



HAL
open science

Topoi with enough points

Morgan Rogers, Ivan Di Liberti

► **To cite this version:**

| Morgan Rogers, Ivan Di Liberti. Topoi with enough points. 2024. hal-04523900v2

HAL Id: hal-04523900

<https://hal.science/hal-04523900v2>

Preprint submitted on 23 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

TOPOI WITH ENOUGH POINTS

IVAN DI LIBERTI AND MORGAN ROGERS

ABSTRACT. We extend Deligne’s original argument showing that locally coherent topoi have enough points, clarified using collage diagrams. We show that our refinement of Deligne’s technique can be adapted to recover every existing result of this kind, including the most recent results about κ -coherent κ -topoi. Our presentation allows us to relax the cardinality assumptions typically imposed on the sites involved. We show that a larger class of locally finitely presentable toposes have enough points and that a closed subtopos of a topos with enough points has enough points.

Keywords. topos theory, categorical logic, completeness theorem, spatial locale. **MSC2020.** 03G30, 03C75, 18B25, 18C50, 18F10, 18F70.

CONTENTS

Introduction	1
1. Collage diagrams	4
2. Improving Deligne	7
3. Recovering completeness results	13
4. Extending the class of topoi with enough points	24
Acknowledgements	29
References	29

INTRODUCTION

General landscape. This paper is concerned with the general problem of assessing whether a topos has enough points, with motivations coming both from geometry and logic. As we will see below, this problem has been influential in both fields up to the present day (see for example the work of Lurie [Lur04, VII, 4.1] and Barwick et al. [BGH18, II.3], for the case of geometry, and Espíndola and Kanalas [EK23] for the case of logic). We shall start by giving an account of the development of this problem to properly frame our contribution and the significance of our work.

In 1964 Pierre Deligne proved a very celebrated theorem in topos theory.

Theorem (Deligne, [BDS06, Exposé VI, 9.0]). Every locally coherent topos has enough points.

The theorem’s original motivation came from algebraic geometry, but after Joyal and Reyes developed the theory of classifying topoi [Rey74], it was observed by Lawvere [Law75] that Deligne’s theorem was essentially the statement of Gödel completeness theorem for first order logic in disguise. This realisation crowned Deligne’s theorem as a major result in categorical logic, and a source of inspiration for finding other completeness-like results using techniques from topos theory.

The first named author received funding from Knut and Alice Wallenberg Foundation, grant no. 2020.0199.

To the present day Deligne’s theorem remains the main argument to show that a wide class of topoi have enough points, and to some extent this paper investigates the limits (and the possibly unexploited potential) of this result. With current technology, we have essentially two different techniques to prove that (locally) coherent topoi have enough points.

- In [Joh77, 7.4], Johnstone gives an account of the original argument by Deligne, making a couple of minor adjustments for the sake of readability and elegance.
- In [LM94, IX.11], Moerdijk and Mac Lane give a more conceptual proof of the theorem, based on an injectivity property of coherent topoi. This idea was further investigated by the first author in [DL22].

These two proof strategies are – to our understanding – different in nature, and seem to rely on different aspects of coherent topoi, with Deligne’s result clearly the one with a greater chance of being generalised to other classes of topoi. Following Deligne’s theorem, it was some time before new results emerged showing that further classes of topoi have enough points. Makkai and Reyes proved that so-called *separable* topoi have enough points.

Theorem (Makkai and Reyes, [MR06, Theorem 6.2.4, page 180]). Let (\mathcal{C}, J) be a site where \mathcal{C} is a countable category with pullbacks and J is generated by a countable family of sieves. Then $\text{Sh}(\mathcal{C}, J)$ has enough points.

This result was inspired by the Fourman-Grayson completeness theorem for the logic L_{ω_1, ω_0} (see [FG82]), and indeed it is almost the translation of it into topos-theoretic language through the bridge of classifying topoi. The proof in [MR06] is a bit sketchy and of model theoretic inspiration, thus it is hard to compare this result to Deligne’s one.

Recent developments. Let us now recount the most recent developments in this topic. Lurie has imported Deligne’s original argument to the world of ∞ -topoi [Lur04, VII, 4.1]. The proof carries over with minor adjustments under the additional assumption that the ∞ -topos is hypercomplete (i.e. that it can be presented in terms of a site). Lurie’s analysis did not make Deligne’s argument sharper in any sense but stands as the best result available for ∞ -topoi in the broad landscape of geometry.

On the logical side, the main advancements are due to Esp ndola [Esp19, Esp20, EK23]. The first two of these papers generalise the result due to Makkai and Reyes to what he calls κ -coherent κ -topos, a notion designed to deliver completeness results for κ -infinitary generalizations of coherent logic. His techniques rely heavily on logic, so the relationship between these results and that of Deligne remained unclear, especially in terms of technology. We recommend [Esp19] for a comprehensive reference of all the completeness results of infinitary logic and their relation to κ -topoi. More recently, together with Kanalas, they have provided a more categorical analysis of the original statement, which also delivers a vast generalization of the original results that Esp ndola had achieved in his PhD thesis [EK23].

Theorem (Esp ndola and Kanalas, [EK23, Thm. 5.4]). Let κ be a regular cardinal. Let (\mathcal{C}, J) be a κ -site, meaning that \mathcal{C} has $< \kappa$ -limits and is of local size $\leq \kappa$ and that J is generated by a family of presieves having cardinality $\leq \kappa$. Then $\text{Sh}(\mathcal{C}, J)$ has enough κ -points.

Esp ndola and Kanalas’ proof, while being categorical in nature, does not seem to follow Deligne’s original strategy beyond inspiration. In parallel with these developments, the community revolving around topos theory has tried to understand the limits of Deligne’s original argument and its possible generalization. Quite

independently the authors of this paper and Tim Campion conjectured that every *locally finitely presentable topos could have enough points*. Some evidence of discussion around this idea can be found in a collection of MathOverflow questions [Cam, DLb, mu, Gin, DLa]. This conjecture finds its motivations in a number of examples, which we may list below:

- coherent topoi are locally finitely presentable ([Joh02b, D3.3.12]), and the gap between these two notions seems quite small, at least by a superficial analysis,
- presheaf topoi are often not coherent and yet they are always locally finitely presentable, and they are the easiest example of topos with enough points,
- Hoffmann-Lawson duality [HL78] from the theory of locales seems to support the conjecture that exponentiable topoi (which include locally finitely presentable topoi) have enough points.

In an early version of [Rog21] the second author found a flawed proof of the conjecture (see the discussion in the comments of [Cam]; Marc Hoyois pointed out the error), and this collaboration stems from an attempt to correct that mistake. Before proceeding further with discussing our contribution, we can summarise that while we were able to prove that a large class of locally finitely presentable toposes have enough points, since the original preprint of this work was released a claimed counterexample to the general conjecture has emerged (see Remark 4.1.11).

Our contribution and structure of the paper. In Section 1 we briefly recall the notion of collage of a profunctor and establish the diagrams featuring in the subsequent discussion, such as the following:

$$\begin{array}{c}
 \cdot \\
 \vdots \\
 p \text{ --- } z \rightarrow Y \begin{array}{l} \xrightarrow{f} \\ \xrightarrow{g} \end{array} X \\
 \vdots \\
 \cdot
 \end{array}$$

Collages offer a convenient setting where points (on the left) can interact with objects (on the right) of a topos in a common category.

Section 2 is the main section of the paper. After introducing the notion of *improvement* (Definition 2.1.2), which is designed to isolate the central idea of Deligne’s proof, we prove our core theorem (which we reproduce below in an easy-to-read form).

Theorem (2.1.5). Let $j : \mathcal{F} \hookrightarrow \mathcal{E}$ be an inclusion of topoi. Suppose that every point p can be improved. If \mathcal{E} has enough points, then \mathcal{F} has enough points.

Section 3 is devoted to recovering all the existing theorems concerning topoi with enough points using our technology. Corollary 3.3.1 recovers the original theorem due to Deligne, Theorem 3.3.9 implies Makkai and Reyes’ result Corollary 3.3.10, while Theorem 3.4.5 generalizes Espindola and Kanalas’ result Corollary 3.4.9. Our generalization involves introducing the notion of $(\kappa\text{-})\textit{warp}$, mildly constrained data generating a Grothendieck topology (Definition 3.3.5) and the corresponding notion of $(\kappa\text{-})\textit{woven}$ site (Definitions 3.3.7 and 3.4.4).

Theorem (3.4.5 and 4.1.1). Let κ be a regular cardinal. Let (\mathcal{C}, J) be a κ -woven site where \mathcal{C} is a category with limits of $< \kappa$ -cochains. Then $\text{Sh}(\mathcal{C}, J)$ has enough κ -points¹.

¹Points preserving limits of size $< \kappa$.

We devote a further subsection (Section 3.5) to discussing some localic implications of our work. There we recover Heckman's theorem concerning countably generated locales having enough points.

We push the boundaries of our results further in Section 4, first showing that the hypothesis of pullbacks can be eliminated from Theorems 3.3.9 and 3.4.5. This section is also where we give our best attempt to prove that locally finitely presentable topoi have enough points.

Theorem (4.1.12). A locally finitely presentable topos whose class of finitely presentable objects is essentially countable has enough points.

Finally we present a 2-dimensional analysis of our core theorem (and of Deligne's argument in general). Besides offering a more conceptual interpretation of the theorem, we show that topoi with enough points are closed under closed embeddings:

Theorem (4.2.1). A closed subtopos of a topos with enough points has enough points.

1. COLLAGE DIAGRAMS

This section recalls the notion of *collage*, which we will employ as a technical framework in which to organise the proof of our core theorem (2.1.5). After the first definitions, we take the opportunity to tailor the discussion and present some examples in order to introduce the reader to our point of view on Deligne's theorem.

1.1. Definitions.

1.1.1 (Evaluation pairings). Recall that for \mathcal{E} a topos, we have an evaluation pairing

$$\text{ev}_{\mathcal{E}} : \text{pt}(\mathcal{E}) \times \mathcal{E} \rightarrow \text{Set}$$

which is defined by $(p, X) \mapsto p^*(X)$. This pairing will play an important role in the paper, so let us spell out more in detail a concrete example. Let $\mathcal{E} = [\mathcal{C}^{\text{op}}, \text{Set}]$. Then $\text{pt}(\mathcal{E})$ is (equivalent to) the ind-completion $\text{Ind}(\mathcal{C}^{\text{op}})$. Throughout, we shall denote by

$$\begin{aligned} \mathfrak{J} &: \mathcal{C} \rightarrow [\mathcal{C}^{\text{op}}, \text{Set}] \\ \mathfrak{L} &: \mathcal{C} \rightarrow \text{Pro}(\mathcal{C}) = \text{Ind}(\mathcal{C}^{\text{op}})^{\text{op}} \subseteq [\mathcal{C}, \text{Set}]^{\text{op}} \end{aligned}$$

the Yoneda embedding and (a restriction of) its dual. Restricting along these embeddings, we find that $\text{ev}(\mathfrak{L}(C), \mathfrak{J}(D)) \cong \mathcal{C}(C, D)$, recovering the hom profunctor on \mathcal{C} .

Definition 1.1.2 (Collage). Let $P : \mathcal{F}^{\text{op}} \times \mathcal{E} \rightarrow \text{Set}$ be a profunctor between locally small categories. Recall that the **collage** of P is the category $\text{Coll}(P)$ whose collection of objects is $\text{ob}(\mathcal{F}) \sqcup \text{ob}(\mathcal{E})$ and whose collection of morphisms is $\text{mor}(\mathcal{F}) \sqcup \text{mor}(\mathcal{E}) \sqcup \bigsqcup_{F, E} P(F, E)$.

Definition 1.1.3 (Collage diagrams). Let P be as in Definition 1.1.2. A **collage diagram** for P is simply a diagram in $\text{Coll}(P)$, which we present as an ordinary diagram equipped with a vertical dotted line separating the diagram into a left and right part such that:

- To the left of the dotted line are objects and morphisms of \mathcal{F} ,
- To the right of the dotted line are objects and morphisms of \mathcal{E} ,
- Morphisms from an object F on the left to an object E on the right correspond to elements of $P(F, E)$, and
- No morphisms from right to left are permitted.

Unless stated otherwise, a diagram implicitly carries the assertion that it **commutes**, in the sense that each region on each side of the dotted line is a commuting diagram in the respective categories, while a region crossing the dotted line commutes in the sense that applying the relevant components of natural transformations on each side of the line to the elements presented by the arrows crossing the line produces identical results.

Remark 1.1.4. Let $P : \mathcal{F}^{\text{op}} \times \mathcal{E} \rightarrow \text{Set}$ be a profunctor between locally small categories, and call $\text{Coll}(P)$ its collage, then we get two canonical functors as described below.

$$\begin{array}{ccc} \mathcal{F} & & \mathcal{E} \\ & \searrow & \swarrow \\ & i_{\mathcal{F}} & \text{Coll}(P) & \leftarrow i_{\mathcal{E}} \end{array}$$

It is easy to see that these functors are both fully faithful. Moreover, $i_{\mathcal{F}}$ creates limits and $i_{\mathcal{E}}$ creates colimits.

While we have given a general definition here, we will exclusively work with collage diagrams for $\text{ev}_{\mathcal{E}}$, viewed as a profunctor $(\text{pt}(\mathcal{E})^{\text{op}})^{\text{op}} \times \mathcal{E} \rightarrow \text{Set}$, for various choices of \mathcal{E} .

Example 1.1.5 (Visualizing some diagrams in the collage). Continuing 1.1.1, let $\mathcal{E} = [\mathcal{C}^{\text{op}}, \text{Set}]$. Given a commuting square in \mathcal{C} (as in the LHS below), we have a corresponding collage diagram for $\text{ev}_{\mathcal{E}}$ (as in the RHS below):

$$\begin{array}{ccc} C \xrightarrow{x} D & & \mathfrak{L}(C) \xrightarrow{x} \mathfrak{L}(D) \\ f \downarrow & & \mathfrak{L}(f) \downarrow \\ C' \xrightarrow{y} D' & \mapsto & \mathfrak{L}(C') \xrightarrow{y} \mathfrak{L}(D') \end{array} \quad \begin{array}{c} \vdots \\ \downarrow \\ \vdots \end{array}$$


where we write $\mathfrak{L}(f)$ for the natural transformation between the representable points corresponding to f and by a mild abuse of notation we denote use the same symbol x to denote an element of $\text{ev}(\mathfrak{L}(C), \mathfrak{L}(D))$ and the corresponding morphism $C \rightarrow D$ (and similarly for y). In this context, all of the arrows in the collage diagram represent morphisms of essentially the same type, but viewed from different perspectives according to the extension of \mathcal{C} into which they have been promoted.

Remark 1.1.6 (Alternative handedness). We have chosen a site-theoretic perspective on the topos and its category of points. Unfortunately, this is dual to the ‘semantic’ perspective in which $\text{pt}(\mathcal{E})$ is treated as a category of models of a theory and \mathcal{E} is presented as a subtopos of the category of copresheaves on a category of finitely presentable models. From that perspective, working in the dual of the collages we present would be more intuitive. As such, we caution the reader to take care when working in the left-hand side of our collage diagrams, and invite them to dualize when needed.

1.2. Expressing conditions in the collage. To accustom the reader to reasoning in $\text{Coll}(\text{ev}_{\mathcal{E}})$, we identify how properties of sheaves and points manifest themselves in this setting, beginning with limits and colimits.


Lemma 1.2.1 (Lexity). Let \mathcal{E} be a topos and $p \in \text{pt}(\mathcal{E})$. Let \mathbb{D} be a small category, $F : \mathbb{D} \rightarrow \mathcal{E}$ a diagram, X an object of \mathcal{E} and $\lambda : X \Rightarrow F$ a cone over F with apex X .

Then p sends λ to a limit cone if and only if for every cone x_{\bullet} over F with apex p in $\text{Coll}(\text{ev}_{\mathcal{E}})$ there exists a unique factoring element $x^* \in p^*(X)$:




In particular, since points preserve finite limits by definition, this is always true for finite \mathbb{D} . It follows that finite limits in \mathcal{E} remain finite limits in the collage $\text{Coll}(\text{ev}_{\mathcal{E}})$. We provide a weaker dual result.

Lemma 1.2.2 (Cocontinuity). Let \mathcal{E} , p , \mathbb{D} , F and X be as in Lemma 1.2.1. Let $\gamma : F \Rightarrow X$ be a cone under F with nadir X . Then p^* sends γ to a jointly epic cone if and only if for any $x \in p^*(X)$ there exists an index $d \in \mathbb{D}$ such that x factors through γ_d in the collage diagram:



For any point p in $\text{pt}(\mathcal{E})$ we know that p^* preserves small colimits and hence jointly epic families, which means that p is always cone-projective in $\text{Coll}(\text{ev}_{\mathcal{E}})$ for jointly epic families in \mathcal{E} .

Corollary 1.2.3 (Points as a weak-projectives). Let (\mathcal{C}, J) be a site for the topos \mathcal{E} . Then an object p of $\text{Ind}(\mathcal{C}^{\text{op}})$ restricts to a point of \mathcal{E} if and only if for each J -covering sieve $S = \{f_j : C \rightarrow C_j\}$ and element $x \in p^*(\mathcal{Y}(C))$, there exists an index k and an element $x' \in p^*(\mathcal{Y}(C_k))$ completing the diagram:



Proof. By Diaconescu's theorem, a point p of $[\mathcal{C}^{\text{op}}, \text{Set}]$ restricts to a point of $\text{Sh}(\mathcal{C}, J)$ if and only if, when p^* is composed with the Yoneda embedding $\mathcal{C} \rightarrow [\mathcal{C}^{\text{op}}, \text{Set}]$ to produce a flat functor, the result is J -flat, meaning that J -covering sieves are sent to jointly epimorphic families. Applying Lemma 1.2.2 produces the stated condition. \square

However, for greater generality, we can simplify this by considering the image of the covering family, which is to say the sieve viewed as a subobject $S \twoheadrightarrow \mathfrak{J}(C)$.

Corollary 1.2.4 (Points as projectives). Let (\mathcal{C}, J) be a site for the topos \mathcal{E} . Then an object p of $\text{Ind}(\mathcal{C}^{\text{op}})$ restricts to a point of \mathcal{E} if and only if for each J -covering sieve $S = \{f_j : C \rightarrow C_j\}$ and element $x \in p^*(\mathfrak{J}(C))$, there exists a (necessarily unique) element $x' \in p^*(S)$ completing the diagram:

$$\begin{array}{ccc}
 & & S \\
 & \nearrow x' & \downarrow \\
 p & \xrightarrow{x} & \mathfrak{J}(C)
 \end{array}$$

More generally, for $j : \mathcal{F} \twoheadrightarrow \mathcal{E}$ a point $p \in \text{pt}(\mathcal{E})$ factors through \mathcal{F} if and only if for every j -dense² monomorphism $m : S \twoheadrightarrow T$ and element $x \in p^*(X)$, there is a factorization of x through m in $\text{Coll}(\text{ev}_{\mathcal{E}})$:

$$\begin{array}{ccc}
 & & S \\
 & \nearrow x' & \downarrow m \\
 p & \xrightarrow{x} & T
 \end{array}$$

1.2.5. Given a subtopos $\mathcal{F} \twoheadrightarrow \mathcal{E}$ the projectivity condition of Corollary 1.2.4 ensures that we get a fully faithful functor between collages,

$$\text{Coll}(\text{ev}_{\mathcal{F}}) \hookrightarrow \text{Coll}(\text{ev}_{\mathcal{E}}),$$

where fullness on the point side comes from the fact that inclusions of topoi are representably fully faithful in the 2-category of topoi [Rog24, Lemma 2.1.3].

2. IMPROVING DELIGNE

2.1. Improving points. Deligne's proof is difficult for several reasons. Our contribution is to decompose it into two layers which can independently be understood and generalized using well-established categorical techniques.

It is straightforward to express the condition of \mathcal{E} having enough points in the context of the collage $\text{Coll}(\text{ev}_{\mathcal{E}})$.

²Beware that we conflate the inclusion of topoi and the local operator in \mathcal{E} corresponding to it, since there is little risk of confusion; a monomorphism m is called j -dense if j_*j^*m is an isomorphism.

Theorem 2.1.5 (Core theorem). Let $j : \mathcal{F} \rightarrow \mathcal{E}$ be an inclusion of topoi. Suppose that for every point p and data (x, y, m, w) as in Definition 2.1.2, there exists an improvement of p with respect to (x, y, m, w) . Assuming Zorn's Lemma and that \mathcal{E} has enough points, \mathcal{F} has enough points.

Proof. Since \mathcal{E} has enough points, given any parallel pair $f, g : j_*(Y) \rightrightarrows j_*(X)$ there is some point p separating them as in (1). Fix corresponding elements $x \neq y \in p^*(j_*(X))$. Without loss of generality, we need only consider the case where m is the monomorphism $J \rightarrow \Omega$ classifying j -dense monomorphisms.

Using Zorn's Lemma, put a well-ordering on the elements $w \in p^*(\Omega)$; denote the indexing limit ordinal κ . For each $\alpha \leq \kappa$, we define by induction a point p_α , a morphism $n_\alpha : p_\alpha \rightarrow p$ in $\mathbf{pt}(\mathcal{E})^{\text{op}}$ and an element $v_\alpha \in p_\alpha^*(J)$ which will be an improvement of p with respect to (x, y, m, w_β) for all $\beta < \alpha$

- We take $p_0 := p$ and $n_\alpha = \text{id}_p$.
- For $\alpha = \beta + 1$, we take p_α to be the improvement of p_β with respect to $(n_\beta; x, n_\beta; y, m, n_\beta; w_\beta)$, and $n_\alpha = n; n_\beta$, where n is the morphism appearing in the improvement diagram (2), appealing to Lemma 2.1.3.
- For α a limit ordinal, we first construct the colimit p_α in $\mathbf{pt}(\mathcal{E})$ of the $\{p_\beta \mid \beta < \alpha\}$; this remains an improvement of p with respect to (x, y, m, w_β) for all $\beta < \alpha$ because $p_\alpha^*(j_*(X))$ and $p_\alpha^*(S)$ are computed as the directed colimit of the corresponding values of p_β^* .

The result of this procedure, $p_{(1)} := p_\kappa$, is thus an improvement of p with respect to (x, y, m, w) for all $w \in p^*(\Omega)$. However, *a priori* there may be some elements of $p_{(1)}^*(\Omega)$ for which projectivity fails. To resolve this, it suffices to repeat this procedure of iterative improvement to obtain $p_{(2)}, p_{(3)}, \dots$, since all elements of the colimit $p_{(\omega)}$ of this sequence of points will be represented by elements of $p_{(n)}$ for some n . This $p_{(\omega)}$ is a point of \mathcal{F} by Corollary 1.2.4 which distinguishes f, g . Since x, y were generic, we conclude that \mathcal{F} has enough points, as required. \square


The proof of Theorem 2.1.5 (rather than its statement) will serve as a foundation for the results in the remainder of the paper. In each proof, the principles involved are the same: demonstrate the existence of a class of improvements using hypotheses on the subtopos or its presentation and iteratively construct points of the subtopos from those improvements. We will not bother to explicitly mention the parallel pair f, g to be distinguished in proofs from here on.

Remark 2.1.6 (Pullbacks provide canonical improvements). If the pullback of m along w exists in $\text{Coll}(\text{ev}_\mathcal{E})$ and *some* improvement of p with respect to (x, y, m, w) exists, then the pullback is automatically an improvement, and in fact the terminal improvement with respect to (x, y, m, w) independently of x and y ; this is in particular the case when p is projective. Moreover, n will be a monomorphism (and hence an epimorphism in $\mathbf{pt}(\mathcal{E})$ by Remark 1.1.4). We expect that pullbacks across the gluing boundary in $\text{Coll}(\text{ev}_\mathcal{E})$ occur only very rarely, but we can use a weaker variant of this hypothesis to eliminate Zorn's Lemma from Theorem 2.1.5.

Corollary 2.1.7. Let $j : \mathcal{F} \rightarrow \mathcal{E}$ be an inclusion of topoi. Suppose that for every point p and data (x, y, m, w) as in Definition 2.1.2, there exists a **terminal** improvement of p with respect to (x, y, m, w) . If \mathcal{E} has enough points, so does \mathcal{F} .

Proof. With the set-up of x, y and m as in the proof of Theorem 2.1.5, consider the directed poset of finite subsets of $p^*(\Omega)$. Given such a subset $W = \{w_1, \dots, w_k\}$,

consider the corresponding diagram,



where p_W is the terminal improvement of p with respect to the indicated data. Composing with projection morphisms on the right, we deduce that p_W is also an improvement of w_i for each i by Lemma 2.1.4, and hence we get a canonical morphism $p_W \rightarrow p_{W'}$ whenever $W' \subseteq W$. Thus these subsets form a codirected diagram of improvements whose colimit in $\text{pt}(\mathcal{E})$ we can substitute for $p_{(1)}$ in the proof of Theorem 2.1.5. The remainder of the proof continues as before. \square

Remark 2.1.8 (Alternative orderings). We selected the classifying j -dense monomorphism for convenience in the proofs of Theorem 2.1.5 and Corollary 2.1.7, but it is worth noting that the set $p^*(\Omega)$ that we impose a well-ordering on in the proof of Theorem 2.1.5 already has a natural ordering inherited from the internal order relation on Ω . However, there is no relation *a priori* between the ordering on elements and their respective improvements. At best, it is possible to construct a diagram of improvements indexed by this ordering under the hypotheses of Corollary 2.1.7: for each w , we consider the filtered diagram of finite subsets of elements below w and proceed to construct a joint improvement with respect to all elements below w via the construction in the proof of Corollary 2.1.7. It may be that the ordering on Ω can be better exploited to obtain a constructive result under weaker hypotheses.

Exploitation of Corollary 2.1.7 will depend on a more systematic investigation of the features of collages than we have time or space for here.

2.2. Making it easier to find improvements. The existence of improvements with respect to arbitrary data seems demanding. Fortunately, we can relax it in a natural way.

Lemma 2.2.1 (Generators suffice). Let $j : \mathcal{F} \rightarrow \mathcal{E}$ be an inclusion of topoi, p a point of \mathcal{E} , $x \neq y \in p^*(j_*(X))$. Then p admits an improvement with respect to (x, y, m, w) for all j -dense $m : S \rightarrow T$ and elements $w \in p^*(T)$ if and only if this is true for T in a generating set of objects of \mathcal{E} .

Proof. This is an immediate consequence of Lemma 1.2.2: fixing a generating family, every object T of \mathcal{E} admits a covering family of morphisms whose domains are in this family, and any element of $p^*(T)$ must factor through one of the legs of the covering family. An improvement with respect to the factored morphisms is an improvement with respect to the original morphisms by Lemma 2.1.4. \square

2.2.2 (Moving away from the generic j -dense morphism). Examining the proof of Theorem 2.1.5 in light of Lemma 2.2.1, we find that we can get away with fewer improvements in this way. For instance, suppose $\mathcal{E} = \text{PSh}(\mathcal{C})$ and p is a representable point (this special case will be the focus of Section 3). Then elements of $\mathcal{L}(C)^*(\Omega)$ are in bijection with the set of sieves on C in \mathcal{C} ; if \mathcal{C} is a category with countably many morphisms then there could be uncountably many such elements,

but $\coprod_{C, D \in \mathcal{E}} \mathcal{A}(C)^*(\mathcal{A}(D))$ is countable by assumption, and these are sufficient for constructing improvements.

In fact, we can restrict not only the codomains of cospans featuring in improvements, but the class of cospans itself using Lemmas 2.1.4 and the following definition and lemma which generalize the site-theoretic notions of presieve and multicomposite of sieves.

Definition 2.2.3 (Generating presieve and multicomposites). Suppose we are given an inclusion $j : \mathcal{F} \rightarrow \mathcal{E}$ and \mathcal{T} a generating family of objects for \mathcal{E} . A **generating presieve** for a j -dense monomorphism $m : S \rightarrow T$ with $T \in \mathcal{T}$ is a collection of morphisms $\{g_i : T_i \rightarrow T\}$ where $T_i \in \mathcal{T}$ whose joint image is m . A monomorphism $m' : S' \rightarrow T$ is a **multicomposite over** m if there exists a generating presieve for m as above and j -dense monomorphisms $m_i : S_i \rightarrow T_i$ such that m' is the joint image of $m_i; g_i$.

Lemma 2.2.4. For $j : \mathcal{F} \rightarrow \mathcal{E}$, \mathcal{T} and $m : S \rightarrow T$ as in Definition 2.2.3, a multicomposite m' over m is j -dense.

Proof. If m' is a multicomposite over m , by definition we have a diagram:

$$(4) \quad \begin{array}{ccccc} S_i & \cdots & S_k & \longrightarrow & S' \\ \downarrow m_i & & \downarrow m_k & & \\ T_i & \cdots & T_k & & \\ & \searrow & \swarrow & & \\ & S & & & \\ & \downarrow m & & & \\ & T & & & \end{array}$$

(Note: The diagram shows a curved arrow from S' to T labeled m' , and a horizontal line between T_i and T_k with arrows pointing towards S .)

where the horizontal line indicates a jointly epimorphic family; since j^* preserves images and epimorphic families and sends m and each m_i to an isomorphism, $j^*(m')$ is the image of a jointly epimorphic family, so is an isomorphism. Thus m' is j -dense as claimed. \square

Proposition 2.2.5 (Minimizing the necessary number of improvements). Let $j : \mathcal{F} \rightarrow \mathcal{E}$ be an inclusion of topoi, p a point of \mathcal{E} , $x \neq y \in p^*(j_*(X))$. Consider the preorder where:

- elements are cospans (m, w) in $\text{Coll}(\text{ev}_{\mathcal{E}})$ whose codomain T is in a generating set of objects of \mathcal{E} , where $m : S \rightarrow T$ is j -dense and $w \in p^*(T)$;
- we have $(m, w) \leq (m', w')$ if there exists a pair (u, t) as in (3) such that $m; u = t; m'$ and $w' = w; u$.

Then for fixed $n : p' \rightarrow p$ in $\text{Coll}(\text{ev}_{\mathcal{E}})$, there exists v with $(p', n, v) \circ_p (x, y, m, w)$ for all (m, w) if and only if holds for a subset of the (m, w) whose closure under pullbacks and multicomposites provides a lower bound for every element with respect to this ordering.

Proof. By the first point of Lemma 2.1.4, if (m, w) is a lower bound for (m', w') in the preorder, then $(p', n, v) \circ_p (x, y, m, w)$ entails $(p', n, v; t) \circ_p (x, y, m', w')$, whence a collection of lower bounds is sufficient. By the second point of Lemma 2.1.4 it is sufficient for a pullback of (m, w) to be a lower bound. Finally, given a multicomposite m' of morphisms for which p is projective as in (4) and an element $w \in p^*(T)$ we can construct the required lift along m' by working up the diagram. It is worth noting that we only actually use projectivity with respect to *one* of the m_i in this construction; we will exploit this fact later on. \square

Our main use case for Proposition 2.2.5 will be for minimizing the number of iterations needed in the induction of Theorem 2.1.5. The following lemma will (under suitable hypotheses) enable us to avoid the secondary induction procedure in the proof.

Lemma 2.2.6 (Folding improvements into a colimit). Let \mathcal{E} be a topos. Suppose that we are given a point $p \in \text{pt}(\mathcal{E})$, elements $x \neq y \in p^*(Z)$ and a morphism $m : S \rightarrow T$ in \mathcal{E} . Suppose that we have constructed a cofiltered diagram $D : \mathbb{I} \rightarrow \text{Coll}(\text{ev}_{\mathcal{E}})/p$ in the slice category on objects $n : p' \rightarrow p$ such that $n; x \neq n; y$ (prospective improvements of p). Suppose that for each $n : p' \rightarrow p$ in the diagram and element $w \in p'^*(T)$ there is a morphism $n' : p'' \rightarrow p'$ in the diagram and an element $v \in p''^*(S)$ such that $(p'', n', v) \circ_{p'}(n; x, n; y, m, w)$. Then the limit q of this diagram (the colimit in $\text{pt}(\mathcal{E})$) is projective with respect to m .

Proof. Every element of $q^*(T)$ is represented by some element $v \in p'^*(T)$, and the existence of the improvement p'' in the diagram provides the morphism required for projectivity. \square

3. RECOVERING COMPLETENESS RESULTS

To find improvements in specific cases, we finally need to start exploiting presentations of topoi as sheaves on sites. We develop some auxiliary results for this setting in Section 3.1 before applying them in Section 3.3. We progressively strengthen hypotheses on the site in the remainder of the section, assuming that the site has further limits in Section 3.4 and then that it is a frame in 3.5.

3.1. Presheaf case preliminaries. The simplest situation, which will be sufficient for recovering all known completeness results, is to consider an inclusion of topoi of the form $\text{Sh}(\mathcal{C}, J) \rightarrow \text{PSh}(\mathcal{C})$. This situation is particularly convenient: not only does any presheaf topos have enough points, but the most natural class of enough points are the representables, whose inverse image functors preserve all limits. Indeed, all objects of $\text{pt}(\text{PSh}(\mathcal{C})) \simeq \text{Ind}(\mathcal{C}^{\text{op}})$ are filtered colimits of representable points, so that properties of $\text{pt}(\text{PSh}(\mathcal{C}))$ can be deduced from those of \mathcal{C} , as we demonstrate now.

Proposition 3.1.1. Let \mathcal{C} be an idempotent-complete small category. The following are equivalent:

- (1) \mathcal{C} has pushouts,
- (2) $\text{Ind}(\mathcal{C})$ has pushouts.

Proof. The implication (2) \Rightarrow (1) follows from the fact that when \mathcal{C} is idempotent complete, the Yoneda embedding $\mathcal{C} \rightarrow \text{Ind}(\mathcal{C})$ creates any finite colimits that exist in $\text{Ind}(\mathcal{C})$, since \mathcal{C} can be recovered as the full subcategory on the finitely presentable objects, which are closed under finite colimits.

To prove (1) \Rightarrow (2), let \ulcorner be the walking span. We can present the fact that \mathcal{C} has pushouts via the existence of a left adjoint

$$\text{colim}_{\ulcorner} : \mathcal{C}^{\ulcorner} \rightarrow \mathcal{C}$$

to the diagonal functor $\Delta : \mathcal{C} \rightarrow \mathcal{C}^{\ulcorner}$.

Applying the Ind -completion to this adjunction, we obtain an adjunction between $\text{Ind}(\mathcal{C})$ and $\text{Ind}(\mathcal{C}^{\ulcorner})$. Because the walking span is a well founded category in the sense of [Hen23, 3.1], it follows from [Hen23, Thm. 1.3] that we have an equivalence of categories, $E : \text{Ind}(\mathcal{C}^{\ulcorner}) \xrightarrow{\cong} \text{Ind}(\mathcal{C})^{\ulcorner}$. Moreover, in the diagram below, it is straightforward to see that the solid triangle commutes. Thus, it follows that $\Delta_{\text{Ind}(\mathcal{C})}$ has a

left adjoint, as required.

$$\begin{array}{ccc}
 \text{Ind}(\mathcal{C})^\Gamma & & \\
 \uparrow E & \swarrow \Delta_{\text{Ind}(\mathcal{C})} & \\
 \text{Ind}(\mathcal{C}^\Gamma) & \xleftarrow{\text{Ind}(\Delta_{\mathcal{C}})} & \text{Ind}(\mathcal{C}) \\
 & \searrow \text{Ind}(\text{colim}_{\mathcal{C}}) & \\
 & &
 \end{array}$$

□

Remark 3.1.2. The forward implication of \mathcal{C} having pushouts implies $\text{Ind}(\mathcal{C})$ has pushouts does not depend on \mathcal{C} being idempotent complete, and this is the direction we will use in the following; more specifically, we deduce that if \mathcal{C} has pullbacks, then $\text{pt}(\text{PSh}(\mathcal{C}))$ has pushouts. The presence of pullbacks is imposed in the hypotheses of all well-known completeness results; we examine how it can be weakened in Section 4.1.

While the idempotent-completeness hypothesis is very mild, in that $\text{PSh}(\mathcal{C})$ is equivalent to the topos of presheaves on the idempotent-completion of \mathcal{C} , we can drop even this condition in the pursuit of sufficient structure to perform the construction of improvements: see Corollary 4.1.4.

3.2. Through the looking glass. The other convenient feature of $\text{Coll}(\text{ev}_{\text{PSh}(\mathcal{C})})$ is that we can unconditionally slide representables (presheaves and points) across the gluing line. To take advantage of this, when considering a J -covering sieve S on C , rather than presenting it as a dense monomorphism $S \rightarrow \mathfrak{Y}(C)$, we instead consider its image under the Yoneda embedding as a collection of morphisms between representables in $\text{PSh}(\mathcal{C})$. Indeed, this allows us to restrict to a weakly final subdiagram or ‘generating presieve’ for S when needed.

Definition 3.2.1. Let Ψ be a class of diagrams and \mathcal{C} a category. We say a sieve S over C in \mathcal{C} is **generated by a presieve in Ψ** if, when S is viewed as a subcategory of \mathcal{C}/C , there exists a weakly final subcategory in Ψ . That is, there is a collection of morphisms $\{f_i : C_i \rightarrow C\}$ generating S and compatible morphisms between these forming a diagram in Ψ .

The value of this definition is that for any presheaf $F : \mathcal{C}^{\text{op}} \rightarrow \text{Set}$ we have,

$$\lim_{(f:C' \rightarrow C) \in S} F(C') \hookrightarrow \lim_{f_i:C_i \rightarrow C} F(C_i),$$

where on the right-hand side we compute a limit whose shape is in Ψ .

The last ingredient for constructing improvements from these ideas is to focus on points whose inverse image functors preserve limits indexed by these generating presieves.

Proposition 3.2.2 (Preserving limits provides improvements). Let Φ be a class of (limit) diagrams. Consider the inclusion $j : \text{Sh}(\mathcal{C}, J) \rightarrow \text{PSh}(\mathcal{C})$ where \mathcal{C} has pullbacks and J -covering sieves are generated by presieves in Φ^{op} . Let p be a point of $\text{PSh}(\mathcal{C})$ such that p^* preserves Φ -limits and $x \neq y \in p^*(j_*(X))$. Then p has an improvement p' with respect to (x, y, m, w) for all j -dense $m : S \rightarrow T$ and $w \in p^*(T)$ such that p'^* also preserves Φ -limits.

Remark 3.2.3 (Deligne’s blueprint). While it is more general, the proof of Proposition 3.2.2 essentially follows the strategy of part of Deligne’s original proof [Joh77, Lemma 7.43] translated into our collage diagram set-up and exploiting the presence of pullbacks more efficiently through Proposition 3.1.1. We believe this presentation

makes it much easier to follow. From a pedagogical perspective, we consider this to be an important contribution of our paper.

Proof. By Lemma 2.2.1, we need only construct improvements with respect to the dense monomorphisms corresponding to J -covering sieves. Therefore, suppose we are given a J -covering sieve $S = \{f_i : C_i \rightarrow C\}$ and an element $w \in p^*(\mathfrak{L}(C))$ (we omit x, y from the diagram):

$$\begin{array}{c}
 \cdot \\
 \vdots \\
 \mathfrak{L}(C_i) \quad \cdots \quad \mathfrak{L}(C_k) \\
 \searrow \mathfrak{L}(f_i) \quad \swarrow \mathfrak{L}(f_k) \\
 \mathfrak{L}(C) \\
 \xrightarrow{p} \quad \xrightarrow{w} \mathfrak{L}(C)
 \end{array}$$

Since we are working in $\text{Coll}(\text{ev}_{\text{PSh}(\mathcal{C})})$, we can translate the representable presheaves across the diagram.

$$\begin{array}{c}
 \cdot \\
 \vdots \\
 \mathfrak{L}(C_i) \quad \cdots \quad \mathfrak{L}(C_k) \\
 \searrow \mathfrak{L}(f_i) \quad \swarrow \mathfrak{L}(f_k) \\
 \mathfrak{L}(C) \\
 \xrightarrow{p} \quad \xrightarrow{w} \mathfrak{L}(C)
 \end{array}$$

By Proposition 3.1.1, $\text{pt}(\text{PSh}(\mathcal{C}))^{\text{op}} \simeq \text{Ind}(\mathcal{C}^{\text{op}})^{\text{op}}$ inherits pullbacks from \mathcal{C} , so we can construct the pullbacks,

$$(5) \quad \begin{array}{c}
 p_i \quad \cdots \quad p_k \quad \xrightarrow{\quad} \quad \mathfrak{L}(C_i) \quad \cdots \quad \mathfrak{L}(C_k) \\
 \searrow g_i \quad \swarrow g_k \quad \searrow \mathfrak{L}(f_i) \quad \swarrow \mathfrak{L}(f_k) \\
 p \quad \xrightarrow{w} \quad \mathfrak{L}(C)
 \end{array}$$

Since $j_*(X)$ is a sheaf, for any morphism $v : D \rightarrow C$ in \mathcal{C} we have,

$$(6) \quad \begin{aligned}
 \text{ev}(\mathfrak{L}(D), j_*(X)) &\cong \lim_{f_i \in S^{\text{op}}} \text{ev}(\mathfrak{L}(\text{dom}(v^*(f_i))), j_*(X)) \\
 &\hookrightarrow \lim_{f_i \in S'^{\text{op}}} \text{ev}(\mathfrak{L}(\text{dom}(v^*(f_i))), j_*(X)),
 \end{aligned}$$

where S' is a generating sieve for S of shape in Φ^{op} .

Let $\text{Exp}(w)$ be the category of factorizations of w through representables. That is, the category whose objects are pairs of composable morphisms

$$p \xrightarrow{u} \mathfrak{L}(D) \xrightarrow{v} \mathfrak{L}(C)$$

such that $u;v = w$ (we denote this object $(u, \mathfrak{L}(D), v)$), and whose morphisms are morphisms between the middle objects making the resulting triangles commute.

The dual of $\text{Exp}(w)$ is an initial subcategory of the filtered diagram defining p , so colimits over it still commute with limits in Φ . By the hypothesis on S . Thus we have:

$$\begin{aligned}
p^*(j_*(X)) &\cong \text{colim}_{(u, \mathfrak{A}(D), v) \in \text{Exp}(w)^{\text{op}}} \text{ev}(\mathfrak{A}(D), j_*(X)) \\
(7) \quad &\cong \text{colim}_{(u, \mathfrak{A}(D), v) \in \text{Exp}(w)^{\text{op}}} \lim_{f_i \in S^{\text{op}}} \text{ev}(\mathfrak{A}(\text{dom}(v^*(f_i))), j_*(X)) \quad \text{by (6)} \\
(8) \quad &\hookrightarrow \text{colim}_{(u, \mathfrak{A}(D), v) \in \text{Exp}(w)^{\text{op}}} \lim_{f_i \in S'^{\text{op}}} \text{ev}(\mathfrak{A}(\text{dom}(v^*(f_i))), j_*(X)) \\
(9) \quad &\cong \lim_{f_i \in S'^{\text{op}}} \text{colim}_{(u, \mathfrak{A}(D), v) \in \text{Exp}(w)^{\text{op}}} \text{ev}(\mathfrak{A}(\text{dom}(v^*(f_i))), j_*(X)) \\
&\cong \lim_{f_i \in S} p_i^*(j_*(X)).
\end{aligned}$$

where the injection (8) is constructed as a colimit of the injections in (6), and remains injective to the fact that $\text{Exp}(w)$ is filtered, so this colimit preserves monomorphisms.

In particular, there must exist some index i for which $g_i(x) \neq g_i(y)$ in $p_i^*(j_*(X))$, and this p_i is thus the desired improvement. Note that p_i^* preserves Φ -limits because it too is presented by a diagram of shape $\text{Exp}(w)$. \square

Remark 3.2.4 (On the sharpness of the assumptions). We don't need either of the intermediate steps in (7) and (9) to be bijections; injections would be enough. As such, this argument would still work with:

- A J -separated presheaf in place of the sheaf $j_*(X)$; as such, when we construct enough points of $\text{Sh}(\mathcal{C}, J)$ using this result we will actually get a set of points which are faithful on the wider class of J -separated presheaves. This is not surprising: the sheafification functor is always faithful on separated presheaves, so this observation cannot lead to a stronger result.
- p only weakly preserving Φ -limits in the sense that the comparison map is injective. Little enough is known (for the time being) about the sufficient conditions for colimits to weakly commute with limits in this way that it is unclear whether this expands our base of examples.

Remark 3.2.5 (The main special case). The principal special case of Proposition 3.2.2 is when Φ is a class of discrete diagrams, so that (6) becomes the injection of the limit appearing in the sheaf condition into a product (the special case of finite products is used in [Joh77, Lemma 7.43]). Complementary examples are challenging to come by, since the inverse image functor p^* always preserves finite limits, and for simple candidates of Φ (such as chains, cofiltered diagrams or connected diagrams of a bounded cardinality), preserving finite limits and Φ -limits is equivalent to preserving finite limits and sufficiently large products.

3.3. Deligne and Makkai-Reyes. By either directly applying the results derived so far or manipulating their proofs, we can recover all known completeness results for topoi. The most straightforward is Deligne's original completeness result, exploiting the special case of Proposition 3.2.2 where Φ is the class of finite diagrams.

Corollary 3.3.1 (Deligne, [BDS06, Exposé VI, 9.0]). Let (\mathcal{C}, J) be a site where \mathcal{C} has pullbacks and J -covering sieves are finitely generated (in other words, $\text{Sh}(\mathcal{C}, J)$ is a locally coherent topos). Then assuming Zorn's Lemma, $\text{Sh}(\mathcal{C}, J)$ has enough points.

Proof. By Proposition 3.2.2, since all points preserve finite limits, all improvements exist, so we may apply Theorem 2.1.5 to the inclusion $\text{Sh}(\mathcal{C}, J) \hookrightarrow \text{PSh}(\mathcal{C})$ to deduce the existence of enough points of $\text{Sh}(\mathcal{C}, J)$. \square


3.3.2 (Not all points are representable!). As we observed in Section 3.1, we can always take the starting collection of enough points of $\text{PSh}(\mathcal{C})$ to be the representables, whose inverse image functors preserve all small limits. If improvements are constructed using Proposition 3.2.2, this remains true for the points p_1, p_2, \dots constructed in the induction procedure in the proof of Theorem 2.1.5 (the points p_i constructed in (5) will also be representables). The obstacle to greater generality, therefore, is that the point p_ω constructed at the first limit ordinal need not preserve the same limits that the p_n do. There are two solutions to this problem. The first is to constrain the site so that p_ω is the end of the procedure. To do so efficiently, we can take advantage of Proposition 2.2.5.

3.3.3 (From topologies to bases). Consider the special case of Proposition 2.2.5 where $\mathcal{E} = \text{PSh}(\mathcal{C})$, $\mathcal{F} = \text{Sh}(\mathcal{C}, J)$ and $p = \mathfrak{z}(C)$ is a representable point. Then for any $w : C \rightarrow D$ and J -covering sieve $m : S' \twoheadrightarrow \mathfrak{z}(D)$, we can construct the diagram:

$$\begin{array}{ccccc}
 & & \cdot & & \\
 & & \vdots & & \\
 p' & \overset{\text{---} v \text{---}}{\dashrightarrow} & S & \longrightarrow & S' \\
 \downarrow n & & \downarrow w^*(m) & \downarrow j & \downarrow m \\
 \mathfrak{z}(C) & \overset{\text{---}}{=} & \mathfrak{z}(C) & \xrightarrow{w} & \mathfrak{z}(D) \\
 & \searrow x & & \searrow y & \\
 & & j_*(X) & &
 \end{array}$$

whence $(\text{id}_C, w^*(m))$ is a lower bound for (w, m) in the preorder of Proposition 2.2.5. As such, we can always construct a lower set determined by J -covering sieves over C , whence the following definition.

Definition 3.3.4 (Cotree and multicomposites). For our purposes, a **cotree**³ is a poset with a maximal element in which every interval $[x, y] = \{z \mid x \leq z \leq y\}$ is (empty or) isomorphic to a finite total order, as illustrated in the following Hasse diagram:



Note that there is no restriction on the width of the cotree. A **cotree in a category** \mathcal{C} is a diagram indexed by a cotree. For each object in such a diagram, we call the edges entering it its *predecessors*. The **multicomposite** of a cotree is the presieve obtained by taking the composites of the maximal branches in the cotree, or if the cotree is trivial (of height 0), consists of the identity morphism on the root.

Definition 3.3.5 (Warp). Let (\mathcal{C}, J) be a site. A **warp** for J is a collection $\mathcal{B}(C)$ of J -covering presieves over C for each C such that:

- the pullback of a sieve generated by a presieve in \mathcal{B} contains a presieve in \mathcal{B} , and

³We call this a cotree so that our later terminology will be compatible with that of Espindola and Kanalas in [EK23].

- every J -covering sieve S' over C' contains a presieve constructed as the multicomposite of a cotree in which the predecessors of each node form a presieve in \mathcal{B} .

Note that we can identify the multicomposites appearing in Definition 3.3.5 with those defined earlier in Definition 2.2.3 by identifying a presieve with the sieve it generates (as a subobject in $\text{PSh}(\mathcal{C})$).

Remark 3.3.6 (Comparison with pretopologies). Our Definition 3.3.5 has weaker conditions than usual notions of basis or pretopology (see [LM94, Definition III.2.2]): we do not demand that elements of a warp be closed under pullbacks or multicomposition since closing under these operations recovers J . It is hard to claim novelty here. The important thing is that this reduces the number of sieves needed to verify projectivity, which was the goal!

Definition 3.3.7 (ω -woven sites). A site (\mathcal{C}, J) is ω -woven if J has a warp \mathcal{B} such that $\mathcal{B}(C)$ is countable at every object C .

3.3.8 (Weaving nomenclature). The choice of terminology comes from textile production, where the *warp* constitutes the longitudinal threads around which the *weft* is *woven*. We have deliberately selected terms diverging from existing terminology (from topology, for example) to avoid potential confusion.

Theorem 3.3.9. Let (\mathcal{C}, J) be an ω -woven site where \mathcal{C} has pullbacks. Then assuming the axiom of *dependent countable choice*, $\text{Sh}(\mathcal{C}, J)$ has enough points.

Proof. Let $\xi : \omega \times \omega \rightarrow \omega$ be a bijection such that $\xi(a, b) \geq a$ for every b .

For a given representable point $\mathfrak{A}(C)$ and $x \neq y \in \mathfrak{A}(C)(j_*(X))$, we can proceed with the inductive construction in the proof of Theorem 2.1.5 where improvements are obtained using Proposition 3.2.2, as follows. Set $\mathfrak{A}(C_0) := \mathfrak{A}(C)$ and $n_{0,0} := \text{id}_{\mathfrak{A}(C)}$. At the $(k+1)$ th step, let $(a, b) = \xi^{-1}(k)$ and define $\mathfrak{A}(C_{k+1})$ and $n_{k+1,k}$ via the improvement diagram:

$$\begin{array}{ccc}
 \mathfrak{A}(C_{k+1}) & \overset{\text{---} v \text{---}}{\dashrightarrow} & S_b \\
 \downarrow n_{k+1,k} & & \downarrow m_b \\
 \mathfrak{A}(C_k) & \xrightarrow{n_{k,a}} & \mathfrak{A}(C_a) \\
 & \searrow & \nearrow \\
 & & j_*(X)
 \end{array}$$

$\begin{array}{c} \vdots \\ \vdots \\ \vdots \end{array}$

where $n_{k,a}$ is the composite of $n_{k,k-1}; n_{k-1,k-2}; \dots; n_{a+1,a}$ and m_b is the b th sieve in the warp of J at $\mathfrak{A}(C_a)$. We need dependent countable choice to complete the construction, since Proposition 3.2.2 requires us to make a choice of improvement at each step.

We now show that this indexing acrobatics ensures that the colimit p_ω of the resulting inductively constructed diagram is projective with respect to all J -covering sieves. Suppose we are given a J -covering sieve $m' : S' \twoheadrightarrow \mathfrak{A}(C')$ and an element of $p_\omega^*(\mathfrak{A}(C'))$; the latter is represented by a morphism $w : \mathfrak{A}(C_a) \rightarrow \mathfrak{A}(C')$. By assumption, the pullback $w^*(m')$ contains a sieve m which can be presented as the multicomposite of a cotree of presieves in \mathcal{B} . We proceed by induction on the height of the multicomposite as follows: if the cotree has height 0, m is the maximal sieve, so projectivity is trivial. Otherwise, suppose that we have proved the result for cotrees of height h and we are given a cotree of height $h+1$. the first layer of the

cotree is the b th sieve in $\mathcal{B}(C_a)$, whence the diagram contains an improvement (C_γ , i say) with respect to $(n_{a,0}; x, n_{a,0}; y, m_b, \text{id}_{C_a})$. To continue the induction, we take the pullback layer by layer of the subtree corresponding to the predecessor chosen by this improvement to produce a cotree over C_γ , substituting the first pulled back sieve for one in $\mathcal{B}(C_\gamma)$ before proceeding to the next layer; this yields a suitable cotree of smaller height over C_γ , so we are done by the induction hypothesis. This is all that is required by Lemma 2.2.6. \square

As a special case, we recover the theorem of Makkai and Reyes regarding so-called ‘separable topoi’.

Corollary 3.3.10 ([MR06, Theorem 6.2.4, page 180]). Let (\mathcal{C}, J) where \mathcal{C} is a countable category with pullbacks and J is generated by a countable family of sieves. Then $\text{Sh}(\mathcal{C}, J)$ has enough points.

Proof. It suffices to observe that the collection of all pullbacks of the generating sieves at each object is countable, so forms the countable warp required to apply Theorem 3.3.9. \square

3.4. Espíndola-Kanalias. The second of the options alluded to in 3.3.2 for finding points which preserve sufficient limits is to construct the diagram in the proof of Theorem 2.1.5 using limits in \mathcal{C} rather than colimits in $\text{pt}(\mathcal{C})$. As long as \mathcal{C} has the required limits, we can do this recursively until the diagram of points produced is sufficiently filtered for their colimit in $\text{pt}(\text{PSh}(\mathcal{C}))$ to have inverse image functors which preserve the desired limits. For reasons that we will explain in Example 3.4.6, this extension requires us to abandon the linear induction, instead constructing a cotree of potential improvements from which we only select an overall improvement after the fact. First we need to generalize the ingredients encountered in the last section.

Definition 3.4.1 (Regular cardinal). Recall that a cardinal κ is **regular** if it is equal to its own cofinality, which is to say that any chain of ordinals below κ whose colimit is κ has cardinality κ .

For our purposes, the relevant consequences are that:

- a κ -indexed diagram is κ -filtered, meaning that κ -indexed colimits commute with limits of size $< \kappa$, and
- we can construct a bijection $\xi : \kappa \times \kappa \rightarrow \kappa$ such that $\xi(\alpha, \beta) \geq \alpha$ for all β , as we did for the case $\kappa = \omega$.

Definition 3.4.2 (κ -cotree and multicomposites). For any cardinal κ , a **κ -cotree** is a poset with a maximal element 1 in which every interval $[x, y] = \{z \mid x \leq z \leq y\}$ is either empty or isomorphic to the dual of an ordinal $\lambda < \kappa$. The cotrees of Definition 3.3.4 are ω -cotrees with this convention.

A κ -cotree in a category \mathcal{C} is a diagram indexed by a κ -cotree such that the image of a limit ordinal in any branch is the limit of the images of objects below it⁴.

The **multicomposite** of a κ -cotree in \mathcal{C} having limits of $< \kappa$ -cochains is the sieve over the root constructed from the limits of maximal branches in the cotree.

Definition 3.4.3 (κ -warp). A **κ -warp** for a site (\mathcal{C}, J) where \mathcal{C} has limits of $< \kappa$ -indexed cochains is a collection $\mathcal{B}(C)$ of J -covering presieves over C for each C such that:

- the pullback of a sieve generated by a presieve in \mathcal{B} contains a presieve in \mathcal{B} , and

⁴Diagrams satisfying this limit condition are usually called *continuous* diagrams but we drop the adjective since we will not consider more general cotree-indexed diagrams.

- every J -covering sieve S' over C' contains a presieve constructed as the multicomposite of a κ -cotree in which the predecessors of each node form a presieve in \mathcal{B} , and conversely multicomposites of such κ -cotrees are J -covering.


Definition 3.4.4 (κ -woven sites). A site (\mathcal{C}, J) is κ -**woven** if J has a κ -warp such that $\mathcal{B}(C)$ contains at most κ presieves for every C .

The following is a strict improvement of Theorem 3.3.9, reducing to it in the case $\kappa = \omega$. We take inspiration from the approach of Espindola and Kanalas in [EK23, Thms 2.4, 2.5].

Theorem 3.4.5. Let κ be a regular cardinal. Let (\mathcal{C}, J) be a κ -woven site where \mathcal{C} is a category with κ -small connected limits⁵. Then assuming κ -dependent choice and the law of excluded middle, $\text{Sh}(\mathcal{C}, J)$ has enough κ -points (points preserving κ -small limits).

Proof. Let $C_0 \in \mathcal{C}$. We construct a tree of height κ in \mathcal{C} with root C_0 inductively as follows. Let $\xi : \kappa \times \kappa \rightarrow \kappa$ as usual. The predecessors of a node C at height γ^+ (with $\xi^{-1}(\gamma) = (\alpha, \beta)$) form the $\mathcal{B}(C)$ -presieve⁶ contained in $n_{\gamma, \alpha}^*(S_\beta)$, where $n_{\gamma, \alpha}$ is the (possibly transfinite) composite of the morphisms in the diagram from C to its successor C' at height α and S_β is the β th presieve in $\mathcal{B}(C')$. The nodes at a limiting height $\lambda < \kappa$ are the limits in \mathcal{C} of the branches in the tree below them.

The colimit in $\text{pt}(\text{PSh}(\mathcal{C}))$ of any branch of this cotree constitutes a κ -point of $\text{Sh}(\mathcal{C}, J)$. Indeed, any element of $p^*(\mathfrak{L}(C'))$ is represented by an element $w \in \mathfrak{L}(C)^*(\mathfrak{L}(C'))$, and by assumption any J -covering sieve on C' can be presented as the multicomposite of a $<\kappa$ -cotree in \mathcal{B} . We can verify projectivity by inducting up the $<\kappa$ -cotree: supposing that we have shown that there is $\mathfrak{L}(C_\beta)$ in the branch which is projective with respect to the subcotree truncated to height β . This selects a branch in the truncated cotree:



We can pull back the \mathcal{B} -presieve of predecessors of this node; by construction of the original diagram, there exists C_{β^+} in the branch which completes a square with one of the members of this presieve. At a limit ordinal λ , the limit of the objects $\{C_\beta\}_{\beta < \lambda}$ must complete a square with the limit of the corresponding branch of the cotree on the right-hand side, by the universal property of the latter. With this, we are done by Lemma 2.2.6.

Now given $x \neq y \in \mathfrak{L}(C_0)(j_*(X))$, by the assumption that the multicomposites of $<\kappa$ -cotrees are J -covering, we conclude that at least one of the branches in the diagram must consist of an improvement of $\mathfrak{L}(C_0)$ for (x, y, m, w) (for any J -covering sieve m and element w). Otherwise, in each branch of length κ there is some $\lambda < \kappa$ such that the composite from height λ down to C_0 equalizes x, y , and

⁵Equivalently with pullbacks and limits of cochains of size $<\kappa$.


⁶Should we care to complete this construction within the constraints of κ -dependent choice, we must assume that assignment from the pullback of a \mathcal{B} -presieve to a \mathcal{B} -presieve contained within it is functorial, hence “the $\mathcal{B}(C)$ -presieve”.

these assemble into a κ -cotree that fails to distinguish x from y , contradicting the hypothesis that the multicomposite of such a tree in $\mathcal{B}(C_0)$ must be J -covering. \square

Example 3.4.6 (An application of Theorem 3.4.5 where Theorem 2.1.5 fails). Consider a category \mathcal{C} constructed as follows:

- Objects are of the form C_w where w is a finite or ω -indexed binary word, plus a further object D ;
- In addition to identity morphisms, there is a morphism $C_w \rightarrow C_{w'}$ whenever w' is a finite prefix of w , so the C_w form an ω^+ cotree with branches of length ω^+ ;
- There is a pair of morphisms $C_w \rightrightarrows D$ whenever w is finite, or infinite and eventually constant at 1, respected by the previous class of morphisms;
- There is a single morphism $C_w \rightarrow D$ for all other infinite w .

Part of the structure of this category is sketched here:



We can equip \mathcal{C} with the Grothendieck topology J having as ω_1 -warp \mathcal{B} the pairs $C_{w0} \rightarrow C_w \leftarrow C_{w1}$. This makes the ‘minimal’ sieve of morphisms into C_w from objects indexed by infinite words J -covering. Moreover, the parallel pair of morphisms $C_\epsilon \rightrightarrows D$ remains distinct in $\text{Sh}(\mathcal{C}, J)$, since composition with the morphism from any ‘eventually 1’ objects distinguishes them. If we attempt to construct a point of $\text{Sh}(\mathcal{C}, J)$ starting at a representable point $\mathfrak{L}(C_w)$ of $\text{PSh}(\mathcal{C})$ using linear induction as in the proof of Theorem 2.1.5, we get stuck at the first limit ordinal, since there is no guarantee that the limit (taken in \mathcal{C}) of a cochain of improvements will be an improvement, as witnessed by the leftmost branch starting from any C_w . Yet there does exist an improvement in the minimal sieve, namely the morphism coming from $C_{w111\dots}$; we know how to find it explicitly here by construction of \mathcal{C} , but we relied on the law of excluded middle to construct it in the proof of Theorem 3.4.5 so it may not be possible to construct it in general.

3.4.7 (Returning to induction). A stronger hypothesis on the category \mathcal{C} (independent of the choice of Grothendieck topology J) that allows us to recover the inductive construction is to ask that the canonical map

$$(10) \quad \text{colim}_{\alpha < \lambda} \mathcal{C}(C_\alpha, C) \rightarrow \mathcal{C}(\lim_{\alpha < \lambda} C_\alpha, C)$$

is injective for all C and cochains indexed by $\lambda < \kappa$. Example 3.4.6 does not satisfy that hypothesis; we leave an elaboration of the proof of the corresponding weaker result to the dedicated reader. However, there is a further convenient consequence of imposing this property.

Corollary 3.4.8 (Presheaf collapse). In the assumptions of Theorem 3.4.5, if \mathcal{C} has limits of chains of length equal to κ satisfying (10) then $\text{Sh}(\mathcal{C}, J)$ is equivalent to a presheaf topos.

Proof. Under these stronger hypotheses, the point constructed in the proof of Theorem 3.4.5 embeds into the representable point obtained by instead taking the limit in \mathcal{C} , so we have a representable point which restricts to $\text{Sh}(\mathcal{C}, J)$. This amounts to saying that the corresponding representable presheaves are J -sheaves, and hence tiny objects in $\text{Sh}(\mathcal{C}, J)$. That there are *enough* points of this form precisely means that these tiny objects form a generating family for $\text{Sh}(\mathcal{C}, J)$, and thus the topos is of presheaf type. In fact, we can deduce that $\text{Sh}(\mathcal{C}, J) \simeq \text{PSh}(\mathcal{C}')$, where \mathcal{C}' is the full subcategory of \mathcal{C} on objects C such that $\mathcal{J}(C)$ is a J -sheaf. \square

Corollary 3.4.9 (Esp ndola and Kanalas, [EK23, Thm. 5.4]). Let κ be a regular cardinal. Let (\mathcal{C}, J) be a site where $\mathcal{C}(C, C')$ has cardinality $\leq \kappa$ for all C, C' , \mathcal{C} has κ -small limits and J is generated under pullbacks and multicomposites of $< \kappa$ -cotrees by at most κ sieves. Then $\text{Sh}(\mathcal{C}, J)$ has enough κ -points.

Proof. Just as in Corollary 3.3.10, the given data is enough to deduce that we have a κ -warp of J obtained by taking pullbacks of the generating sieves. Thus we conclude the result by applying Theorem 3.4.5. \square

Remark 3.4.10 (Connection to forcing). As noted by Esp ndola and Kanalas in [EK23] and independently pointed out to us by Dianthe Basak, the hypotheses of Theorem 3.3.9 and its proof resemble those of the Rasiowa–Sikorski Lemma. There is a similar resemblance between Theorem 3.4.5 and Martin’s Axiom, although the hypotheses ensure that we need only appeal to choice in the construction and no such stronger (or logically independent) axiom need be invoked. This suggests a relationship between constructibility of points and extensions of ZFC that we do not have space to explore in depth here.

3.4.11 (Large cardinals and posetal collapse). In [Esp20], Esp ndola invokes the condition $\kappa^{<\kappa} = \kappa$ on cardinals, which is shorthand for the existence of a bijection $\kappa^\lambda \cong \kappa$ for any $\lambda < \kappa$. This condition is implied by inaccessibility. While extensive use is made of this condition, its negation actually has strong consequences if we impose constraints on the size of hom-sets in \mathcal{C} . This illustrates how removing size constraints in Theorem 3.4.5 may represent a significant improvement of the result of Esp ndola and Kanalas, although compelling concrete examples are harder to come by.

Proposition 3.4.12 (Posetal collapse for cardinals $\kappa < \kappa^{\kappa}$). Suppose (\mathcal{C}, J) is a site satisfying the conditions of Corollary 3.4.9 but that there exists $\lambda < \kappa$ such that $\kappa^\lambda > \kappa$. Then \mathcal{C} is a poset.

Proof. This is a version of the classic argument of Freyd⁷, [Fre66, Chapter 3.D]. First, since κ is not inaccessible, there exists λ' with $2^{\lambda'} \geq \kappa$. Given a parallel pair of arrows $f, g: A \rightrightarrows B$ in \mathcal{C} , if $f \neq g$ then there must be $2^{\lambda'} \geq \kappa$ morphisms from A to $B' := \prod_{\lambda'} B$ in \mathcal{C} . Now taking λ as in the statement, we conclude that there are at least $\kappa^\lambda > \kappa$ morphisms from A to $\prod_\lambda B'$, a contradiction. Thus we must have had $f = g$, as required. \square

3.5. Spatial locales. In the previous sections, we have discussed a number of assumptions on a site (\mathcal{C}, J) ensuring that the topos of sheaves on this site will have enough points. Some of these results have a logical interpretation: we know that coherence is linked to first order logic, and that *separability* is somehow connected to Fourman–Grayson’s completeness theorem. *What about their geometric meaning?* This subsection aims to shed some light on this topic, both to provide a more geometric intuition on this theory, and to bring some clarity where the literature seems to have created some confusion. Let us start with this last point.

⁷Citation is not endorsement.

In this subsection, when we mention a notion from general topology, we shall refer to [Mun14, Sie20, SSS78].

Definition 3.5.1 (Gradations of separability [Sie20, Chap. III]). According to the classics, a topological space is:

- (1) separable if it has a countable dense subset,
- (2) first countable, if each point has a countable neighbourhood basis,
- (3) second countable, or completely separable, if it has a countable case.

It is straightforward to prove that $3 \Rightarrow 1 \wedge 2$.

Proposition 3.5.2 (A site description of second countability). Let X be a topological space. The following are equivalent:


- X is second countable.
- $\text{Sh}(X)$ has a countable (posetal) site with finite limits.
- $\text{Sh}(X)$ has a site with countably many objects.

Proof. If X is second-countable, the *basis* B – in the classical sense of general topology – is exactly a lex subposet $B \hookrightarrow \mathcal{O}(X)$ that is dense.

$2 \Rightarrow 3$ is immediate. Finally, if $\text{Sh}(X)$ has a countable site, then the supports of representable sheaves form a generating family for $\text{Sh}(X)$, which is a *prebasis* for X (a subposet of $\mathcal{O}(X)$ whose closure under finite intersections is a basis). Since the closure of such a prebasis under finite intersections remains countable, X is second countable. \square

3.5.3 (Comparing naming conventions). Makkai and Reyes justified naming their class of topoi ‘separable’ because of a very special case, namely that the topos of sheaves over a Hausdorff space X is separable if and only if X is a complete separable metric space [MR06, Proposition 6.2.3]; we find this justification wanting. The classical notion of *separable* or *completely separable* space from general topology are much more general than the (very) special case they are inspired by and this choice creates confusion on how such a geometric object may behave (especially given the hypothesis of Hausdorffness).

In the diagram below we collected all the notions that appeared in this paper or that seem relevant to the discussion.



We chose to represent them in different columns because they ostensibly axiomatize different features of the space. Let us comment on them from a site-theoretic perspective.

- compactness and ω -compactness (often called Lindelöf) tell us that covering families have subfamilies of a certain size.
- second countability is a bound on the size of the site itself.
- woven-ness puts a bound on the complexity of the class of covering families. Notice that this is *a priori* very different from any form of compactness, even though it may seem to have a similar flavour.
- Makkai-Reyes separability is a mix of woven-ness and second countability.
- separability and first countability are strange notions from a locale-theoretic perspective, because they refer to points.

Of course, it happens that these columns do get mixed. For example, second countability is such a strong requirement that it implies ω -compactness, precisely because it cuts down the number of opens quite drastically, and thus forces some compactness-like behaviour into the space.

We finish this subsection by recalling a result due to Heckmann which is a further consequence of 3.3.9. Heckmann's proof provides a new insight into the topic via a mild generalization of the Baire category theorem to prove his main theorem [Hec15, Corollary 3.15].

Corollary 3.5.4 (Heckmann). A locale L whose frame of opens $\mathcal{O}(L)$ can be presented with countably many relations (inequalities between upward-closed subsets of $\mathcal{P}_{\text{fin}}(G)$) is spatial.

Proof. Being presented by countably many relations means that there is an inclusion $j : \text{Sh}(L) \rightarrow \text{PSh}(\mathcal{P}_{\text{fin}}(G))$. The relations corresponds to inclusions of subterminal subobjects $\text{PSh}(\mathcal{P}_{\text{fin}}(G))$ which generate the Grothendieck topology on $\mathcal{P}_{\text{fin}}(G)$ corresponding to j , and by hypothesis this generating family is countable. For each finite subset $U \in \mathcal{P}_{\text{fin}}(G)$, we obtain a countable warp of covers of U by pulling back these covers; this produces only countably many sieves over U because the subterminal objects form poset. Thus we may apply Theorem 3.3.9. \square

4. EXTENDING THE CLASS OF TOPOI WITH ENOUGH POINTS

4.1. Eliminating pullbacks. Depending on one's disposition, the requirement that the categories featuring in the sites of the last section should have pullbacks may seem rather demanding. We can eliminate this requirement from Theorem 3.4.5 (and hence from Theorem 3.3.9) by extending a site without pullbacks to an equivalent one having them.

Theorem 4.1.1. Let κ be a regular cardinal. Let (\mathcal{C}, J) be a κ -woven site where \mathcal{C} has limits of cochains of size $< \kappa$ (but need *not* have pullbacks). Then assuming κ -dependent choice and the law of excluded middle, $\text{Sh}(\mathcal{C}, J)$ has enough κ -points.

Proof. Let $\bar{\mathcal{C}}$ be the closure of \mathcal{C} in $\text{PSh}(\mathcal{C})$ under pullbacks. The Grothendieck topology \bar{J} on $\bar{\mathcal{C}}$ such that $\text{Sh}(\bar{\mathcal{C}}, \bar{J}) \simeq \text{Sh}(\mathcal{C}, J)$ has a κ -warp extending the one for J by adding to each of the objects in $\bar{\mathcal{C}}$ the presieve of all morphisms from objects of \mathcal{C} (this generates the maximal sieve for objects in \mathcal{C}). It is easily checked that the pullback of such a presieve contains the corresponding presieve and that every \bar{J} -covering sieve is the multicomposite of a κ -cotree in which the first layer is this presieve. As such, we may apply Theorem 3.4.5 to deduce the result. \square

4.1.2 (Pullbacks resist being eliminated altogether). It is not possible to eliminate pullbacks quite as directly in Proposition 3.2.2 due to the same obstacle that led us to Theorem 3.3.9. Namely, as soon as we reach the limit ordinal ω in the induction of Theorem 2.1.5, we lose the ability to improve points with respect to infinite sieves. With some more work, we will be able to improve Deligne's original result, though.

Careful analysis of Deligne's proof leads one to conclude that we only ever use the universal property of pushouts in $\text{pt}(\mathcal{E})$ to deduce 'functorial amalgamation'.

Definition 4.1.3 (Functorial amalgamation). We say that \mathcal{C} has **functorial amalgamation** if there merely exists a functor $\text{amalg} : \mathcal{C}^{\Gamma} \rightarrow \mathcal{C}$ together with a natural transformation $\text{id}_{\mathcal{C}^{\Gamma}} \Rightarrow \Delta_{\mathcal{C}} \circ \text{amalg}$.

Functorial amalgamation is a significant weakening of the requirement of pushouts, where even more than the universal property of the pushout is lost: for example, we do not impose that the amalgamation of a cospan where both maps are identities at

an object C returns C itself. Nonetheless, the structure of the proof of Proposition 3.1.1 can still be applied to deduce the following corollary.

Corollary 4.1.4 (Functorial amalgamation in $\text{Ind}(\mathcal{C})$). Let \mathcal{C} be any small category. The following are equivalent:

- (1) \mathcal{C} has functorial amalgamation.
- (2) $\text{Ind}(\mathcal{C})$ has a functorial amalgamation such that amalg preserves finitely presentable objects.

Proof. The proof that (1) \Rightarrow (2) proceeds exactly as it did in Proposition 3.1.1.

To deduce (2) \Rightarrow (1), observe that the hypothesis provides functorial amalgamation in the idempotent completion of \mathcal{C} (which is equivalent to the category of finitely presentable objects of $\text{Ind}(\mathcal{C})$). Every object in the idempotent completion is canonically a retract of one in \mathcal{C} , so given the amalgam of a cospan in \mathcal{C} taken in the idempotent completion, we can extend by the section of the retraction map to obtain an amalgam in \mathcal{C} . \square


Remark 4.1.5 (Eliminating idempotent completeness). While $\text{Ind}(\mathcal{C})$ having pushouts only implies that the idempotent completion of \mathcal{C} has pushouts by Proposition 3.1.1, we can deduce from Corollary 4.1.4 that this still forces \mathcal{C} to have functorial amalgamation.

Example 4.1.6 (Adhesiveness provides functorial amalgamation). An interesting example of functorial amalgamation is the following. Suppose \mathcal{C} has pushouts and let $\mathcal{C}_{\mathcal{M}} \rightarrow \mathcal{C}$ be a wide subcategory (identified with its class of maps $\mathcal{M} \subset \mathcal{C}^{\rightarrow}$) such that \mathcal{M} is stable under pushouts. This means precisely that in the diagram below, we can construct the dashed functor, which will automatically be a functorial amalgamation. Yet, in full generality this construction will not provide pushouts for $\mathcal{C}_{\mathcal{M}}$, since we do not demand that the pushout comparison map belongs to \mathcal{M} .

$$\begin{array}{ccc} \mathcal{C}_{\mathcal{M}}^{\Gamma} & \xrightarrow{i^{\Gamma}} & \mathcal{C}^{\Gamma} \\ \downarrow & & \downarrow \\ \mathcal{C}_{\mathcal{M}} & \xrightarrow{i} & \mathcal{C} \end{array}$$

Some concrete instances of this construction are as follows:

- (1) The simplest is given the case where \mathcal{C} is Fin , the category of finite sets, and $\mathcal{C}_{\mathcal{M}}$ is its wide subcategory of monomorphisms. The induced amalgamation is not a pushout, since (using n to denote an n -element set), the dashed arrow in the following diagram in Fin is not a monomorphism.



- (2) More generally, monomorphisms are stable under pushout in any *adhesive category* by definition, so we can perform this construction.

Definition 4.1.7 (Functorial coamalgamation). We say that a category \mathcal{C} has **functorial coamalgamation** if \mathcal{C}^{op} has functorial amalgamation. We say that a site (\mathcal{C}, J) has **functorial coamalgamation** if this is true for \mathcal{C} and for a J -covering

presieve $\{f_i : C_i \rightarrow C\}_{i \in I}$ and any morphism $h : D \rightarrow C$, the coamalgamation squares:

$$(11) \quad \begin{array}{ccc} D_i & \xrightarrow{h_i} & C_i \\ g_i \downarrow & & \downarrow f_i \\ D & \xrightarrow{h} & C \end{array}$$

are such that $\{g_i : D_i \rightarrow D\}_{i \in I}$ forms a J -covering presieve. (If the coamalgamation squares are pullbacks, this is automatic).

Proposition 4.1.8 (Constructing improvements without pullbacks). Let Φ be a class of limit diagrams. Consider the inclusion $j : \text{Sh}(\mathcal{C}, J) \rightarrow \text{PSh}(\mathcal{C})$ where (\mathcal{C}, J) be a site with functorial coamalgamation in the sense of Definition 4.1.7 and J -covering sieves are generated by presieves in Φ^{op} . Let p be a point of $\text{PSh}(\mathcal{C})$ such that p^* preserves Φ -limits and $x \neq y \in p^*(j_*(X))$. Then p has an improvement p' with respect to (x, y, m, w) for all j -dense $m : S \rightarrow T$ and $w \in p^*(T)$ such that p'^* also preserves Φ -limits.

Proof. Checking the proof of Proposition 3.2.2, we see that everything works identically for a site with functorial coamalgamation after replacing the pullbacks $v^*(f_i)$ with the coamalgams g_i of (11). \square

Corollary 4.1.9 (Deligne without pullbacks). Let (\mathcal{C}, J) be a site with functorial coamalgamation in which every J -covering sieve contains a finite covering presieve; this corresponds to taking Φ to be the class of finite diagrams in Proposition 4.1.8. Then $\text{Sh}(\mathcal{C}, J)$ has enough points.

Remark 4.1.10 (Not all locally finitely presentable topoi). The topoi which admit a site (\mathcal{C}, J) such that every J -covering sieve contains a finite covering presieve are exactly the *locally finitely generated* topoi (or *compactly generated* topoi in the terminology of [Rog21]): those which have a generating set of finitely generated (also called ‘quasi-compact’) objects. The class of locally finitely generated topoi notably includes the locally finitely presentable ones.

We cannot hope to completely remove functorial coamalgamation from Corollary 4.1.9, since any atomic topos is locally finitely generated, and there are well-known examples of atomic topoi without points, such as the classifying topos for the theory of uncountable complete binary trees due to Malitz, see [Joh02b, Example D3.4.14].

On the other hand, this does not *a priori* determine whether or not there exists a locally finitely presentable topos without enough points (assuming the axiom of choice); see Remark 4.1.11.

Remark 4.1.11. Since the preprint version of the present paper appeared, Jérémie Marquès claims to have shown that the example of an atomic topos with no points [Joh02b, Example D3.4.14] is actually locally finitely presentable [Mar24]. It would follow that being locally finitely presentable is not sufficient to admit *any* points at all, even assuming the axiom of choice!

In spite of this, we can extend Corollary 3.3.1 at least a little by constraining the local size of the site.

Corollary 4.1.12. A locally finitely presentable topos whose subcategory of finitely presentable objects has essentially countable slice categories has enough points.

Proof. Locally finitely presentable toposes are precisely those admitting a site (\mathcal{C}, J) in which J is generated by sieves having finite final presieves⁸, which corresponds to

⁸This is essentially the content of [DLRG20, Prop. 5.5], although the condition of finality of the presieve was forgotten in the formulation of the theorem there.

taking Φ to be the class of finite diagrams in Proposition 4.1.8 but without assuming functorial coamalgamation. The stated hypothesis ensures that we can choose a site in which there are at most countably many morphisms with a given object as codomain. This means that there are at most countably many finite J -covering families over each object, whence the collection of all such is an ω -warp for J and (\mathcal{C}, J) satisfies the hypotheses of Theorem 4.1.1. \square

Remark 4.1.13 (Not quite Hofmann-Lawson). It is worth noting that if we specialize Corollary 4.1.12 to the localic case, we do not directly recover the well-established result that continuous locales have enough points, proved by Hofmann and Lawson [HL78]. While parts of the proof have a similar flavour to improvements, they critically use Zorn’s lemma in a way that cannot be directly generalized beyond frames, so cannot be presented as a variant of Deligne’s strategy.

4.1.14 (Morphisms and comorphisms of sites). Of course, if a topos \mathcal{G} admits a geometric surjection from a topos with enough points, then it too has enough points. From a site-theoretic perspective, we can construct geometric morphisms using either morphisms or comorphisms of sites. Caramello established necessary and sufficient conditions for when the resulting geometric morphisms are surjections [Car19]. In principle one could characterize the sites (\mathcal{D}, J) admitting a surjection-inducing morphism of sites to (resp. a surjection-inducing comorphism of sites from) a site (\mathcal{C}, K) satisfying the conditions of Theorem 4.1.1 or Corollary 4.1.9. Lacking compelling examples, we leave that effort to future work.

4.2. Zooming out: the 2-category of topoi. There is a more abstract perspective on the construction performed in the proof of Theorem 2.1.5. We considered a parallel pair of morphisms $f, g : Y \rightrightarrows X$ in \mathcal{F} and a point p such that $p^* j_* f \neq p^* j_* g$; from these we constructed a point p' of \mathcal{F} equipped with a morphism $p \Rightarrow p'; j$ in $\text{pt}(\mathcal{E})$:

$$\begin{array}{ccc} & & \mathcal{F} \\ & \nearrow^{p'} & \downarrow f \\ \text{Set} & \xrightarrow{p} & \mathcal{E}. \end{array}$$

However, we can more properly summarize the result by presenting a collection of ‘enough’ points of \mathcal{E} as a geometric morphism from Set^K for some set K . Thus we construct a square:

$$\begin{array}{ccc} \text{Set}^{K'} & \xrightarrow{P'} & \mathcal{F} \\ \text{Set}^\pi \downarrow & \nearrow & \downarrow f \\ \text{Set}^K & \xrightarrow{P} & \mathcal{E} \end{array}$$

where on the left π is the map sending each point constructed in Theorem 2.1.5 to the one it was constructed from. In general we will need many points p' for each point p , but under some circumstances at most one is needed, as the following result shows.

Theorem 4.2.1. The class of topoi with enough points is closed under closed inclusions. That is, letting $f : \mathcal{F} \rightarrow \mathcal{E}$ be a closed inclusion such that \mathcal{E} has enough points, it follows that \mathcal{F} has enough points.

Proof. Observe that surjections are stable by pullback along closed inclusions. Indeed, consider a pullback square of geometric morphisms

$$\begin{array}{ccc} \mathcal{G} & \xrightarrow{q} & \mathcal{F} \\ g \downarrow & \lrcorner & \downarrow f \\ \mathcal{S} & \xrightarrow{p} & \mathcal{E}, \end{array}$$

in which p is a surjection and f is a proper inclusion. By Theorem [MV00, I.6.1], the Beck-Chevalley condition $p^*f_* \cong g_*q^*$ holds. Since p^* and f_* are faithful by assumption, this forces q^* to be faithful, and hence q is a surjection. Moreover, closed inclusions are stable under pullback, so g is a closed inclusion too. In particular, considering the case where \mathcal{E} has enough points, we may take \mathcal{S} to be Set^K for K a sufficient set of points, at which point \mathcal{F} admits a surjection from a (closed) subtopos of Set^K , which is necessarily of the form $\text{Set}^{K'}$ for some subset $K' \subseteq K$. Hence \mathcal{F} has enough points. \square

Remark 4.2.2 (Johnstone proved the localic case). One might compare the theorem with [Joh02a, C1.2.6(b)], which is the localic case of this result. Johnstone remarks there that closed sublocales of spatial locales being spatial is a non-constructive result, and it is worth isolating the non-constructive aspects of our proof, which lies in the fact that a subtopos of the copower Set^K is necessarily of the form $\text{Set}^{K'}$. For a more generic topos \mathcal{S} , a closed subtopos of \mathcal{S}^K will instead be a K -indexed coproduct of closed subtopoi of \mathcal{S} . Thus we arrive at the following more general result.

Corollary 4.2.3 (Relative version). Let \mathcal{S} be any base topos. Given a topos \mathcal{E} over \mathcal{S} which has enough \mathcal{S} -valued points, any closed subtopos \mathcal{F} of \mathcal{E} has enough (\mathcal{S} -valued) ‘closed sub-points’, meaning that the collection of geometric morphisms coming from closed subtopoi of \mathcal{S} have jointly conservative inverse image functors.

4.2.4 (Beck-Chevalley is the key). The essential feature exploited in the proof of Theorem 4.2.1 is the Beck-Chevalley condition. Indeed, we only needed the *weak* Beck-Chevalley condition, where we merely have a componentwise monic transformation, which amounts to g being **proper** by [Joh02b, Corollary C3.2.22]. For subtopoi this adds nothing to Theorem 4.2.1, however, since an inclusion is proper if and only if it is tidy, if and only if it is closed.

We can also attain the weak Beck-Chevalley condition by imposing dual hypotheses, but we shall see that this doesn’t tell us quite so much.

Lemma 4.2.5 (Open points). Suppose that \mathcal{E} has enough *open* points, which amounts to the existence of an open surjection $\text{Set}^K \rightarrow \mathcal{E}$ for some set K . Then any subtopos of \mathcal{E} also has enough open points.

Proof. The argument proceeds identically to Theorem 4.2.1, this time deducing the weak Beck-Chevalley condition from the fact that the geometric morphism $\mathcal{S} \rightarrow \mathcal{E}$ is open by [Joh02b, Theorem C3.1.27]. We again rely on the fact that every subtopos of Set^K is of the form $\text{Set}^{K'}$. \square

This result is redundant since by [Joh02b, Lemmas C3.5.1, C3.5.2] a topos admitting an open surjection from Set^K is necessarily Boolean, so all of its subtopoi are closed and we may thus apply Theorem 4.2.1 to them. Its extension to general topoi may be non-trivial, however.

Corollary 4.2.6. Let \mathcal{S} be any base topos. Given a topos \mathcal{E} over \mathcal{S} which has enough open \mathcal{S} -valued points, any subtopos \mathcal{F} of \mathcal{E} has enough (\mathcal{S} -valued) sub-points, meaning that the collection of geometric morphisms coming from subtopoi of \mathcal{S} have

jointly conservative inverse image functors (and in fact the sub-points which are open suffice).

Remark 4.2.7 (Not quite Deligne). Theorem 4.2.1 is not as general as it might seem at first sight. Indeed, the inclusion of a coherent topos into the category of presheaves on its coherent objects is seldom a closed inclusion, although it preserves filtered colimits (and so is *relatively tidy* in the sense of Moerdijk and Vermeulen, [MV00]). Conversely, since we have not been able to verify that every locally finitely presentable topos has enough points, we do not know whether the relativised weak Beck-Chevalley condition that applies to relatively tidy morphisms via the construction of the ‘comma topos’ is sufficient to witness the existence of enough points of relatively tidy subtopoi more generally.

Remark 4.2.8 (A logical interpretation). Theorem 4.2.1 is easier to interpret than Deligne’s theorem from a logical point of view. Following [Car18, 4.2.2.2], let \mathbb{E} be a geometric theory with enough set models⁹, and let ψ be a geometric propositional formula (a geometric sequent with no free variables formed from finitary conjunctions, arbitrary disjunctions and existential quantification). Then the quotient theory $\mathbb{E} \cup (\psi \vdash \perp)$ has enough Set-models too. In other words, adding a finite number of negated propositional formulas to a geometric theory cannot break the existence of enough models.

ACKNOWLEDGEMENTS

The authors started working on this paper in late 2021, and this project has been a collection of ups and downs, attempts and failures. We are indebted to our friends and colleagues *Axel Osmond* and *Joshua Wrigley* who have encouraged us throughout the journey, as well as colleagues in our respective institutions who have patiently listened to vague outlines of the work during its long gestation. A significant portion of this research was carried out during two invitations: Ivan Di Liberti was invited by Université Sorbonne Paris Nord (Paris) in 2022 and Morgan Rogers was invited by Stockholm University (Stockholm) in 2023. Both authors are grateful to these institutions for the financial support that made the research visits possible.

REFERENCES

- [BDS06] N Bourbaki, Pierre Deligne, and Bernard Saint-Donat. *Théorie des Topos et Cohomologie Etale des Schémas. Séminaire de Géométrie Algébrique du Bois-Marie 1963-1964 (SGA 4): Tome 2*, volume 270. Springer, 2006.
- [BGH18] Clark Barwick, Saul Glasman, and Peter Haine. Exodromy. *arXiv preprint arXiv:1807.03281*, 2018.
- [Cam] Tim Champion. Does every exponentiable (∞ -)topos have enough points? MathOverflow. URL:<https://mathoverflow.net/q/401429> (version: 2021-08-11).
- [Car18] Olivia Caramello. *Theories, sites, toposes: relating and studying mathematical theories through topos-theoretic ‘bridges’*. Oxford University Press, 2018.
- [Car19] Olivia Caramello. Denseness conditions, morphisms and equivalences of toposes. *arXiv preprint arXiv:1906.08737*, 2019.
- [DLa] Ivan Di Liberti. A list of locally finitely presentable topoi that are not coherent. MathOverflow. URL:<https://mathoverflow.net/q/347107> (version: 2020-06-15).
- [DLb] Ivan Di Liberti. A list of proofs of ‘coherent topoi have enough points’. MathOverflow. URL:<https://mathoverflow.net/q/338862> (version: 2020-06-15).
- [DL22] Ivan Di Liberti. The geometry of coherent topoi and ultrastructures. *arXiv preprint arXiv:2211.03104*, 2022.
- [DLRG20] Ivan Di Liberti and Julia Ramos González. Gabriel–ulmer duality for topoi and its relation with site presentations. *Applied Categorical Structures*, 28(6):935–962, 2020.

⁹Which means that $\mathbb{E} \vdash \phi$ if and only if ϕ is true in all Set-valued models.

- [EK23] Christian Espíndola and Kristóf Kanalas. Every theory is eventually of presheaf type. *arXiv preprint arXiv:2312.12356*, 2023.
- [Esp19] Christian Espíndola. Infinitary first-order categorical logic. *Annals of Pure and Applied Logic*, 170(2):137 – 162, 2019.
- [Esp20] Christian Espíndola. Infinitary generalizations of deligne’s completeness theorem. *The Journal of Symbolic Logic*, 85(3):1147–1162, 2020.
- [FG82] Michael P Fourman and Robin J Grayson. Formal spaces. In *Studies in Logic and the Foundations of Mathematics*, volume 110, pages 107–122. Elsevier, 1982.
- [Fre66] Peter Freyd. *Abelian Categories*. Harper & Row and John Wetherhill Inc., 1966.
- [Gin] Harry Gindi. Deligne’s theorem for n -topos. MathOverflow. URL:<https://mathoverflow.net/q/371473> (version: 2020-09-12).
- [Hec15] Reinhold Heckmann. Spatiality of countably presentable locales (proved with the baire category theorem). *Mathematical Structures in Computer Science*, 25(7):1607–1625, 2015.
- [Hen23] Simon Henry. When does $\text{Ind}_\kappa(C^I) \simeq \text{Ind}_\kappa(C)^I$? *arXiv preprint arXiv:2307.06664*, 2023.
- [HL78] Karl H Hofmann and Jimmie D Lawson. The spectral theory of distributive continuous lattices. *Transactions of the American Mathematical Society*, 246:285–310, 1978.
- [Joh77] Peter T. Johnstone. Topos theory, volume 10 of. *London Mathematical Society Monographs*, 1977.
- [Joh02a] Peter T. Johnstone. *Sketches of an elephant: a Topos theory compendium*. Oxford logic guides. Oxford Univ. Press, New York, NY, 2002.
- [Joh02b] Peter T. Johnstone. *Sketches of an Elephant: A Topos Theory Compendium: 2 Volume Set*. Oxford University Press UK, 2002.
- [Law75] F William Lawvere. Continuously variable sets; algebraic geometry= geometric logic. In *Studies in Logic and the Foundations of Mathematics*, volume 80, pages 135–156. Elsevier, 1975.
- [LM94] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in Geometry and Logic*. Springer New York, 1994.
- [Lur04] Jacob Lurie. *Derived algebraic geometry*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [Mar24] Jérémie Marquès. The Malitz–Gregory pointless atomic topos is locally finitely presentable. *Unpublished, private communication*, 2024.
- [MR06] Michael Makkai and Gonzalo E Reyes. *First order categorical logic: model-theoretical methods in the theory of topoi and related categories*, volume 611. Springer, 2006.
- [mu] mathoverflow user1034242. Topos with enough points but not coherent. MathOverflow. URL:<https://mathoverflow.net/q/417780> (version: 2022-03-09).
- [Mun14] James Munkres. *Topology, Second Edition*. Pearson Education Limited, 2014.
- [MV00] Ieke Moerdijk and Jacob Johan Caspar Vermeulen. *Proper maps of toposes*, volume 705. American Mathematical Soc., 2000.
- [Rey74] Gonzalo Reyes. From sheaves to logic. *Studies in Algebraic Logic, A. Daigneault, ed. MAA Studies in Math*, 9:143–134, 1974.
- [Rog21] Morgan Rogers. On supercompactly and compactly generated toposes. *arXiv preprint arXiv:2101.04015*, 2021.
- [Rog24] M. Rogers. Topological endomorphism monoids of models of geometric theories. *Theory and Applications of Categories*, 2024.
- [Sie20] Waclaw Sierpinski. *General topology*. Courier Dover Publications, 2020.
- [SSS78] Lynn Arthur Steen, J Arthur Seebach, and Lynn A Steen. *Counterexamples in topology*, volume 18. Springer, 1978.

IVAN DI LIBERTI

DEPARTMENT OF PHILOSOPHY, LINGUISTICS AND THEORY OF SCIENCE
 UNIVERSITY OF GOTHENBURG
 GOTHENBURG, SWEDEN
diliberti.math@gmail.com

MORGAN ROGERS:

LABORATOIRE D’INFORMATIQUE PARIS NORD
 UNIVERSITÉ SORBONNE PARIS NORD
 PARIS, FRANCE
rogers@lipn.univ-paris13.fr