

WHAT DOES THE CLASSIFYING SPACE OF A CATEGORY CLASSIFY ?

MICHAEL WEISS

Let \mathcal{C} be a small category. Contravariant functors from \mathcal{C} to the category of sets, and natural transformations between them, will be called \mathcal{C} -sets and \mathcal{C} -maps, respectively. The category of \mathcal{C} -sets shares many good properties with the category of sets. (In short, it is a *topos*. See [4] or [6]. Here we will not make any explicit use of this fact.) The \mathcal{C} -sets which are of the form $b \mapsto \text{mor}_{\mathcal{C}}(b, c)$ for fixed $c \in \mathcal{C}$, and any isomorphic ones, are called *representable*. By the Yoneda lemma, the representable \mathcal{C} -sets form a full subcategory of the category of all \mathcal{C} -sets which is equivalent to \mathcal{C} .

We will be concerned with *sheaves* of \mathcal{C} -sets on a topological space X . For such a sheaf, and $x \in X$, the *stalk* \mathcal{F}_x is again a \mathcal{C} -set. It is the direct limit of the \mathcal{C} -sets $\mathcal{F}(U)$ where U runs through the open neighborhoods of x .

Theorem 0.1. *The classifying space BC classifies sheaves of \mathcal{C} -sets with representable stalks.*

Notation, terminology, clarifications. Let \mathcal{F} be any sheaf of \mathcal{C} -sets on X . We may regard \mathcal{F} as a contravariant functor $(c, U) \mapsto \mathcal{F}^{(c)}(U)$ in two variables (where $c \in \text{ob}(\mathcal{C})$ and U is open in X). Specializing one of the variables, we obtain $\mathcal{F}^{(c)}$, a sheaf of sets on X , and $\mathcal{F}(U)$, a \mathcal{C} -set.

Let \mathcal{L} be any sheaf of sets on X . The *espace étalé* of \mathcal{L} , denoted $\text{Spé}(\mathcal{L})$, is the (disjoint) union of the stalks \mathcal{L}_x , suitably topologized. See [2, II.1, ex.1.13] for details. The sheaf \mathcal{L} can be identified with the sheaf of continuous (partial) sections of the projection

$$\text{Spé}(\mathcal{L}) \rightarrow X.$$

The projection $\text{Spé}(\mathcal{L}) \rightarrow X$ is an étale map alias *local homeomorphism*. [But it happens frequently that X is Hausdorff while $\text{Spé}(\mathcal{L})$ is not.] The construction $\text{Spé}(\mathcal{L})$ leads to a good notion of *pullback* of sheaves: for a map $v: Y \rightarrow X$, the pullback $v^*\mathcal{L}$ is defined in such a way that $\text{Spé}(v^*\mathcal{L}) = v^*\text{Spé}(\mathcal{L})$. More generally, for a sheaf \mathcal{F} of \mathcal{C} -sets on X and $v: Y \rightarrow X$, the pullback $v^*\mathcal{F}$ is defined in such a way that $\text{Spé}((v^*\mathcal{F})^{(c)}) = v^*\text{Spé}(\mathcal{F}^{(c)})$.

Let \mathcal{F} and \mathcal{G} be sheaves of \mathcal{C} -sets on X , both with representable stalks. Let $e_0, e_1: X \rightarrow X \times [0, 1]$ be given by $e_0(x) = (x, 0)$ and $e_1(x) = (x, 1)$. The sheaves \mathcal{F} and \mathcal{G} are *concordant* if there exists a sheaf of \mathcal{C} -sets \mathcal{H} on $X \times [0, 1]$, again with representable stalks, such that $e_0^*\mathcal{H} \cong \mathcal{F}$ and $e_1^*\mathcal{H} \cong \mathcal{G}$.

The precise meaning of theorem 0.1 is as follows. Suppose that X has the homotopy type of a *CW*-space. There is a natural bijection from the homotopy set $[X, BC]$ to the set of concordance classes of sheaves of \mathcal{C} -sets on X with representable stalks.

Partially supported by the Royal Society.

Remark. Suppose that \mathcal{C} is a group. To be more precise, suppose that \mathcal{C} has just one object c and $\text{mor}(c, c)$ is a group. Let \mathcal{F} be a sheaf of \mathcal{C} -sets on a space X . If the stalks of \mathcal{F} are all representable, then the projection $\text{Spé}(\mathcal{F}) \rightarrow X$ is a principal $\text{mor}(c, c)$ -bundle. Indeed any choice of an open U and $s \in \mathcal{F}^{(c)}(U)$ determines a bundle chart

$$\text{Spé}(\mathcal{F}|U) \cong \text{mor}(c, c) \times U.$$

In this situation, concordant sheaves of \mathcal{C} -sets on X (with representable stalks) are isomorphic, because “concordant” implies “isomorphic” for principal $\text{mor}(c, c)$ -bundles.

Remark. The question in the title undoubtedly has many correct answers and a few have already been given elsewhere. Moerdijk [6, Introd.] has a result like theorem 0.1 in which the representability condition on stalks is replaced by a weaker condition. (He also replaces \mathcal{C} by \mathcal{C}^{op} at that point, but for simplicity I undo that.) A \mathcal{C} -set S satisfies Moerdijk’s condition if and only if every \mathcal{C} -map from a *finitely presented* \mathcal{C} -set to S factors through a representable \mathcal{C} -set. (A finitely presented \mathcal{C} -set T is described by finitely many generators $x_i \in T(c_i)$ and finitely many relations of the form $f_r^*(x_i) = g_r^*(x_j)$, where $f_r: b_r \rightarrow c_i$ and $g_r: b_r \rightarrow c_j$ are morphisms in \mathcal{C} .) Any retract of a representable \mathcal{C} -set satisfies Moerdijk’s condition.

Moerdijk [6] takes the discussion much further by considering topological categories, which I have not attempted to do.

Another precursor of theorem 0.1 is due to tom Dieck (1972, unpublished). He used a notion of \mathcal{C} -bundle defined in terms of a bundle atlas. His result was rediscovered in [5, thm 4.1.2].

1. THE CANONICAL SHEAF ON $B\mathcal{C}$

We are going to construct a sheaf \mathcal{E} of \mathcal{C} -sets on $B\mathcal{C}$ which will eventually turn out to be “universal”. Recall to begin with that $B\mathcal{C}$ is the geometric realization of the simplicial set whose n -simplices are the contravariant functors $[n] \rightarrow \mathcal{C}$, where $[n]$ is the linearly ordered set $\{0, 1, 2, \dots, n\}$. (There are historical reasons for insisting on *contravariant* functors $[n] \rightarrow \mathcal{C}$; the formulae for boundary operators look more familiar in the case where \mathcal{C} is a group or monoid.) Now suppose that U is open in $B\mathcal{C}$ and c is an object of \mathcal{C} .

Definition 1.1. An element of $\mathcal{E}^{(c)}(U)$ is a “function” which to every $\alpha: [n]^{\text{op}} \rightarrow \mathcal{C}$ and $x \in B[n]^{\text{op}} \cong \Delta^n$ with $\alpha_*x \in U$ assigns a morphism $s(\alpha, x): c \rightarrow \alpha(0)$ in \mathcal{C} . The function is required to be

- locally constant in the second variable, so that for $y \in \Delta^n$ sufficiently close to x , with $\alpha_*y \in U$, we have $s(\alpha, y) = s(\alpha, x)$;
- natural in the first variable. That is, for an order-preserving $g: [m] \rightarrow [n]$ and $y \in \Delta^m$, we have

$$s(\alpha, g_*y) = \alpha(0, g(0)) \circ s(\alpha g, y)$$

where $\alpha(0, g(0)): \alpha(g(0)) \rightarrow \alpha(0)$ is the morphism in \mathcal{C} induced by the unique morphism $0 \rightarrow g(0)$ in $[n]$.

The contravariant dependence of $\mathcal{E}^{(c)}(U)$ on U and c is obvious. The sheaf property is also obvious. Because of the naturality condition, an element s of $\mathcal{E}^{(c)}(U)$ is determined by its values $s(\alpha, x)$ for *nondegenerate* $\alpha: [n]^{\text{op}} \rightarrow \mathcal{C}$ and $x \in \Delta^n \setminus \partial\Delta^n$. Then α and x are determined by $\alpha_*x \in U$; in particular n is the dimension of the cell (in the canonical CW–decomposition of BC) to which α_*x belongs. (Regarding *cells*, the convention used here is that the cells of a CW–space are pairwise disjoint, and each cell is homeomorphic to some euclidean space. This is in agreement with [1], for example.)

Lemma 1.2. *Fix a nondegenerate $\beta: [m]^{\text{op}} \rightarrow \mathcal{C}$ and $y \in \Delta^m \setminus \partial\Delta^m$. The stalk of \mathcal{E} at β_*y is the contravariant functor $c \mapsto \text{mor}(c, \beta(0))$.*

Proof. Any point of BC can be uniquely written as α_*z where $\alpha: [n]^{\text{op}} \rightarrow \mathcal{C}$ is nondegenerate and $z \in \Delta^n \setminus \partial\Delta^n$. If α_*z is sufficiently close to β_*y , then some degeneracy of β will be a face of α . That is, there are an order–preserving surjection $f: [k] \rightarrow [m]$ and an order–preserving injection $g: [k] \rightarrow [n]$ such that $\alpha g = \beta f$. And moreover, there will be $w \in \Delta^k$ such that $f_*w = y$ and z is close to g_*w . For s in the stalk of $\mathcal{E}^{(c)}$ at $\beta_*y \in BC$, we then have

$$\begin{aligned} s(\alpha, z) &= s(\alpha, g_*w) = \alpha(0, g(0)) \circ s(\alpha g, w) , \\ s(\alpha g, w) &= s(\beta f, w) = s(\beta, f_*w) = s(\beta, y) , \end{aligned}$$

so that $s(\alpha, z) = \alpha(0, g(0)) \circ s(\beta, y)$. Hence s is determined by $s(\beta, y) \in \text{mor}(c, \beta(0))$. To establish the existence of a germ s with prescribed value $s(\beta, y)$, we proceed differently. Suppose inductively that the values of s at points near β_*y and in the $(n-1)$ –skeleton of BC have already been determined consistently, for some fixed $n > m$. For an n –simplex $\alpha: [n]^{\text{op}} \rightarrow \mathcal{C}$ we have an attaching map from $\partial\Delta^n$ to the $(n-1)$ –skeleton of BC . Hence $s(\alpha, x)$ is already determined for x in some open subset V of $\partial\Delta^n$. As a function on V , denoted informally $s|_V$, it satisfies the continuity and naturality conditions of definition 1.1 (*mutatis mutandis*). We now have to find an open $W \subset \Delta^n$ such that $V = W \cap \partial\Delta^n$ and an extension of $s|_V$ from V to W . This is easy. For example, Δ^n can be identified with a cone on $\partial\Delta^n$ and W could then be defined as the cone on V minus the cone point. Then $s|_V$ has a unique extension from V to W . \square

For an object c of \mathcal{C} , let $(c \downarrow \mathcal{C})$ be the “under” category associated with c . The objects of $(c \downarrow \mathcal{C})$ are the morphisms in \mathcal{C} with source c , and the morphisms of $(c \downarrow \mathcal{C})$ are morphisms in \mathcal{C} *under* c . The classifying space $B(c \downarrow \mathcal{C})$ is contractible since $(c \downarrow \mathcal{C})$ has an initial object. The forgetful map $B(c \downarrow \mathcal{C}) \rightarrow BC$ has a canonical factorization

$$B(c \downarrow \mathcal{C}) \xrightarrow{\lambda_c} \text{Spé}(\mathcal{E}^{(c)}) \xrightarrow{\text{proj.}} BC .$$

Indeed, any point of BC can be uniquely written as α_*z where $\alpha: [n]^{\text{op}} \rightarrow \mathcal{C}$ is nondegenerate and $z \in \Delta^n \setminus \partial\Delta^n$. Lifting α_*z to $B(c \downarrow \mathcal{C})$ amounts to specifying a morphism $c \rightarrow \alpha(n)$ in \mathcal{C} ; lifting α_*z to $\text{Spé}(\mathcal{E}^{(c)})$ amounts to specifying a morphism $c \rightarrow \alpha(0)$ in \mathcal{C} . Clearly a morphism $c \rightarrow \alpha(n)$ determines a morphism $c \rightarrow \alpha(0)$ by composition with $\alpha(0, n): \alpha(n) \rightarrow \alpha(0)$.

The map λ_c will be useful in the proof of

Proposition 1.3. *The space $\text{Spé}(\mathcal{E}^{(c)})$ is weakly contractible.*

Proof. Let \bar{BC} be the fat realization of the nerve of \mathcal{C} , obtained by ignoring the degeneracy operators. The quotient map $q: \bar{BC} \rightarrow BC$ is a quasifibration with

contractible fibers. To see this, note that the fat realization of any simplicial set Z can be described as the ordinary realization of another simplicial set \bar{Z} whose n -simplices are triples (k, f, x) where $x \in Z_k$ and $f: [n] \rightarrow [k]$ is an order-preserving surjection. The forgetful simplicial map $\bar{Z} \rightarrow Z$ is a Kan fibration with contractible fibers; hence the induced map of (lean) geometric realizations, $|\bar{Z}| \rightarrow |Z|$, is a quasifibration with contractible fibers. See [3].

Let $E = \text{Spé}(\mathcal{E}^{(c)})$, let $r: E \rightarrow BC$ be the projection, and let $\bar{E} = q^*E$. In the pullback square

$$\begin{array}{ccc} \bar{E} & \xrightarrow{\bar{r}} & \bar{B} \\ r^*q \downarrow & & \downarrow q \\ E & \xrightarrow{r} & B \end{array}$$

the map q is a quasifibration with contractible fibers and r is a local homeomorphism. It follows that r^*q is a quasifibration with contractible fibers, and consequently a weak homotopy equivalence.

It remains to prove that \bar{E} is contractible, or equivalently, that the canonical map $\bar{\lambda}_c: \bar{B}(c \downarrow C) \rightarrow \bar{E}$, the fat version of λ_c , is a weak homotopy equivalence. Suppose therefore that Y is any finitely generated Δ -set (= ‘‘simplicial set without degeneracy operators’’) and let

$$f: |Y| \rightarrow \bar{E}$$

be any map. We want to show that, up to a homotopy, f lifts to $\bar{B}(c \downarrow C)$. The argument has two parts.

- (i) *If $\bar{r}f: |Y| \rightarrow \bar{B}$ is induced by a map of the underlying Δ -sets, then f admits a unique factorization through $\bar{B}(c \downarrow C)$.*
- (ii) *Modulo iterated barycentric subdivision of $|Y|$, and a homotopy of f , the composition $\bar{r}f$ is indeed induced by a map of the underlying Δ -sets.*

For the proof of (i), we may assume that Y is generated by a single n -simplex, so $|Y| = \Delta^n$. Suppose that $\bar{r}f: \Delta^n \rightarrow BC$ is the characteristic map of an n -simplex $\alpha: [n]^{\text{op}} \rightarrow \mathcal{C}$ in the nerve of \mathcal{C} . The extra information contained in f amounts to compatible morphisms $u_i: c \rightarrow \alpha(i)$ for $i = 0, 1, \dots, n$; clearly all u_i are determined by u_n . Together, u_n and α determine an n -simplex in the nerve of $(c \downarrow \mathcal{C})$.

For the proof of (ii), we note that the first barycentric subdivision of \bar{BC} can be described as \bar{BC}' for another category \mathcal{C}' . An object of \mathcal{C}' is a simplex of the nerve of \mathcal{C} ; a morphism from an m -simplex α to an n -simplex β is an injective order-preserving $v: [m] \rightarrow [n]$ with $v^*\beta = \alpha$. The functor $\mathcal{C}' \rightarrow \mathcal{C}$ given by $\alpha \mapsto \alpha(0)$ induces a Δ -map from the nerve of \mathcal{C} to the nerve of \mathcal{C}' , and then a map

$$\varphi_1: \bar{BC} \rightarrow \bar{BC}'.$$

This map is not a homeomorphism. There is of course another (well-known) map $\varphi_0: \bar{BC}' \rightarrow \bar{BC}$ which is a homeomorphism. What is important here is that φ_0 and φ_1 are homotopic in an obvious way, by a homotopy $(\varphi_t)_{t \in [0,1]}$. (Each track of the homotopy is a straight line segment, or a single point, in a simplex of \bar{BC} .) The homotopy $(\varphi_t \varphi_0^{-1})_{t \in [0,1]}$, from the identity of \bar{BC} to $\varphi_1 \varphi_0^{-1}$, has a unique lift to a homotopy

$$(\psi_t: \bar{E} \times [0,1] \rightarrow \bar{E})_{t \in [0,1]}$$

with $\psi_0 = \text{id}$. (To verify this claim, compare the pullbacks of \bar{E} under the maps

$$\bar{BC} \times [0,1] \xrightarrow{\varphi_0 \varphi_0^{-1}} \bar{BC}, \quad \bar{BC} \times [0,1] \xrightarrow{\text{proj.}} \bar{BC}.$$

They are homeomorphic as spaces over $\bar{BC} \times [0, 1]$.) We can similarly look at iterated barycentric subdivisions of \bar{BC} . They all have two canonical maps $\varphi_0, \varphi_1: \bar{BC} \rightarrow \bar{BC}$, one being a homeomorphism and the other being simplicial, and these two maps are homotopic by a homotopy $(\varphi_t)_{t \in [0, 1]}$. Again, the homotopy $(\varphi_t \varphi_0^{-1})_{t \in [0, 1]}$ has a unique lift to a homotopy $(\psi_t: \bar{E} \times [0, 1] \rightarrow \bar{E})_{t \in [0, 1]}$ with $\psi_0 = \text{id}$. Coming back now to maps $|Y| \rightarrow \bar{E}$, any such map is homotopic to a map f such that $\bar{r}f$ is induced by a Δ -map from some iterated barycentric subdivision of Y to some iterated barycentric subdivision of the nerve of \mathcal{C} . Compose f with ψ_1 from the above homotopy. Then $\bar{r}\psi_1 f$ is induced by a Δ -map from Y to the nerve of \mathcal{C} . \square

2. RESOLUTIONS

The previous section gives us a method to “convert” a map $f: X \rightarrow BC$ into a sheaf of \mathcal{C} -sets on X with representable stalks, by $f \mapsto f^* \mathcal{E}_{\mathcal{C}}$. Going in the opposite direction is more difficult. From a sheaf \mathcal{F} of \mathcal{C} -sets on X , we shall construct a “resolution” $p_{\mathcal{F}}: X_{\mathcal{F}} \rightarrow X$ and a map $\pi_{\mathcal{F}}: X_{\mathcal{F}} \rightarrow BC$. It turns out that $p_{\mathcal{F}}$ is a homotopy equivalence if \mathcal{F} has representable stalks and X is a CW -space. Then we can choose a homotopy inverse $p_{\mathcal{F}}^{-1}$ and obtain a map $\pi_{\mathcal{F}} p_{\mathcal{F}}^{-1}: X \rightarrow BC$, well defined up to homotopy.

Let $\mathcal{O}(X)$ be the poset of open subsets of a space X , ordered by inclusion. Let \mathcal{F} be a sheaf of \mathcal{C} -sets on X . We can regard \mathcal{F} as a contravariant functor from $\mathcal{O}(X) \times \mathcal{C}$ to sets. The functor \mathcal{F} determines a *transport category* $\mathcal{T}_{\mathcal{F}}$ whose objects are the triples (U, c, s) consisting of an object U in $\mathcal{O}(X)$, an object c in \mathcal{C} , and $s \in \mathcal{F}^{(c)}(U)$. A morphism from (U, c, s) to (V, d, t) is a morphism $U \rightarrow V$ in $\mathcal{O}(X)$ together with a morphism $f: c \rightarrow d$ in \mathcal{C} such that $f^*(t)|_U = s$. Let τ be the tautological functor (taking $U \in \mathcal{O}(X)$ to the space U) from $\mathcal{O}(X)$ to spaces, and let $\varphi: \mathcal{T}_{\mathcal{F}} \rightarrow \mathcal{O}(X)$ be the forgetful functor. Put

$$X_{\mathcal{F}} := \text{hocolim } \tau\varphi.$$

This comes with a canonical projection $p_{\mathcal{F}}: X_{\mathcal{F}} \rightarrow X$, induced by the obvious natural inclusions $\tau\varphi(U, c, s) \rightarrow X$. There is also a projection $X_{\mathcal{F}} \rightarrow BT_{\mathcal{F}}$ which we can compose with the forgetful map $BT_{\mathcal{F}} \rightarrow BC$. This gives

$$\pi_{\mathcal{F}}: X_{\mathcal{F}} \rightarrow BC.$$

Proposition 2.1. *If \mathcal{F} has representable stalks, then the projection $p_{\mathcal{F}}: X_{\mathcal{F}} \rightarrow X$ is a weak homotopy equivalence.*

The proof relies on a few lemmas which in turn rely on the notion of a microfibration. Recall that a map $p: E \rightarrow B$ is a *Serre fibration* if it has the homotopy lifting property for homotopies $X \times [0, 1] \rightarrow B$, with prescribed “initial” lift $X \rightarrow E$, where X is a CW -space. [It is enough to check this in all cases where X is a disk.] A map $p: E \rightarrow B$ is a *Serre microfibration* if, for any homotopy $h: X \times [0, 1] \rightarrow B$ with prescribed initial lift $\bar{h}_0: X \rightarrow E$, there exist a neighborhood U of $X \times \{0\}$ in $X \times [0, 1]$ and a map $\bar{h}: U \rightarrow E$ such that $p\bar{h} = h|_U$ and $\bar{h}(x, 0) = \bar{h}_0(x)$ for all $x \in X$. In that case the map \bar{h} is a *microlift* of h . [Again it is enough to check the micro-lifting property in all cases where X is a disk.]

Lemma 2.2. *Let $p: E \rightarrow B$ be a Serre microfibration. If p has weakly contractible fibers, then it is a Serre fibration.*

Notes on the proof. This is essentially due to G. Segal [7, A.2]. The hypotheses here are slightly more general, though. There is a short inductive argument as follows. The induction step consists in showing that if $p: E \rightarrow B$ is a Serre microfibration with contractible fibers, then so is the projection $p^I: E^I \rightarrow B^I$. Here $I = [0, 1]$, and the mapping spaces $E^I = \text{map}(I, E)$ and $B^I = \text{map}(I, B)$ come with the compact-open topology. The Serre microfibration property for p^I is obvious, so it is enough to establish the weak contractibility of the fibers of p^I . Suppose therefore given a map $\gamma: I \rightarrow B$ and a map $f: \mathbb{S}^n \times I \rightarrow E$ which covers γ , so that $pf(z, t) = \gamma(t)$ for $z \in \mathbb{S}^n$ and $t \in I$. We must extend f to a map $g: \mathbb{D}^{n+1} \times I \rightarrow E$ which covers γ . But that is easy: Use a sufficiently fine subdivision of I into subintervals $[a_r, a_{r+1}]$ so that partial extensions

$$g_r: \mathbb{D}^{n+1} \times [a_r, a_{r+1}] \rightarrow E$$

of f can be constructed, with $pg_r(z, t) = \gamma(t)$ for $z \in \mathbb{D}^{n+1}$ and $t \in [a_r, a_{r+1}]$. Then improve g_r if necessary, on a small neighborhood of $\mathbb{D}^{n+1} \times \{a_r\}$ in $\mathbb{D}^{n+1} \times [a_r, a_{r+1}]$, to ensure that $g_r(z, a_r) = g_{r-1}(z, a_r)$ for $z \in \mathbb{D}^{n+1}$.

The induction beginning consists in showing that p has the path lifting property. (That is, given a path $\gamma: I \rightarrow B$ and $a \in E$ with $p(a) = \gamma(0)$, there exists a path $\omega: I \rightarrow E$ with $p\omega = \gamma$ and $\omega(0) = a$.) But that is also easy. \square

Lemma 2.3. *Let τ be the tautological functor from $\mathcal{O}(X)$ to spaces and let K be a compact subset of $\text{hocolim } \tau$. Then there exist a finite full sub-poset $\mathcal{P} \subset \mathcal{O}(X)$ and a subfunctor κ of $\tau|_{\mathcal{P}}$ with compact values such that $K \subset \text{hocolim } \kappa$.*

Remarks. The fullness assumption means that $U, V \in \mathcal{P}$ and $U \subset V$ imply $U \leq V$ in \mathcal{P} . By a subfunctor κ of $\tau|_{\mathcal{P}}$ is meant a selection of subspaces $\kappa(U) \subset \tau(U) = U$, one for each $U \in \mathcal{P}$, such that $\kappa(V) \subset \kappa(U)$ if $V \leq U$ in \mathcal{P} .

The lemma is closely related to an observation for which I am indebted to Larry Taylor: *The mapping cylinder C of the inclusion of the open unit interval in the closed unit interval is not homeomorphic to a subset of $[0, 1]^2$.* This is easy to verify, although surprising. The two endpoints of the closed unit interval, viewed as elements of the mapping cylinder C , don't have countable neighborhood bases; hence C is not even metrizable. Equally surprising, and more to the point, is the following. Let K be a compact subset of C . Then there exists a compact subinterval L of the open unit interval such that K is contained in the mapping cylinder of the inclusion $L \rightarrow [0, 1]$. For the proof, exhaust the open unit interval by an ascending sequence of compact subintervals L_i . Suppose if possible that for each i there exists $x_i \in K$ which is not contained in the mapping cylinder of the inclusion $L_i \rightarrow [0, 1]$. Then the x_i form an infinite discrete closed subset of K , which contradicts the compactness of K .

Proof of lemma 2.3. The classifying space $B\mathcal{O}(X)$ is a simplicial complex. This has one n -simplex for each subset of $\mathcal{O}(X)$ of the form $\{U_0, U_1, \dots, U_n\}$ where U_{i-1} is a proper subset of U_i , for $i \in \{1, 2, \dots, n\}$. The image of C under the projection $\text{hocolim } \tau \rightarrow B\mathcal{O}(X)$ is contained in a compact simplicial subcomplex $B\mathcal{P}$, and without loss of generality we can assume that the subcomplex has the form $B\mathcal{P}$ for a finite full sub-poset \mathcal{P} of $\mathcal{O}(X)$. For each simplex S of $B\mathcal{P}$, let $e(S) \subset B\mathcal{P}$ be the "cell" determined by S , so that $e(S)$ is locally closed in $B\mathcal{P}$ and $B\mathcal{P}$ is the disjoint union (but not the coproduct in general) of the $e(S)$ for the simplices S of

$B\mathcal{P}$. Let $U(S)$ be the smallest of the open sets corresponding to the vertices of S . The inverse image of $e(S)$ for the projection

$$\text{hocolim } \tau \longrightarrow \text{BO}(X)$$

is identified with $e(S) \times U(S)$. Its intersection with K is contained in a subset of the form $e(S) \times L(S)$, where $L(S) \subset U(S)$ is compact. (This can be proved as in the remark just above.) Choose such an $L(S)$ for every simplex S in $B\mathcal{P}$. For an element U of \mathcal{P} let

$$\kappa(U) := \bigcup_{S \text{ with } U(S) \subset U} L(S).$$

Then $\kappa(V) \subset \kappa(U)$ for $V, U \in \mathcal{P}$ with $V \subset U$, and each $\kappa(U)$ is compact. \square

Corollary 2.4. *The projection $p_{\mathcal{F}}: X_{\mathcal{F}} \rightarrow X$ is a Serre microfibration.*

Proof. Write $p = p_{\mathcal{F}}$. Let $q: X_{\mathcal{F}} \rightarrow BT_{\mathcal{F}}$ be the standard projection (from the homotopy colimit to the classifying space of the indexing category). As we just discovered, the formula $y \mapsto (p(y), q(y))$ need not define an embedding of $X_{\mathcal{F}}$ in $X \times BT_{\mathcal{F}}$, but it certainly defines an injective map and we can use that to label elements of $X_{\mathcal{F}}$. In particular, let

$$h: \mathbb{D}^i \times [0, 1] \longrightarrow X$$

be a homotopy with an ‘‘initial lift’’ $H_0: \mathbb{D}^i \rightarrow X_{\mathcal{F}}$, so that $h(z, 0) = pH_0(z)$. We need to find $\varepsilon > 0$ and a map $H: \mathbb{D}^i \times [0, \varepsilon] \rightarrow X_{\mathcal{F}}$ such that $pH = h$ on $\mathbb{D}^i \times [0, 1]$ and $H(z, 0) = H_0(z)$ for $z \in \mathbb{D}^i$. The plan is to define H in such a way that

$$(pH(z, t), qH(z, t)) = (h(z, t), qH_0(z)),$$

for $(z, t) \in \mathbb{D}^i \times [0, \varepsilon]$, which means that $qH: \mathbb{D}^i \times [0, 1] \rightarrow BT_{\mathcal{F}}$ is a constant homotopy. By lemma 2.3, the plan is sound, giving a well defined and continuous map $\mathbb{D}^i \times [0, \varepsilon] \rightarrow X_{\mathcal{F}}$ for sufficiently small ε . \square

Lemma 2.5. *The projection $p_{\mathcal{F}}: X_{\mathcal{F}} \rightarrow X$ has contractible fibers.*

Proof. The fiber over $x \in X$ is identified with the homotopy colimit of the (contravariant, set-valued) functor

$$(U, c) \mapsto \mathcal{F}^{(c)}(U)$$

where U runs through the open subsets of X containing x , and c runs through the objects of \mathcal{C} . By a well-known Fubini principle for homotopy colimits, it is homotopy equivalent to the double homotopy colimit

$$\text{hocolim}_c \text{hocolim}_{U \ni x} \mathcal{F}^{(c)}(U).$$

In this expression the inside homotopy colimit is a homotopy colimit of sets (i.e., discrete spaces) taken over a directed poset, and therefore the canonical map

$$\text{hocolim}_{U \ni x} \mathcal{F}^{(c)}(U) \longrightarrow \text{colim}_{U \ni x} \mathcal{F}^{(c)}(U)$$

is a homotopy equivalence. Therefore

$$\text{hocolim}_c \text{hocolim}_{U \ni x} \mathcal{F}^{(c)}(U) \simeq \text{hocolim}_c \mathcal{F}_x^{(c)}$$

where $\mathcal{F}_x^{(c)}$ is the stalk of $\mathcal{F}^{(c)}$ at x . But the stalk functor \mathcal{F}_x is representable by assumption. The homotopy colimit of a representable functor is contractible. \square

Proof of proposition 2.1. Apply lemma 2.2 and note that a Serre fibration with weakly contractible fibers is a weak homotopy equivalence. \square

3. CLASSIFICATION OF SHEAVES UP TO CONCORDANCE

Lemma 3.1. *Let \mathcal{F}_0 and \mathcal{F}_1 be two sheaves of \mathcal{C} -sets on X , both with representable stalks. Let $g: \mathcal{F}_0 \rightarrow \mathcal{F}_1$ be a binatural transformation. Then \mathcal{F}_0 and \mathcal{F}_1 are concordant.*

Proof. Let $e_0: X \rightarrow X \times I$ be given by $e_0(x) = (x, 0)$ and let $p: X \times I \rightarrow X$ be the projection. For an object c in \mathcal{C} and an open subset U of $X \times [0, 1]$, let $U_0 = e_0^{-1}(U)$ and let $\mathcal{G}^{(c)}(U)$ be the set of pairs $(s, t) \in \mathcal{F}_0(U_0) \times p^*\mathcal{F}_1(U)$ such that $gs = e_0^*t$. Now \mathcal{G} is a sheaf of \mathcal{C} -sets on $X \times I$ with representable stalks. Its restrictions to $X \times \{0\}$ and $X \times \{1\}$ are identified with \mathcal{F}_0 and \mathcal{F}_1 , respectively. \square

Corollary 3.2 (to proposition 2.1 and lemma 3.1.). *Let \mathcal{F} be a sheaf of \mathcal{C} -sets on X with representable stalks. Suppose that X is a CW -space. Then*

$$\pi_{\mathcal{F}} p_{\mathcal{F}}^{-1}: X \rightarrow BC$$

is a classifying map for \mathcal{F} . That is, $(\pi_{\mathcal{F}} p_{\mathcal{F}}^{-1})^ \mathcal{E}$ is concordant to \mathcal{F} , with \mathcal{E} as in definition 1.1.*

Proof. Abbreviate $p = p_{\mathcal{F}}$, $\pi = \pi_{\mathcal{F}}$. It is enough to show that the sheaves $\pi^* \mathcal{E}$ and $p^* \mathcal{F}$ on $X_{\mathcal{F}}$ are concordant. By lemma 3.1, it is then also enough to make a map from $\pi^* \mathcal{E}$ to $p^* \mathcal{F}$. That is what we will do, using the “étale” point of view. Therefore let $z \in X_{\mathcal{F}}$. We need to compare the stalk of \mathcal{F} at $p(z) \in X$ with the stalk of \mathcal{E} at $\pi(z) \in BC$. The point z maps to some cell in $BT_{\mathcal{F}}$ which corresponds to a nondegenerate diagram

$$(U_0, c_0) \leftarrow (U_1, c_1) \leftarrow \cdots \leftarrow (U_{k-1}, c_{k-1}) \leftarrow (U_k, c_k)$$

in $\mathcal{O}(X) \times \mathcal{C}$, with $p(z) \in U_k$, and an element $s_0 \in \mathcal{F}^{(c_0)}(U_0)$. The stalk of \mathcal{E} at $\pi(z)$ is then represented by the object c_0 . The germ of s_0 near $p(z)$ amounts to a morphism from c_0 to the object which represents the stalk of \mathcal{F} at $p(z)$; equivalently, by the Yoneda lemma, s_0 determines a \mathcal{C} -map from the stalk of \mathcal{E} at $\pi(z)$ to the stalk of \mathcal{F} at $p(z)$. Letting z vary now, and selecting an object c in \mathcal{C} , we obtain a map over $X_{\mathcal{F}}$ from $\text{Spé}(\pi^* \mathcal{E}^{(c)})$ to $\text{Spé}(p^* \mathcal{F}^{(c)})$. This is continuous (verification left to the reader) and natural in c , and therefore amounts to a map between sheaves of \mathcal{C} -sets on $X_{\mathcal{F}}$, from $\pi^* \mathcal{E}$ to $p^* \mathcal{F}$. \square

Proof of theorem 0.1. Suppose that X is a CW -space. Let $g: X \rightarrow BC$ be any map and put $\mathcal{F} = g^* \mathcal{E}$. We have to show that g is homotopic to πp^{-1} , where $\pi = \pi_{\mathcal{F}}$ and $p = p_{\mathcal{F}}$. We note that $X_{\mathcal{F}}$ also has the homotopy type of a CW -space since it is a homotopy colimit of open subsets of X (all of which have the homotopy type of CW -spaces). Hence p is a homotopy equivalence. Therefore, showing $g \simeq \pi p^{-1}$ amounts to showing that $gp \simeq \pi$.

Now recall that $X_{\mathcal{F}}$ was constructed as the homotopy colimit of a functor $\tau\varphi$ from a certain category $\mathcal{T}_{\mathcal{F}}$ with objects (U, c, s) to the category of spaces. The maps $gp: X_{\mathcal{F}} \rightarrow BC$ and $\pi: X_{\mathcal{F}} \rightarrow BC$ both have a factorization of the following kind:

$$\text{hocolim}_{(U, c, s)} U \xrightarrow{v} \text{hocolim}_c \text{Spé}(\mathcal{E}^{(c)}) \xrightarrow{w} BC.$$

Here v is (in both cases) induced by a natural transformation from the functor $(U, c, s) \mapsto U$ to the functor $(U, c, s) \mapsto \text{Spé}(\mathcal{E}^{(c)})$, given by the maps

$$U \longrightarrow \text{Spé}(\mathcal{E}^{(c)}) \quad ; \quad x \mapsto (x, s).$$

In the case of gp , the second map $w = w_0$ in the factorization is determined by the projections $\text{Spé}(\mathcal{E}^{(c)}) \rightarrow BC$. In the case of π , the map $w = w_1$ is the composition of the projection to $BT_{\mathcal{F}}$ and the forgetful map $BT_{\mathcal{F}} \rightarrow BC$.

Consequently it is now sufficient to show that w_0 and w_1 are *weakly homotopic*, which is to say, $w_0 f \simeq w_1 f$ for any map f from a CW -space to the common source of w_0 and w_1 . It is enough to check this for a particular f which is a weak equivalence. A good choice of such an f is the map

$$\text{hocolim}_c B(c \downarrow \mathcal{C}) \longrightarrow \text{hocolim}_c \text{Spé}(\mathcal{E}^{(c)})$$

induced by the natural maps $\lambda_c: B(c \downarrow \mathcal{C}) \rightarrow \text{Spé}(\mathcal{E}^{(c)})$. By proposition 1.3, this f is indeed a weak homotopy equivalence. The maps $w_0 f$ and $w_1 f$ are easily seen to be homotopic. Indeed, each $(c \downarrow \mathcal{C})$ has two obvious functors to \mathcal{C} , one given by $(c \rightarrow d) \mapsto d$ and the other by $(c \rightarrow d) \mapsto c$. These are related by a natural transformation, which determines a homotopy $h^{(c)}$ between the induced maps $B(c \downarrow \mathcal{C}) \rightarrow BC$. Integrating the homotopies $h^{(c)}$ one obtains a homotopy $w_0 f \simeq w_1 f$. \square

REFERENCES

1. A. Dold, *Lectures on algebraic topology*, Grundlehren, Springer, 1972.
2. R. Hartshorne, *Algebraic geometry*, Graduate Texts in Mathematics, Springer, 1977.
3. J.P.May, *Simplicial objects in algebraic topology*, Chicago Lectures in Math., Univ. of Chicago Press, 1967.
4. S. MacLane and I. Moerdijk, *Sheaves in geometry and logic*, Springer, 1992.
5. I. Madsen and M. Weiss, *The stable moduli space of Riemann surfaces: Mumford's conjecture*, preprint, arXiv:math.AT/0212321, v3, 2002–2004.
6. I. Moerdijk, *Classifying spaces and classifying topoi*, Lecture Notes in Math., vol. 1616, Springer, 1995.
7. G. B. Segal, *Classifying spaces related to foliations*, *Topology* **17** (1978), 367–382.

DEPT. OF MATHEMATICS, UNIVERSITY OF ABERDEEN, AB24 3UE
E-mail address: m.weiss@maths.abdn.ac.uk