(uv)^*S$ is α_{uv}. Since $uv \in S$, the sieve $(uv)^*S$ is the maximal sieve on d: for every $k : e \to d$, the composite uvk lies in S. The element $\alpha(uv) \in F(d)$ restricts to the same matching family, because for every $k : e \to d$,

$$F(k)(\alpha(uv)) = \alpha(uvk) = \alpha_{uv}(k).$$

By uniqueness,

$$x_{uv} = \alpha(uv).$$

Therefore

$$F(v)F(u)(x) = \alpha(uv).$$

On the other hand, by naturality of α,

$$F(v)(\alpha(u)) = \alpha(uv).$$

Hence

$$F(v)F(u)(x) = F(v)(\alpha(u))$$

for every $v \in u^*R$. Since $\rho_{u^*R}^F$ is injective, we conclude that

$$F(u)(x) = \alpha(u).$$

This holds for every $u \in S$, so x restricts to α.

Thus ρ_S^F is surjective. Since injectivity was already proved, ρ_S^F is bijective. □

5. The Topology J_F

We now prove the main theorem.

Theorem 5.1. *For every presheaf $F : A^{\mathrm{op}} \to \mathbf{Set}$, the assignment*

$$a \longmapsto J_F(a)$$

is a Grothendieck topology on A.

Proof. We verify the three Grothendieck topology axioms.

Maximal sieve. Let $a \in A$. We show that the maximal sieve ya belongs to $J_F(a)$.

Let $h : b \to a$ be any arrow. Then

$$h^*(ya) = yb.$$

Therefore

$$\rho^F_{h^*(ya)} = \rho^F_{yb} : F(b) \to \mathrm{Nat}(yb, F)$$

is the Yoneda bijection. Hence ya is stably F-effective, so

$$ya \in J_F(a).$$

Pullback stability. Let $S \in J_F(a)$, and let $f : b \to a$ be an arrow. We show that

$$f^*S \in J_F(b).$$

Let $h : c \to b$ be any arrow. Then

$$h^*(f^*S) = (fh)^*S.$$

Since $S \in J_F(a)$, the map

$$\rho^F_{(fh)^*S} : F(c) \to \mathrm{Nat}((fh)^*S, F)$$

is bijective. Therefore $\rho^F_{h^*(f^*S)}$ is bijective for every $h : c \to b$. Hence

$$f^*S \in J_F(b).$$

Transitivity. Let $R \in J_F(a)$. Let S be a sieve on a. Suppose that for every arrow $f : b \to a$ in R,

$$f^*S \in J_F(b).$$

We must prove that

$$S \in J_F(a).$$

Thus let $h : c \to a$ be any arrow. We need to show that

$$\rho^F_{h^*S} : F(c) \to \mathrm{Nat}(h^*S, F)$$

is bijective.

We apply Lemma 4.1 inside the object c, with the two sieves

$$h^*R \quad \text{and} \quad h^*S$$

on c.

First, since $R \in J_F(a)$, the map

$$\rho^F_{h^*R} : F(c) \to \mathrm{Nat}(h^*R, F)$$

is bijective.

Second, let $g : d \to c$ be an arrow in h^*R. Then $hg : d \to a$ lies in R. By the transitivity hypothesis,

$$(hg)^*S \in J_F(d).$$

In particular, taking the identity arrow on d, the map

$$\rho^F_{(hg)^*S} : F(d) \to \mathrm{Nat}((hg)^*S, F)$$

is bijective. But

$$g^*(h^*S) = (hg)^*S.$$

Thus $\rho^F_{g^*(h^*S)}$ is bijective for every $g \in h^*R$.

Third, let $u : e \to c$ be an arrow in h^*S. Since $R \in J_F(a)$, the map

$$\rho^F_{(hu)^*R} : F(e) \to \mathrm{Nat}((hu)^*R, F)$$

is bijective, hence injective. But

$$u^*(h^*R) = (hu)^*R.$$

Therefore $\rho^F_{u^*(h^*R)}$ is injective for every $u \in h^*S$.

All hypotheses of Lemma 4.1 are satisfied for the pair of sieves h^*R and h^*S. Hence

$$\rho^F_{h^*S} : F(c) \to \mathrm{Nat}(h^*S, F)$$

is bijective.

Since this holds for every $h : c \to a$, the sieve S is stably F-effective. Therefore

$$S \in J_F(a).$$

The three axioms are verified, so J_F is a Grothendieck topology. □

6. Maximality

Theorem 6.1. *Let $F : A^{\mathrm{op}} \to$ **Set** be a presheaf. Then J_F is the largest Grothendieck topology on A for which F is a sheaf. Equivalently, for every Grothendieck topology J on A,*

$$F \in \mathrm{Sh}(A, J) \quad \Longleftrightarrow \quad J \leq J_F.$$

Proof. First we show that F is a J_F-sheaf.

Let $S \in J_F(a)$. Since S is stably F-effective, we may take $h = \mathrm{id}_a$. Then

$$\mathrm{id}_a^* S = S,$$

and therefore

$$\rho^F_S : F(a) \to \mathrm{Nat}(S, F)$$

is bijective. Hence F satisfies the sheaf condition for every J_F-covering sieve. Thus

$$F \in \mathrm{Sh}(A, J_F).$$

Now let J be any Grothendieck topology on A such that F is a J-sheaf. We prove that

$$J \leq J_F.$$

Let $S \in J(a)$. To show $S \in J_F(a)$, let $h : b \to a$ be any arrow. Since J is a Grothendieck topology, it is stable under pullback, so

$$h^*S \in J(b).$$

Since F is a J-sheaf, the map

$$\rho^F_{h^*S} : F(b) \to \mathrm{Nat}(h^*S, F)$$

is bijective. This holds for every $h : b \to a$. Hence S is stably F-effective, so

$$S \in J_F(a).$$

Therefore every J-covering sieve is J_F-covering, and so

$$J \leq J_F.$$

Combining the two directions gives the claimed equivalence. Thus J_F is the largest Grothendieck topology for which F is a sheaf. □

7. Families of Presheaves

The single-presheaf construction extends without change to a set of presheaves.

Definition 7.1. Let $\mathcal{F}$ be a set of presheaves on A. Define $J_{\mathcal{F}}(a)$ to be the class of sieves S on a such that, for every $F \in \mathcal{F}$ and every arrow $h : b \to a$, the map

$$\rho^F_{h^*S} : F(b) \to \mathrm{Nat}(h^*S, F)$$

is bijective.

Proposition 7.2. *The assignment $J_{\mathcal{F}}$ is the largest Grothendieck topology on A for which every $F \in \mathcal{F}$ is a sheaf.*

Proof. The proof is the same as the proof of Theorems 5.1 and 6.1, applied simultaneously to all $F \in \mathcal{F}$.

For the topology axioms, maximal sieves are covering because the Yoneda lemma gives

$$F(b) \cong \mathrm{Nat}(yb, F)$$

for every $F \in \mathcal{F}$. Pullback stability is immediate from

$$h^*(f^*S) = (fh)^*S.$$

For transitivity, apply Lemma 4.1 separately to each $F \in \mathcal{F}$.

Thus $J_{\mathcal{F}}$ is a Grothendieck topology.

If $S \in J_{\mathcal{F}}(a)$, then for every $F \in \mathcal{F}$, taking $h = \mathrm{id}_a$ shows that

$$\rho^F_S : F(a) \to \mathrm{Nat}(S, F)$$

is bijective. Hence every $F \in \mathcal{F}$ is a $J_{\mathcal{F}}$-sheaf.

Conversely, if J is a Grothendieck topology for which every $F \in \mathcal{F}$ is a J-sheaf, and if $S \in J(a)$, then for every $h : b \to a$, the pullback sieve h^*S belongs to $J(b)$. Hence every $F \in \mathcal{F}$ satisfies

$$F(b) \xrightarrow{\sim} \mathrm{Nat}(h^*S, F).$$

Therefore $S \in J_{\mathcal{F}}(a)$. Thus $J \le J_{\mathcal{F}}$, so $J_{\mathcal{F}}$ is maximal. □

8. Relation with the Canonical Topology

Recall that the canonical topology on A is the largest Grothendieck topology for which every representable presheaf is a sheaf [1, 3].

Let

$$\mathcal{Y} = \{ya \mid a \in A\}$$

be the set of representable presheaves. By the preceding proposition, $J_{\mathcal{Y}}$ is the largest topology for which every representable is a sheaf. Therefore

$$J_{\mathcal{Y}} = J_{\mathrm{can}}.$$

Equivalently,

$$J_{\mathrm{can}} = \bigwedge_{a \in A} J_{ya},$$

where the meet is computed by intersection of covering-sieve assignments.

Indeed, a topology J is below every J_{ya} if and only if every representable ya is a J-sheaf. This is precisely the defining property of the canonical topology.

9. Examples

9.1. **Discrete categories.** Let A be a discrete category. For each object a, the only arrows into a are identities. Therefore the sieves on a are

$$\emptyset \quad \text{and} \quad ya.$$

Let $F : A^{\mathrm{op}} \to \mathbf{Set}$ be a presheaf. Since there are no nontrivial arrows, F is just a family of sets $(F(a))_{a\in A}$.

The maximal sieve ya is always J_F-covering. The empty sieve $\emptyset$ is J_F-covering on a precisely when

$$F(a) \to \mathrm{Nat}(\emptyset, F)$$

is bijective. But $\mathrm{Nat}(\emptyset, F)$ is a singleton. Hence $\emptyset$ covers a exactly when $F(a)$ is a singleton.

Thus, for a discrete category,

$$J_F(a) = \begin{cases} \{ya, \emptyset\}, & F(a) \text{ is a singleton,} \\ \{ya\}, & \text{otherwise.} \end{cases}$$

9.2. **Posets.** Let P be a poset regarded as a category: there is a unique arrow $q \to p$ exactly when $q \leq p$. A sieve on p is a down-closed subset

$$S \subseteq \downarrow p.$$

A presheaf $F : P^{\mathrm{op}} \to \mathbf{Set}$ assigns a set $F(p)$ to each p and restriction maps

$$F(p) \to F(q)$$

whenever $q \leq p$.

For $q \leq p$, the pullback sieve q^*S is

$$S \cap \downarrow q.$$

Therefore S is J_F-covering exactly when, for every $q \leq p$,

$$F(q) \longrightarrow \lim_{r \in S\cap\downarrow q} F(r)$$

is bijective, where the limit is taken over the full downset $S \cap \downarrow q$.

Thus in a poset, J_F consists exactly of the downsets along which F is locally recoverable, after every restriction $q \leq p$, as the limit of its values on the smaller elements lying in the sieve.

10. Behaviour Under Isomorphisms and Retracts

Proposition 10.1. *If $F \cong G$ in $\widehat{A}$, then*

$$J_F = J_G.$$

Proof. Let $\theta : F \to G$ be an isomorphism. For every sieve S on a and every arrow $h : b \to a$, naturality gives a commutative square

$$\begin{array}{ccc} F(b) & \xrightarrow{\rho^F_{h^*S}} & \mathrm{Nat}(h^*S, F) \\ \downarrow \theta_b & & \downarrow \mathrm{Nat}(h^*S, \theta) \\ G(b) & \xrightarrow{\rho^G_{h^*S}} & \mathrm{Nat}(h^*S, G). \end{array}$$

The vertical maps are bijections. Hence $\rho^F_{h^*S}$ is bijective if and only if $\rho^G_{h^*S}$ is bijective.

Therefore the stably F-effective sieves are exactly the stably G-effective sieves. Thus $J_F = J_G$. $\square$

Proposition 10.2. *Suppose G is a retract of F in $\widehat{A}$. That is, suppose there are natural transformations*

$$G \xrightarrow{s} F \xrightarrow{p} G$$

such that

$$ps = \mathrm{id}_G.$$

Then

$$J_F \leq J_G.$$

Proof. Let $S \in J_F(a)$. We show that $S \in J_G(a)$.

Let $h : b \to a$ be any arrow. Since $S \in J_F(a)$, the map

$$\rho^F_{h^*S} : F(b) \to \mathrm{Nat}(h^*S, F)$$

is bijective.

We prove that

$$\rho^G_{h^*S} : G(b) \to \mathrm{Nat}(h^*S, G)$$

is bijective.

First, injectivity. Let $x, x' \in G(b)$ have the same restriction to h^*S. Applying s_b, the elements $s_b(x)$ and $s_b(x')$ of $F(b)$ also have the same restriction to h^*S. Since $\rho^F_{h^*S}$ is injective,

$$s_b(x) = s_b(x').$$

Applying p_b, and using $p_b s_b = \mathrm{id}_{G(b)}$, we obtain

$$x = x'.$$

Thus $\rho^G_{h^*S}$ is injective.

Now let

$$\alpha : h^*S \to G$$

be a matching family. Compose with s to obtain

$$s\alpha : h^*S \to F.$$

Since $\rho^F_{h^*S}$ is surjective, there exists $y \in F(b)$ whose restriction to h^*S is $s\alpha$.

Set

$$x = p_b(y) \in G(b).$$

For every arrow $k : c \to b$ in h^*S,

$$G(k)(x) = G(k)p_b(y) = p_c F(k)(y) = p_c s_c(\alpha(k)) = \alpha(k).$$

Therefore x restricts to α. Hence $\rho^G_{h^*S}$ is surjective.

Thus $\rho^G_{h^*S}$ is bijective for every $h : b \to a$, and therefore $S \in J_G(a)$. Hence $J_F \leq J_G$. □

There is no corresponding monotonicity statement for arbitrary morphisms $F \to G$. The retract hypothesis is used essentially to transfer both uniqueness and existence of amalgamations from F to G.

11. Conclusion

For a presheaf $F : A^{\mathrm{op}} \to \mathbf{Set}$, the largest Grothendieck topology for which F is a sheaf is obtained by declaring a sieve S on a to cover precisely when all its pullbacks are F-effective:

$$S \in J_F(a) \quad \Longleftrightarrow \quad \forall h : b \to a,\ F(b) \xrightarrow{\sim} \mathrm{Nat}(h^*S, F).$$

This stable formulation is the essential point. The unstabilized condition that

$$F(a) \to \mathrm{Nat}(S, F)$$

be bijective does not in general define a Grothendieck topology, because it need not be preserved by pullback. Once stability is built into the definition, the topology axioms follow directly from the Yoneda lemma and the local transitivity argument.

The construction gives a precise universal property:

$$F \in \mathrm{Sh}(A, J) \quad \Longleftrightarrow \quad J \leq J_F.$$

Thus J_F is exactly the maximal site structure on A under which the given presheaf F satisfies descent.

References

[1] S. Mac Lane and I. Moerdijk, *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*, Springer, 1992.
[2] A. Grothendieck and J.-L. Verdier, *Théorie des topos et cohomologie ètale des schémas*, SGA 4, Exposé II, Lecture Notes in Mathematics 269, Springer, 1972.
[3] P. T. Johnstone, *Sketches of an Elephant: A Topos Theory Compendium*, Vols. 1–2, Oxford University Press, 2002.
[4] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Springer, 1998.
[5] F. Borceux, *Handbook of Categorical Algebra 3: Categories of Sheaves*, Cambridge University Press, 1994.