Adult help needed: factoring geometric morphisms

Let's define a *finite geometric morphism* (an *fgm*) like this.
A functor $f : A \to B$ "is" an fgm iff both A and B are finite categories.
If $f : A \to B$ is an fgm, then $(\circ f) : \mathbf{Set}^B \to \mathbf{Set}^A$ has both adjoints:

$$\mathsf{Lan}_f \dashv (\circ f) \dashv \mathsf{Ran}_f$$

An fgm $f : A \to B$ induces an essential geometric morphism
$f : \mathbf{Set}^A \to \mathbf{Set}^B$,

$$f \equiv (f_! \dashv f^* \dashv f_*) := (\mathsf{Lan}_f \dashv (\circ f) \dashv \mathsf{Ran}_f)$$

When will use the expression "the fgm f" to refer either to $f : A \to B$ or to $(f_! \dashv f^* \dashv f_*)$, depending on the context.

Some terminology:
An fgm f is an *inclusion* when $(f^* \dashv f_*)$ is an inclusion.
An fgm f is a *surjection* when $(f^* \dashv f_*)$ is a surjection.
An fgm f is *dense* when $(f^* \dashv f_*)$ is dense.
An fgm f is *closed* when $(f^* \dashv f_*)$ is closed.

Look at these theorems from part A of Johnstone's "Sketches of an Elephant":
Theorem 4.2.10 Every geometric morphism can be factored, uniquely up to canonical equivalence, as a surjection followed by an inclusion.
Corollary 4.5.20 Any geometric inclusion $\mathcal{E}' \to \mathcal{E}$ has a unique factorization $\mathcal{E}' \to \mathcal{E}'' \to \mathcal{E}$, where $\mathcal{E}' \to \mathcal{E}''$ is dense and $\mathcal{E}'' \to \mathcal{E}$ is closed.

They say that every geometric morphism $\mathcal{E} \to \mathcal{E}'''$ can be factored as:

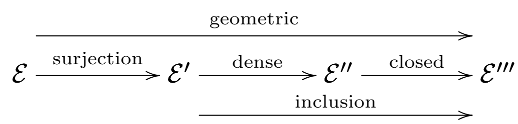

I would like to do something similar for fgms.
If $f : A \to D$ is an fgm, then (I believe that) it factors as:

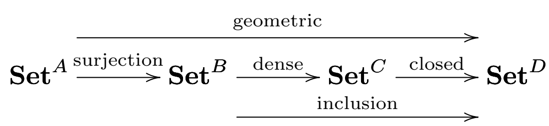

My guess was that $A \overset{\alpha}{\to} B \overset{\iota}{\to} D$ was the epi-monic factorization of the functor $f : A \to B$, and in $B \overset{\beta}{\to} C \overset{\gamma}{\to} D$ one map added just points, while the other one added just arrows (regarding A, B, C, D as graphs).

I asked a friend that I met in a conference a few months ago to check my hypothesis, and in the next day he showed me some diagrams, with very small A, B, C, D, and explained me why I was wrong.
I lost the diagrams he drew on that day (on a napkin, iirc)...
but I just sent an e-mail to him (= Peter Arndt).

Now I know how to do diagrams like the ones in the next page, and my intent is to use that to construct examples with fgms to clarify a good part of section A4 in the Elephant.

From: Peter Arndt
To: Eduardo Ochs
Date: 2015oct04

So, part one of your conjecture is correct (the way I read it). You write that the (surjective,inclusion) factorization of a geometric morphism is induced by the (epi,mono) factorization of the index categories. "Epi-mono" factorization is an uncommon name for functors - what you can use here, and what you probably meant, is the (essentially surjective, fully faithful) factorization for functors. This factorization is unique up to equivalence and can be realized concretely in different ways, e.g. by factoring the given functor through the full subcategory on the objects in the image (the second part of the factorization is then the inclusion of this full subcategory).

Proof that this gives the (surjection,inclusion) factorization of the presheaf categories (maybe you already told me this in Istanbul, I don't remember):

It is enough to check that the (essentially surjective, fully faithful) factorizationof the functor between the index categories induces a factorization into a surjectivegeometric morphism and an inclusion, since this factorization is essentially unique (Moerdijk/MacLane VII, Prop 4/Thm 6).

By definition a geometric morphism (f_*, f^*) is surjective if f^* is faithful. In our case f^* is precomposition. This is faithful for a functor f that is surjective on objects, since natural transformations are determined objectwise and all objects are in the image. So the first part of the factorization indeed yields a surjective geometric morphism.

For the second part we have to show that the essential geometric morphism (g_*, g^*) induced by a fully faithful functor g between the small index categories is an inclusion. By definition this means that g_* is fully faithful. But in our case g_* is right Kan extension along g, and (left or right) Kan extensions along fully faithful functors are fully faithful again.

(end of proof)

Note that this goes through for any index categories, not just preorders. Also note that it was not a priori clear that the middle topos ocurring in the (surjection,inclusion) factorization is a presheaf topos again, or that the two parts of the factorization are essential geometric morphisms again, but we have now seen that this is true.

However, I have some doubts about the corresponding statement for the (dense, closed) factorization.

Your second proposed factorization $B \to C \to D$ you described as "adding just points, then adding just arrows". This factorization probably doesn't exist for general functors. You can factorize a functor by first adjoining the objects from the target category which are not in the image. In the second step you can take care of the Hom-sets, but the maps betwen these will not in general be injective or surjective, i.e. it will not be just adding arrows, but also identifying a few old arrows. For your preorders this doesn't occur, of course, so your terminology makes sense (but it confused me a bit because I was thinking of general small categories).

For a set X let cX denote the constant presheaf with value X. That a geometric morphism (f_*, f^*) is dense, means that for a set X the unit map $cX \to f_* f^* cX$ is a mono. In our case, where the geometric morphism is induced by a functor $f : A \to B$ between small categories, f^* is precomposition with f, so $f^* cX$ is the constant presheaf on the domain category of the functor f. Again f_* is the right Kan extension along f. Looking at the pointwise construction of right Kan extension, by taking limits over some comma category, one can see that for an object b of B we have $f_* f^* cX(b) = X^{M(b)}$ where $M(b)$ is the set of connected components of the comma category $b \downarrow f$. The unit map $cX \to f_* f^* cX$ is objectwise given by the diagonal map $X = cX(b) \to f_* f^* cX(b) = X^{M(b)}$ whenever the set $M(b)$ is nonempty and this map is always a mono. When the set $M(b)$ is empty, then $X^{M(b)}$ is the terminal object and the unit map will not be a mono for X with more than one element. So the condition for being dense is exactly that of forcing the overcategories occurring in the construction of right Kan extensions to be nonempty.

I cannot think of a universal way of enforcing this in general categories, but maybe for your preorders one can, I haven't yet thought about it.

The closed embedding part seems difficult. I have to think more about it...

A counterexample showing that the (dense, closed) factorization is not induced by first adding objects then adding+identifying arrows:

Consider the inclusion of small categories $f : [a \leftarrow b] \hookrightarrow [a \leftarrow b \to c]$. Your proposed factorization goes

$$[a \leftarrow b] \hookrightarrow [a \leftarrow b \quad c] \hookrightarrow [a \leftarrow b \to c]$$

. Call the first arrow i and the second j.

Density part: Let $X := \{0, 1\}$ be the two element set and cX the constant presheaf on $[a \leftarrow b \quad c]$ with value X.

The condition for density is that the unit $cX \to i_* i^* cX$ (which is a natural transformation of presheaves on $[a \leftarrow b \quad c]$) is a mono. The pullback $i^* cX$ is the constant presheaf with value X on the category $[a \leftarrow b]$ (it is just the restriction of the previous constant presheaf to the subcategory $[a \leftarrow b]$). Let's compute its right Kan extension along i.

This is a pointwise Kan extension. To compute the value $i_* i^* cX(z)$ at an object z of the target category (here: $[a \leftarrow b \quad c]$) one does the following: Form the comma category $z \downarrow i$, i.e. the category with objects $z \to i(y)$ for some y in the domain category and arrows commutative triangles with one side of the form "i of some morphism", then forget the morphisms from z and forget the is - this gives us a diagram of objects in the domain category (here: $[a \leftarrow b]$) – an object y of the domain category can occur multiple times in this diagram (one time for each morphism $z \to i(y)$), so this is not a subdiagram of the domain category. To this diagram apply now the given functor that we want to extend (here: constant functor with value X); we get a diagram of sets. Now take the limit of this diagram.

If you have a fully faithful functor along which you want to Kan extend, then the values of the objects in the image are the same as before (because the category over which we take the limit has an initial object)

Hence for the objects in the image of i we get the same value as before, i.e. we have in our example $(i_*i^*cX)(a) = (i^*cX)(a) = X$ and $(i_*i^*cX)(b) = (i^*cX)(b) = X$. The comma category for the object c has no objects, since there is no morphism from c to something in the image of i, so here we take the limit over the empty diagram and get $(i_*i^*cX)(c) = \{*\}$.

Thus the morphism $cX \to i_*i^*cX$, evaluated at $c \in Ob([a \leftarrow b \quad c])$ is the map $\{0, 1\} \to 1$ and therefore not a mono.

Closed inclusion part:

Adding arrows does usually not even give an inclusion, let alone a closed inclusion. Inclusion means that the right adjoint is fully faithful or, equivalently, that the counit $j^*j_*X \to X$ is an isomorphism for all X. In our presheaf case the right adjoint j_* is given by right Kan extension. Consider the case $j : [a \quad b] \hookrightarrow [a' \to b']$ (inclusion of a discrete subcategory, this case would be contained in many other cases). To compute the value of the Kan extension f_*X at a' we have to consider the comma category $ja \leftarrow a' \to jb$, forget the a' and the js and then take the limit of the values in sets. The outcome is: $j_*X(a') = X(a) \times X(b)$. It is easy to see that $j_*X(b') = X(b)$. Now consider the case $X(a) \neq \emptyset$ and $X(b) = \emptyset$ then $X(a) \times X(b) = \emptyset \neq X(a)$. Then the counit, whatever it is cannot be an isomorphism. (The counit that would have to be an iso is a pair of morphisms of sets, given by $id_{X(b)}$ and $pr_2 : X(a) \times X(b) \to X(a)$, if I see it correctly).

Alternatively one can see directly that j_* is not fully faithful: A morphism of presheaves on the discrete category $[a \quad b]$, i.e. a pair of morphisms $(g : X(a) \to Y(a), h : X(b) \to Y(b))$ is sent by j_* to $(g \times h : X(a) \times X(b) \to Y(a) \times Y(b), h : X(b) \to Y(b))$ (I think). Again take the example $Y(b) = \emptyset$, then the map $g : X(a) \to Y(a)$ is simply forgotten.

This example also occurs within the previous example, so it gives a counterexample to whole factorization. I think in general it is a good idea to look at the index category $[a' \to b']$ for the (dense,closed) part. Looking at the definition of closed embedding in terms of Lawvere-Tierney topologies, this is an easy example. There are only three subterminal objects in this topos, and, according to this post http://math.stackexchange.com/a/177894 only three Lawvere-Tierney topologies (should be easy to check), whose corresponding subtoposes of sheaves are equivalent to the one object topos, *Set*, and the whole topos, respectively. So there is not much choice for a factorization there. I think that the nontrivial subtopos is indeed closed.

I didn't look at the general construction of the (dense,closed) factorization, but it is a construction where one takes sheaves and to me it feels unlikely that in general one gets a topos of presheaf type this way. I don't think it will be too hard to settle this question, but I didn't try. It is still a possibility that there is a (dense,closed) factorization *within the world of presheaf toposes*, it might or might not be unique, it might or might not be induced by functors of the small index categories, but I somehow doubt it. But on the other hand for Heyting algebras it should work, I think, since the (dense,closed) factorization for localic (or spatial) toposes doesn't lead out of the world of localic toposes...

So, part one of your conjecture is correct (the way I read it). You write that the (surjective,inclusion) factorization of a geometric morphism is induced by the (epi,mono) factorization of the index categories. "Epi-mono" factorization is an uncommon name for functors - what you can use here, and what you probably meant, is the (essentially surjective, fully faithful) factorization for functors. This factorization is unique up to equivalence and can be realized concretely in different ways, e.g. by factoring the given functor through the full subcategory on the objects in the image (the second part of the factorization is then the inclusion of this full subcategory).

Proof that this gives the (surjection,inclusion) factorization of the presheaf categories (maybe you already told me this in Istanbul, I don't remember):

It is enough to check that the (essentially surjective, fully faithful) factorizationof the functor between the index categories induces a factorization into a surjectivegeometric morphism and an inclusion, since this factorization is essentially unique (Moerdijk/MacLane VII, Prop 4/Thm 6).

By definition a geometric morphism (f_*, f^*) is surjective if f^* is faithful. In our case f^* is precomposition. This is faithful for a functor f that is surjective on objects, since natural transformations are determined objectwise and all objects are in the image. So the first part of the factorization indeed yields a surjective geometric morphism.

For the second part we have to show that the essential geometric morphism (g_*, g^*) induced by a fully faithful functor g between the small index categories is an inclusion. By definition this means that g_* is fully faithful. But in our case g_* is right Kan extension along g, and (left or right) Kan extensions along fully faithful functors are fully faithful again.
(end of proof)

Note that this goes through for any index categories, not just preorders. Also note that it was not a priori clear that the middle topos ocurring in the (surjection,inclusion) factorization is a presheaf topos again, or that the two parts of the factorization are essential geometric morphisms again, but we have now seen that this is true. However, I have some doubts about the corresponding statement for the (dense, closed) factorization.

Your second proposed factorization $B \to C \to D$ you described as "adding just points, then adding just arrows". This factorization probably doesn't exist for general functors. You can factorize a functor by first adjoining the objects from the target category which are not in the image. In the second step you can take care of the Hom-sets, but the maps betwen these will not in general be injective or surjective, i.e. it will not be just adding arrows, but also identifying a few old arrows. For your preorders this doesn't occur, of course, so your terminology makes sense (but it confused me a bit because I was thinking of general small categories).

For a set X let cX denote the constant presheaf with value X. That a geometric morphism (f_*, f^*) is dense, means that for a set X the unit map $cX \to f_* f^* cX$ is a mono. In our case, where the geometric morphism is induced by a functor $f : A \to B$ between small categories, f^* is precomposition with f, so $f^* cX$ is the constant presheaf on the domain category of the functor f. Again f_* is the right Kan extension along f. Looking at the pointwise construction of right Kan extension, by taking limits over some comma category, one can see that for an object b of B we have $f_* f^* cX(b) = X^{M(b)}$ where $M(b)$ is the set of connected components of the comma category $b \downarrow f$. The unit map $cX \to f_* f^* cX$ is objectwise given by the diagonal map $X = cX(b) \to f_* f^* cX(b) = X^{M(b)}$ whenever the set $M(b)$ is nonempty and this map is always a mono. When the set $M(b)$ is empty, then $X^{M(b)}$ is the terminal object and the unit map will not be a mono for X with more than one element. So the condition for being dense is exactly that of forcing the overcategories occurring in the construction of right Kan extensions to be nonempty.

I cannot think of a universal way of enforcing this in general categories, but maybe for your preorders one can, I haven't yet thought about it.

The closed embedding part seems difficult. I have to think more about it...
——— 2nd mail ————————————————-

A counterexample showing that the (dense, closed) factorization is not induced by first adding objects then adding+identifying arrows:

Consider the inclusion of small categories $f : [a \leftarrow b] \hookrightarrow [a \leftarrow b \to c]$. Your proposed factorization goes

$$[a \leftarrow b] \hookrightarrow [a \leftarrow b \quad c] \hookrightarrow [a \leftarrow b \to c]$$

. Call the first arrow i and the second j.

Density part: Let $X := \{0, 1\}$ be the two element set and cX the constant presheaf on $[a \leftarrow b \quad c]$ with value X.

The condition for density is that the unit $cX \to i_* i^* cX$ (which is a natural transformation of presheaves on $[a \leftarrow b \quad c]$) is a mono. The pullback $i^* cX$ is the constant presheaf with value X on the category $[a \leftarrow b]$ (it is just the restriction of the previous constant presheaf to the subcategory $[a \leftarrow b]$). Let's compute its right Kan extension along i.

This is a pointwise Kan extension. To compute the value $i_* i^* cX(z)$ at an object z of the target category (here: $[a \leftarrow b \quad c]$) one does the following: Form the comma category $z \downarrow i$, i.e. the category with objects $z \to i(y)$ for some y in the domain category and arrows commutative triangles with one side of the form "i of some morphism", then forget the morphisms from z and forget the is - this gives us a diagram of objects in the domain category (here: $[a \leftarrow b]$) – an object y of the domain category can occur multiple times in this diagram (one time for each morphism $z \to i(y)$), so this is not a subdiagram of the domain category. To this diagram apply now the given functor that we want to extend (here: constant functor with value X); we get a diagram of sets. Now take the limit of this diagram.

If you have a fully faithful functor along which you want to Kan extend, then the values of the objects in the image are the same as before (because the category over which we take the limit has an initial object)

Hence for the objects in the image of i we get the same value as before, i.e. we have in our example $(i_*i^*cX)(a) = (i^*cX)(a) = X$ and $(i_*i^*cX)(b) = (i^*cX)(b) = X$. The comma category for the object c has no objects, since there is no morphism from c to something in the image of i, so here we take the limit over the empty diagram and get $(i_*i^*cX)(c) = \{*\}$.

Thus the morphism $cX \to i_*i^*cX$, evaluated at $c \in Ob([a \leftarrow b \quad c])$ is the map $\{0,1\} \to 1$ and therefore not a mono.

Closed inclusion part:

Adding arrows does usually not even give an inclusion, let alone a closed inclusion. Inclusion means that the right adjoint is fully faithful or, equivalently, that the counit $j^*j_*X \to X$ is an isomorphism for all X. In our presheaf case the right adjoint j_* is given by right Kan extension. Consider the case $j : [a \quad b] \hookrightarrow [a' \to b']$ (inclusion of a discrete subcategory, this case would be contained in many other cases). To compute the value of the Kan extension f_*X at a' we have to consider the comma category $ja \leftarrow a' \to jb$, forget the a' and the js and then take the limit of the values in sets. The outcome is: $j_*X(a') = X(a) \times X(b)$. It is easy to see that $j_*X(b') = X(b)$. Now consider the case $X(a) \neq \emptyset$ and $X(b) = \emptyset$ then $X(a) \times X(b) = \emptyset \neq X(a)$. Then the counit, whatever it is, cannot be an isomorphism. (The counit that would have to be an iso is a pair of morphisms of sets, given by $id_{X(b)}$ and $pr_2 : X(a) \times X(b) \to X(a)$, if I see it correctly).

Alternatively one can see directly that j_* is not fully faithful: A morphism of presheaves on the discrete category $[a \quad b]$, i.e. a pair of morphisms $(g : X(a) \to Y(a), h : X(b) \to Y(b))$ is sent by j_* to $(g \times h : X(a) \times X(b) \to Y(a) \times Y(b), h : X(b) \to Y(b))$ (I think). Again take the example $Y(b) = \emptyset$, then the map $g : X(a) \to Y(a)$ is simply forgotten.

This example also occurs within the previous example, so it gives a counterexample to whole factorization. I think in general it is a good idea to look at the index category $[a' \to b']$ for the (dense,closed) part. Looking at the definition of closed embedding in terms of Lawvere-Tierney topologies, this is an easy example. There are only three subterminal objects in this topos, and, according to this post http://math.stackexchange.com/a/177894 only three Lawvere-Tierney topologies (should be easy to check), whose corresponding subtoposes of sheaves are equivalent to the one object topos, Set, and the whole topos, respectively. So there is not much choice for a factorization there. I think that the nontrivial subtopos is indeed closed.

I didn't look at the general construction of the (dense,closed) factorization, but it is a construction where one takes sheaves and to me it feels unlikely that in general one gets a topos of presheaf type this way. I don't think it will be too hard to settle this question, but I didn't try. It is still a possibility that there is a (dense,closed) factorization *within the world of presheaf toposes*, it might or might not be unique, it might or might not be induced by functors of the small index categories, but I somehow doubt it. But on the other hand for Heyting algebras it should work, I think, since the (dense,closed) factorization for localic (or spatial) toposes doesn't lead out of the world of localic toposes...

——— 3rd mail ———————————-

Cool, I like your diagrams! Everything is visible at a glance. I just found out how the (dense, closed) factorization works and it is well visualizable! It boils down to the empty/nonemptyness criterion of comma categories from the first mail, and that is something you can really see in your small examples. (by the way, the surjection/dense/closed terminology is indeed standard, and well chosen).

About subtoposes (closed or not, for now): They are all given as sheaves for a Lawvere-Tierney topology. A Lawvere-Tierney topology is an endomorphism j of the presheaf Ω, it is a closure operator (satisfying some conditions) which is on each $\Omega(X)$ given by specifying which are the closed subobjects (then, to an arbitrary subobject this closure operator associates the smallest closed subobject bigger than the given one). Of course this endomorphism is determined by saying what it does on each object, or equivalently, by saying what happens when postcomposing it with maps from representable functors to Ω.

But now, since your index categories are preorders, the representable functors are subterminal objects[1]. The upshot of what I want to write is that in a presheaf category many constructions are determined by what they do the representable functors, and that here the representable functors are subterminal, which gives a good handle on Lawvere-Tierney topologies, since they are about closure of subobjects, in particular those of the terminal object.

Just to be sure: presheaves are contravariant functors for me now. For a subterminal object $A \to 1$ the restriction map is given by $\Omega(1) \to \Omega(A), V \mapsto V \cap A$ and surjective. From the condition that j is a natural transformation we get that the closed subobjects of A are exactly the intersections of closed subobjects of 1 with A:

$$\begin{array}{ccc} \Omega(1) & \xrightarrow{U \mapsto U \cap A} & \Omega(A) \\ {\scriptstyle U \mapsto \bar{U}}\downarrow{\scriptstyle j} & & \downarrow{\scriptstyle j} \\ \Omega(1) & \xrightarrow{\bar{U} \mapsto \bar{U} \cap A} & \Omega(A) \end{array}$$

[1] But there can still be subterminal objects which are not representable

Moreover for a *closed* subobject $\bar{A}$ of 1 we have that the closed subobjects of $\bar{A}$ are exactly the subobjects which are closed in $\Omega(1)$ and are contained in $\bar{A}$: We know that the closed subobjects of $\bar{A}$ are those of the form $\bar{A} \cap \bar{U}$ for some closed subobject $\bar{U}$ of 1, but seeing $\bar{A} \cap \bar{U}$ as a subobject of 1 we get $\bar{A} \cap \bar{U} = j(\bar{A}) \cap j(\bar{U})$ (since both are j-closed), then $j(\bar{A}) \cap j(\bar{U}) = j(\bar{A} \cap U)$ (properties of j), hence altogether $\bar{A} \cap \bar{U} = j(\bar{A} \cap U)$.

Finally, considering an arbitrary subterminal object and its closure $A \hookrightarrow \bar{A} \hookrightarrow 1$, the restriction map $\Omega_{cl}(\bar{A}) \to \Omega_{cl}(A)$, which we know to be surjective, is also injective (I denote by Ω_{cl} the collection of closed subobjects). To see this, suppose we have two subobjects $\bar{A} \cap \bar{V}$, $\bar{A} \cap \bar{W}$ of $\bar{A}$ which are mapped to the same subobject of A, i.e. such that $A \cap \bar{V} = A \cap \bar{W}$. Now apply the map $\Omega(A) \to \Omega(1)$ which considers them as subobjects of 1, the apply j. We get $\bar{A} \cap \bar{V} = j(A) \cap j(\bar{V}) = j(A \cap \bar{V}) = j(A \cap \bar{W}) = j(A) \cap j(\bar{W}) = \bar{A} \cap \bar{W}$.

Altogether this shows that Lawvere-Tierney topologies for presheaves on a preorder are in bijection to the closure operators $j : \Omega(1) \to \Omega(1)$ satisfying $j(A \cap B) = j(A) \cap j(B)$. Now $\Omega(1)$ itself can be easily determined[2] and then probably also such closure operators are not hard to work out completely for very small preorders. Put differently, this means classifying all subtoposes of such a topos.

Seeing which subtoposes among them are closed (i.e. which Lawvere-Tierney topologies are of the required kind) is even easier: Just take all the subobjects and apply the recipe for the construction of a *closed* Lawvere-Tierney topology, i.e. $j : V \mapsto U \cup V$.

(Another approach would be to follow the article by Olivia Caramello I mentioned last time - taking that kind of perspective one would have to work out what kind of theory a presheaf topos of your type classifies and what is the lattice of theories coming with it.)

What about the sheafification functor for a closed Lawvere-Tierney topology?

If the topology is given by the subterminal object U, i.e. some functor that maps certain objects of the index category to $1 := \{*\}$ and the others to $\emptyset$, then by the elephant, Prop. A.4.5.3(ii) an object A is sheaf iff the projection $A \times U \cong U$ is an isomorphism (it is not necessary to talk about the projection here; since U is subterminal there is at most morphism to U anyway, so it must be the projection). The presheaf $A \times U$ has value $\emptyset$ at those objects where U has value $\emptyset$ and has the value of A at the other objects. I.e. the passage from A to $A \times U$ "deletes" the values where U has value $\emptyset$ and keeps the others. Thus the requirement for A to be a sheaf becomes: A must have value 1 wherever U has value 1 (and it can have arbitrary values otherwise).

[2]The terminal object is of course the constant functor with value $1 := \{*\}$. Its subobjects are exactly the functors which map the objects of your index preorder either to the empty set or the one point set (this can of course not happen arbitrarily but has to satisfy the constraint that an object mapped to the empty set cannot receive a map from an object mapped to the one point set, but that is the only constraint).

Accordingly, the sheafification is the functor that enforces this condition, i.e. it keeps the values of A at those objects where U produces the empty set and resets them to 1 at those objects where U produces the value 1 — this is what Prop. A.4.5.3(iii) says.

This description of sheaves makes it apparent that a closed suptopos of a presheaf topos Set^C is actually a presheaf topos again — it can be described as the topos of presheaves on the full subcategory of C whose objects are mapped to $\emptyset$ by the subterminal object U (in one direction of the equivalence one can forget the other objects, in the other direction one adds the value 1 at all the additional objects)[3].

So your (dense, closed) factorization actually stays in the world of presheaf toposes, it seems. The next question is when the geometric morphisms occurring in the factorization are induced by functors between the index categories. This is true for the second one, because all closed inclusions between presheaf categories are of this form:

Proof: Let U be the subterminal object of Set^C inducing the closed Lawvere-Tierney topology j_U and let D be the full subcategory of C whose objects are mapped to $\emptyset$ by U.

We need to establish a commutative (up to natural iso) triangle of geometric morphisms

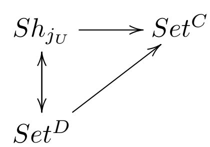

where the vertical morphism is the equivalence, the diagonal morphism is the (precomposition, right Kan extension) adjunction and the horizontal morphism is the (sheafification, inclusion) adjunction.

It is enough to show a natural isomorphism just between the left adjoints (or just the right adjoints) by uniqueness of adjoints. But this is clear: $Set^C \to Sh_{j_U}$ resets the values at the objects of $Ob(C)\setminus Ob(D)$ to 1, then $Sh_{j_U} \to Set^D$ forgets those values — and this forgetting is exactly restriction to the objects of D, i.e. precomposition with $D \hookrightarrow C$.

(I had first checked it on right adjoints. That is also instructive, you have to show that the right Kan extension produces just the value 1 at the objects of $Ob(C) \setminus Ob(D)$. By the formula for pointwise Kan extensions this happens exactly when the undercategories are empty and this is true exactly because the objects of our full subcategory are given by a subterminal object...) (end of proof)

---

[3]One can maybe also use the techniques of Olivia Caramello's article to show this, but I didn't see a statement in the relevant places, e.g. section 7.2 or section 14. Another article of hers that might be useful for a logic based proof is http://arxiv.org/abs/1404.4610, in particular section 6.4. But there she doesn't mention closed subtoposes.

Now this tells us how to find the (dense, closed) factorization. The mental image for this is that from topology, i.e. the dense part should map the domain to the closure of its image, and the closed part is the inclusion of this closure. I suppose this is a valid intuition for toposes too, and that the topos occuring in the middle should be the smallest closed subtopos through which the given geometric morphism factors ("smallest" meaning: the full subcategory with the fewest objects; sheaves always form a full subcategory, so that makes sense).

So in our situation we start with a geometric inclusion $(h_*, h^*) : Set^A \hookrightarrow Set^C$ induced by a fully faithful functor $h : A \to C$ (since that is the situation after we have done the (surjection, inclusion) factorization). Our right adjoint is given by right Kan extension. We look for a full subcategory through which h_* factors and which is given by demanding that its objects take value 1 on some specified objects. Hence we need to look which objects c of C get value 1 after performing the Kan extension – we know that these are the ones whose comma category $c \downarrow h$ is empty. So a careless first guess could be that the index category in the middle (in $A \to D \to C$) is the full subcategory of C whose objects are the c with $c \downarrow h$ empty.

One might think that we need a small correction to this guess: We cannot just take any set of objects where we want our functors to take the value 1; the set of objects must be specified by a subterminal object as before, in order to give a closed subtopos. So we need to take the smallest *such* set of objects. This extra condition is sort of a hereditariness condition: If an object is in our subcategory and another one receives a morphism from it, then it also needs to be in the subcategory. However, if $c \downarrow h$ is empty and there is a morphism $c \to c'$, then $c' \downarrow h$ is also empty, otherwise there would be some $c' \to h(a)$, hence also $c \to c' \to h(a)$. So our guess was actually not careless and the extra condition is automatic.

By construction (resp. by the observation in the last paragraph) the inclusion $Set^D \to Set^C$ is closed. As for the other part of the factorization, i.e. $Set^A \to Set^D$: This is also the one induced by right Kan extension along $A \to D$ and it is dense by the criterion we saw at the end of my first mail: This criterion was exactly that the comma categories $d \downarrow h$ need to be nonempty and we achieved that by construction (since we put all the other objects in the second step). Yay! That is the proof that we get a (dense, closed) factorization this way!

Ok, I just took a look at the elephant. The proof of Lemma A.4.5.19(i) (together with the paragraph after and A.4.5.20) says how to construct the closure in the general case: You start with a geometric inclusion. The domain of this then corresponds to the topos of sheaves for some Lawvere-Tierney topology j. You take the smallest subobject of 1, i.e. $0 \hookrightarrow 1$, then form its closure $j(0 \hookrightarrow 1)$. This is another subobject of 1 which you then use to form a *closed* Lawvere-Tierney topology and thus a closed subtopos. We could compare this with the above construction. We would have to figure out how to obtain the j from the geometric morphism, it probably as simple as chasing the initial object back and forth through the adjunction...

Kan extensions: a general case and a particular case.

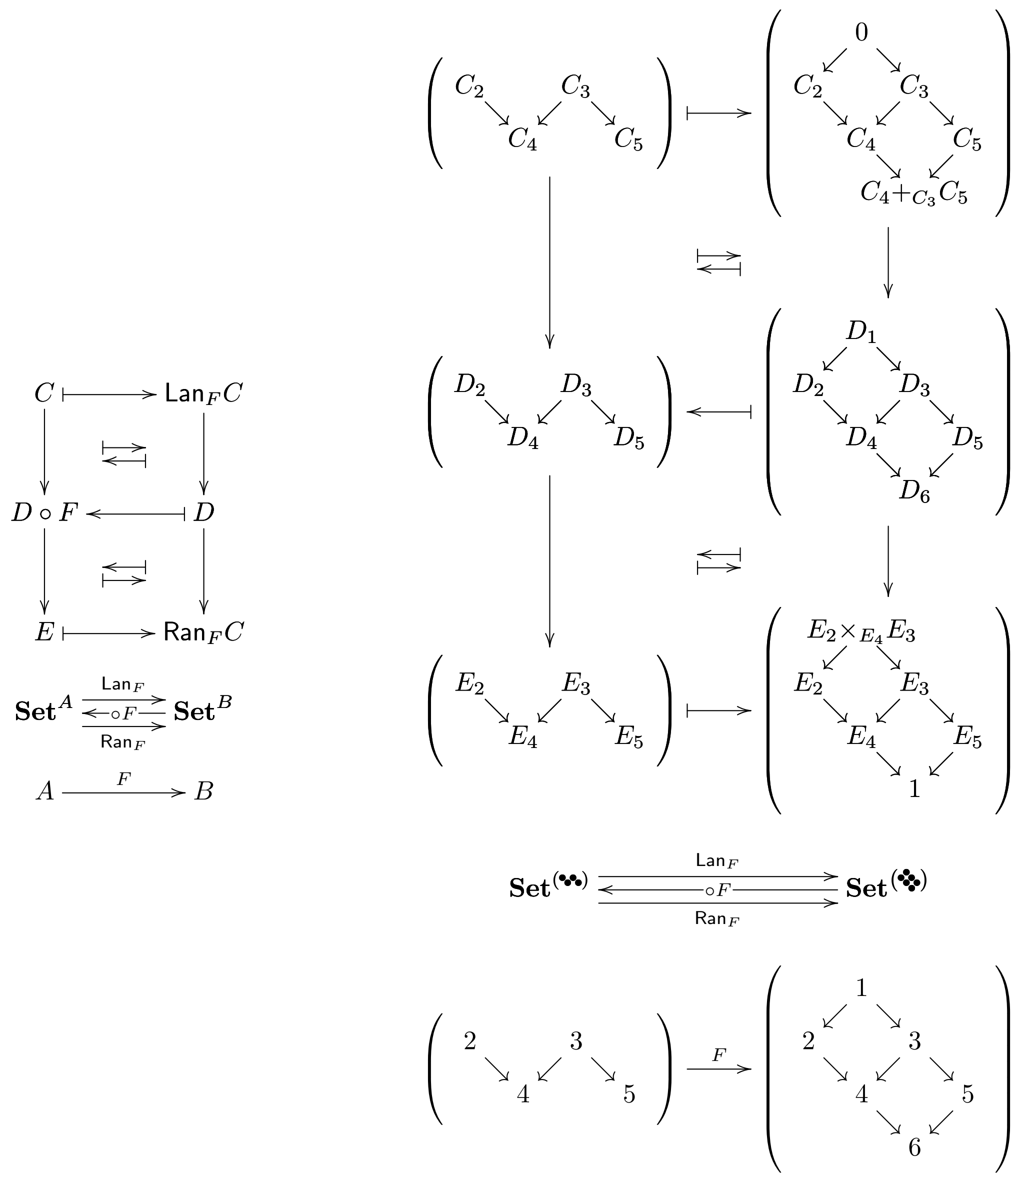

Equivalent conditions

$$\begin{array}{ccc} A' & \longmapsto & \alpha_! A' \\ \downarrow & & \downarrow \\ \alpha^* B & \underset{\text{faithful}}{\longleftarrow\!\!\shortmid} & B \\ \downarrow & & \downarrow \\ A & \longmapsto & \alpha_* A \end{array}$$

$$\mathbf{Set}^{\mathbf{A}} \xrightarrow[\text{surjective}]{\alpha} \mathbf{Set}^{\mathbf{B}}$$

$$\mathbf{A} \xrightarrow[\substack{\text{essentially} \\ \text{surjective}}]{\alpha} \mathbf{B}$$

$$\begin{array}{ccc} B' & \longmapsto & \iota_! B' \\ \downarrow & & \downarrow \\ \iota^* D & \longleftarrow\!\!\shortmid & D \\ \downarrow & & \downarrow \\ B & \underset{\text{fully faithful}}{\longmapsto} & \iota_* B \end{array}$$

$$\mathbf{Set}^{\mathbf{B}} \xrightarrow[\text{inclusion}]{\iota} \mathbf{Set}^{\mathbf{D}}$$

$$\mathbf{B} \xrightarrow[\text{fully faithful}]{\iota} \mathbf{D}$$

$$\begin{array}{ccc} B' & \longmapsto & \beta_! B' \\ \downarrow & & \downarrow \\ \beta^* \mathrm{k} C & \longleftarrow\!\!\shortmid & \mathrm{k} C \\ {\scriptstyle \text{id}}\downarrow & & \downarrow{\scriptstyle \text{mono}} \\ \beta^* \mathrm{k} C & \longmapsto & \beta_* \beta^* \mathrm{k} C \end{array}$$

$$\mathbf{Set}^{\mathbf{B}} \xrightarrow[\text{dense}]{\beta} \mathbf{Set}^{\mathbf{C}}$$

$$\mathbf{B} \xrightarrow[?]{\beta} \mathbf{C}$$

$$\begin{array}{ccc} C' & \longmapsto & \gamma_! C' \\ \downarrow & & \downarrow \\ \gamma^* \gamma_* C & \longleftarrow\!\!\shortmid & \gamma_* C \\ {\scriptstyle \text{iso}}\downarrow & & \downarrow{\scriptstyle \text{id}} \\ C & \longmapsto & \gamma_* C \end{array}$$

$$\mathbf{Set}^{\mathbf{C}} \xrightarrow[\text{closed}]{\gamma} \mathbf{Set}^{\mathbf{D}}$$

$$\mathbf{C} \xrightarrow[?]{\gamma} \mathbf{D}$$

Peter's counterexample, 1: the Kan extensions

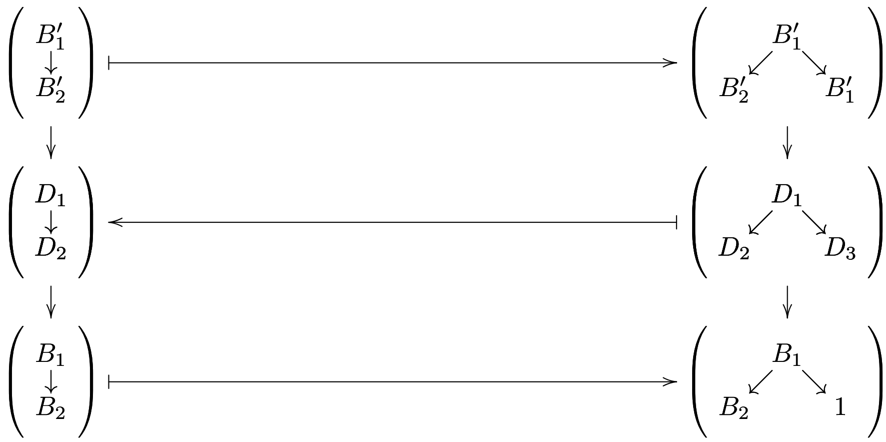

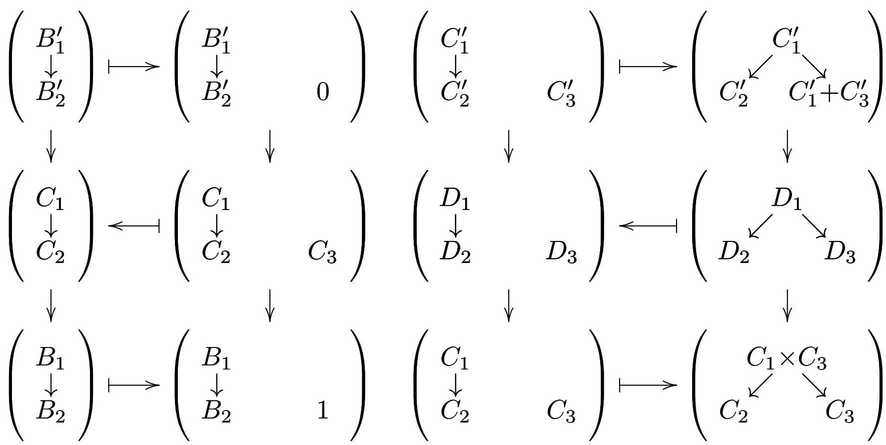

Peter's counterexample, 2: the equivalent conditions fail

$$
\begin{array}{ccccccc}
\begin{pmatrix} C \\ \downarrow \\ C \end{pmatrix} & \longleftarrow\!\!\mid & \begin{pmatrix} C & \\ \downarrow & \\ C & C \end{pmatrix} & & \begin{pmatrix} C_1 \times C_3 & \\ \downarrow & \\ C_2 & C_3 \end{pmatrix} & \longleftarrow\!\!\mid & \begin{pmatrix} & C_1 \times C_3 & \\ \swarrow & & \searrow \\ C_2 & & C_3 \end{pmatrix} \\
{\scriptstyle \mathrm{id}}\downarrow & & \downarrow{\scriptstyle \text{not monic}} & & {\scriptstyle \text{not iso}}\downarrow & & \downarrow{\scriptstyle \mathrm{id}} \\
\begin{pmatrix} C \\ \downarrow \\ C \end{pmatrix} & \longmapsto & \begin{pmatrix} C & \\ \downarrow & \\ C & 1 \end{pmatrix} & & \begin{pmatrix} C_1 & \\ \downarrow & \\ C_2 & C_3 \end{pmatrix} & \longmapsto & \begin{pmatrix} & C_1 \times C_3 & \\ \swarrow & & \searrow \\ C_2 & & C_3 \end{pmatrix}
\end{array}
$$