THE BROWN REPRESENTABILITY THEOREM, OLD AND NEW

MARK MACERATO AND SAAD SLAOUI

"Go beyond the shadow on the wall."

Contents

1. Introduction

Following the precise formulation of a set of axioms characterizing generalized cohomology theory by Eilenberg and Steenrod, the late 1950s and early 1960s were marked by the intensified study of such theories, particularly real and complex topological K-theory and various cobordism theories. In 1962, Edgar Brown published a paper [1] in which he gave a sufficient set of conditions under which a set-valued functor out of the homotopy category of CW complexes was representable, i.e. naturally isomorphic to a functor of the form $[-, K]$ for some based CW complex K. Strikingly, the hypotheses of the theorem were weak enough to be satisfied by any reduced generalized cohomology theory as axiomatized by Eilenberg and Steenrod. Brown's insight thereby provided an impulse to the study of stable homotopy theory, for reasons we describe briefly in Section 4.

The main objective of these notes is to give a comprehensive and accessible treatment of the proof of the Brown representability theorem. We mainly follow the argument given in Section 3.4 of Kochman [4], paying special attention to the details. We also take some time to describe the modern version of this theorem and its proof, based on Lurie's work in [2] and [3], with the hope that this juxtaposition will highlight the formal nature of Brown's original proof on the one hand, and the classical motivations underlying Lurie's modern approach on the other.

2. Definitions and statement of the theorem

Definition 2.1. A functor $h : \mathrm{h}\mathbf{CW}_*^{op} \to \mathbf{Set}_*$ from the homotopy category of based CW complexes to the category of pointed sets is said to be a *Brown functor* if it satisfies the following two properties:

- (i) **(Wedge axiom)** Given a wedge of spaces $\bigvee_i X_i$, the based inclusion maps $j_i : X_i \hookrightarrow \bigvee_i X_i$ induce a bijection:
$$h(\textstyle\bigvee_i X_i) \xrightarrow{(j_i^*)_i} \prod_i h(X_i),$$
under the convention $j_i^* := h(j_i)$.
- (ii) **(Mayer-Vietoris axiom)** Given a based CW complex X and subcomplexes $A, B \subset X$ such that the union of the interiors of A and B covers X, the induced commutative square

$$\begin{array}{ccc} h(X) & \xrightarrow{l^*_{A,X}} & h(A) \\ {\scriptstyle l^*_{B,X}}\downarrow & & \downarrow{\scriptstyle l^*_{A\cap B,A}} \\ h(B) & \xrightarrow[l^*_{A\cap B,B}]{} & h(A\cap B), \end{array} \tag{2.2}$$

satisfies the following condition: whenever $x \in h(A)$, $y \in h(B)$ are such that $l^*_{A\cap B,A}(x) = l^*_{A\cap B,B}(y)$ in $h(A \cap B)$, there is a $z \in h(X)$ such that $l^*_{A,X}(z) = x$ and $l^*_{B,X}(z) = y$. In other words, the canonical map
$$h(X) \to h(A) \times_{h(A\cap B)} h(B)$$
is surjective.

Remarks 2.3.
- The inclusion $\mathrm{pt} \xrightarrow{j_1} \mathrm{pt} \vee \mathrm{pt}$ of the one point space into its wedge product with itself is a homeomorphism. Thus, by the wedge axiom, we have a commutative triangle:

$$\begin{array}{ccc} h(\mathrm{pt}) \cong h(\mathrm{pt}\vee\mathrm{pt}) & \xrightarrow{\cong} & h(\mathrm{pt}) \times h(\mathrm{pt}) \\ & \searrow^{j_1^*}_{\cong} & \downarrow{\scriptstyle \pi_1} \\ & & h(\mathrm{pt}) \end{array}$$

where both the horizontal and diagonal arrows are bijections. It follows that the projection map on the first factor $h(\mathrm{pt}) \times h(\mathrm{pt}) \xrightarrow{\pi_1} h(\mathrm{pt})$ must be a bijection, so that $h(\mathrm{pt}) = \{v\}$ must be a one-point set (and likewise for any contractible space by homotopy invariance).
- Squares such as 2.2 satisfying stated property are sometimes called *weak pullback squares*.
- When talking about the image of an element $x \in h(X)$ under the map induced by an inclusion $A \xrightarrow{l} X$, we shall often use the notation $x|_A \in h(A)$, intuitively thought of as the "restriction" of x to the subspace A. In this language, the Mayer-Vietoris axiom states that whenever two elements $x \in h(A)$, $y \in h(B)$ restrict to the same element in $h(A \cap B)$, they can be expressed as the restrictions of a common element in $h(X)$.

Examples 2.4. (1) For any given based CW complex Y, we claim that the representable functor
$$[-, Y] : \mathrm{h}\mathbf{CW}_*^{op} \to \mathbf{Set}_*$$
is a valid Brown functor:

Firstly, the wedge axiom is immediately satisfied by the defining universal property of coproducts, namely that to specify a map out of $\bigvee_i X_i$ is equivalent to specifying a map out of each X_i. Next, start with a based CW complex X and subcomplexes $A, B \subset X$ as in Definition 2.1 (ii), and suppose we are given representatives $A \xrightarrow{f} Y$, $B \xrightarrow{g} Y$ of classes of maps in $[A, Y]$, resp. $[B, Y]$, with the property that the restrictions $f|_{A\cap B} \simeq g|_{A\cap B}$ are homotopic. By the gluing lemma, it suffices to exhibit a map $f' \simeq f$ with the property that $f'|_{A\cap B} = g|_{A\cap B}$ on the nose, in which case we know that we can construct a map $G : X \to Y$ such that $G|_A = f' \simeq f$, $G|_A = g$, as needed. Now, let $h : A \cap B \times I \to Y$ be the provided homotopy between $f|_{A\cap B}$ and $g|_{A\cap B}$. Since inclusions of subcomplexes are cofibrations, we may use the homotopy extension property on the following diagram:

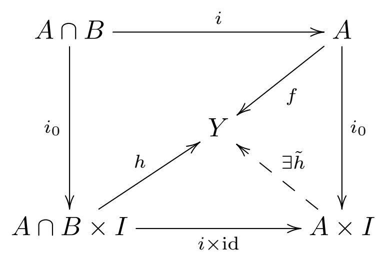

to get a homotopy $\tilde{h} : A \times I \to Y$ from f to some map $f' := \tilde{h}|_{A\times\{1\}}$ such that $f'|_{A\cap B} = h|_{(A\cap B)\times\{1\}} = g|_{A\cap B}$, completing the proof.

(2) As mentioned in the introduction, any functor $h^n : \mathrm{h}\mathbf{CW}_*^{op} \to \mathbf{Set}_*$ appearing as the n^{th}component of a reduced generalized cohomology theory (for some $n \geq 0$) is a Brown functor. The wedge axiom is one of the defining axioms for h^n, and the Mayer Vietoris axiom can be deduced from the remaining Eilenberg-Steenrod axioms.

Theorem 2.5. ***(Brown Representability Theorem)*** *Any Brown functor is representable. That is, given a Brown functor $h : \mathrm{h}\boldsymbol{CW}_*^{op} \to \boldsymbol{Set}_*$, there exists a based CW complex K and an element $u \in h(K)$ such that the assignment:*

$$T_u : [-, K] \longrightarrow h(-)$$

given by sending a map $X \xrightarrow{f} K$ to $f^*(u) \in h(X)$ is a natural isomorphism. Furthermore, such a CW complex K is unique up to homotopy equivalence.

Remarks 2.6. (1) The uniqueness part of the above theorem is a formal consequence of the Yoneda lemma, which guarantees that any two objects K, L with naturally isomorphic represented functors $h_K \simeq h_L$ are themselves isomorphic (in our settings, homotopy equivalent).

(2) The assignment T_u defined above is natural for any choice of $K \in \mathbf{CW}_*$, $u \in h(K)$. Indeed, given a map $X \xrightarrow{g} Y$, we may look at the square:

$$\begin{array}{ccc} [X, K] & \xleftarrow{\circ g} & [Y, K] \\ {\scriptstyle T_u}\downarrow & & \downarrow{\scriptstyle T_u} \\ h(X) & \xleftarrow[g^*]{} & h(Y) \end{array} ,$$

where the top horizontal map is given by precomposition. Then, picking an arbitrary element $[f] \in [Y K]$, we see that:

$$g^* T_u([f]) = g^* f^*(u) = (f \circ g)^*(u) = T_u([f \circ g]),$$

proving commutativity of the above square and hence naturality.

3. Proof of the Brown representability theorem

Throughout the following, we fix a Brown functor $h\colon \mathrm{h}\mathbf{CW}_*^{op} \to \mathbf{Set}_*$.

Lemma 3.1. *Let $f\colon X \to Y$ be a based map between pointed CW complexes, with associated cofiber sequence $X \xrightarrow{f} Y \xrightarrow{j} Cf$. Then the induced sequence of pointed sets:*

$$h(Cf) \xrightarrow{j^*} h(Y) \xrightarrow{f^*} h(X)$$

is exact (with respect to the convention that $\ker f^*$ *is the set of elements $y \in h(Y)$ mapped under f^* to the basepoint of $h(X)$).*

Proof. First recall that Cf may be obtained as the pushout:

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ \downarrow & & \downarrow{\scriptstyle j} \\ CX & \longrightarrow & Cf, \end{array}$$

where CX is the cone on X, a contractible space, so that by Remarks 2.3 (1) $h(CX) = \{v\}$ is a one-point set and we obtain a commutative square:

$$\begin{array}{ccc} h(X) & \xleftarrow{f^*} & h(Y) \\ \uparrow & & \uparrow{\scriptstyle j^*} \\ \{v\} & \longleftarrow & h(Cf). \end{array}$$

Hence, the composite $f^* j^*$ factors through the zero map, implying that $\mathrm{im} j^* \subset \ker f^*$.

Next, we check that $\ker f^* \subset \mathrm{im} j^*$. Start with an element $y \in h(Y)$ such that $f^*(y) = 0$, the basepoint in $h(X)$. Write $Cf = \mathrm{int}A \cup \mathrm{int}B$, where we set $A := X \times [0, 3/4]$ and $B := (X \times [1/4, 1]) \cup_f Y$, so that we have homotopies $A \simeq \{\mathrm{pt}\}$, $B \simeq Y$, $A \cap B \simeq X$ induced by the obvious inclusions. Our plan is to exploit what we know about the Cartesian square induced by A, B, $A \cap B$ and Cf via the Mayer-Vietoris axiom to get information about the square induced by the corresponding homotopy equivalent spaces Y, $\{pt\}$, X, Cf respectively. We start by observing that we can form a cube, commutative up to homotopy:

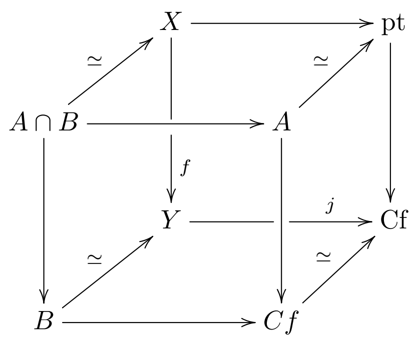

Then, by homotopy invariance, applying $h(-)$ results in a commutative diagram of pointed sets:

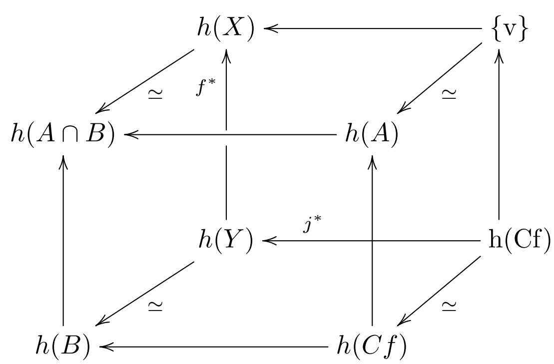

where the front square is a weak pullback square. It then follows by diagram chasing that, since the elements $y \in h(Y), v \in h(\mathrm{pt})$ satisfy $y|_X = v|_X = 0$, we can find an element $x \in h(Cf)$ such that $j^*(x) = y$, as needed. ☐

Remark 3.2. In what follows, we will make implicit use of the above method, whereby we deduce that a given square is a weak pullback square by means of pointwise homotopies.

We now start moving towards a proof of the main theorem.

Definition 3.3. Given a based CW complex K, say that $u \in h(K)$ is *n-universal* if $T_u : [S^t, K] \to h(S^t)$ is an isomorphism for $t < n$ and an epimorphism for $t = n$. If $u \in h(K)$ is n-universal for all n, say that u is *universal*.

Remark 3.4. The condition for being (-1)-universal is vacuous, so that any element $u \in h(K)$ is (-1)-universal.

Lemma 3.5. *Let K be a based CW complex and $u_n \in h(K)$ an n-universal element (for some $n \geq -1$). Then there exists a based CW complex $L \supset K$ obtained from K by attaching $(n+1)$-cells, together with an $(n+1)$-universal element $u_{n+1} \in h(L)$ such that $u_{n+1}|_K = u_n$.*

Proof. Since u_n is n-universal, we may look at the epimorphism:

$$T_{u_n} : [S^n, K] = \pi_n(K) \longrightarrow h(S^n).$$

Let B be a system of representatives for the classes in $\ker(T_{u_n}: \pi_n(K) \to h(S^n))$, so that each $\beta \in B$ is a map $S^n \xrightarrow{\beta} K$ corresponding to some class in the kernel. For the case $n = -1$, set B to be the empty set.

Now, form the based spaces $X := \bigvee_{\beta \in B} S^n$, $Y := K \vee \bigvee_{\alpha \in h(S^{n+1})} S^{n+1}$. We may consider the composite:

$$f\colon X \xrightarrow{\bigvee_{\beta \in B} \beta} K \xrightarrow{l} Y, \tag{3.6}$$

where $K \xrightarrow{l} Y$ is given by inclusion. We then have an associated cofiber sequence:

$$X \xrightarrow{f} Y \xrightarrow{j} Cf, \tag{3.7}$$

and we define $L := Cf$ to be the mapping cone of f. Observe that L was constructed by trivially attaching one $(n+1)$-cell to K for each $\alpha \in h(S^{n+1})$, as well as one $(n+1)$-cell for each element in $\ker(T_{u_n}: \pi_n(K) \to h(S^n))$ using the associated map $S^n \xrightarrow{\beta} K$ as attaching map. To complete the proof, it suffices to exhibit an $(n+1)$-universal element $u_{n+1} \in h(L)$ such that $u_{n+1}|_K = u_n$.

First, observe that by Lemma 3.1 the cofiber sequence 3.7 induces an exact sequence:

$$h(L) \xrightarrow{j^*} h(Y) \xrightarrow{f^*} h(X), \tag{3.8}$$

where $h(Y) \cong h(K) \times \prod_{\alpha \in h(S^{n+1})} h(S^{n+1})$ and $h(X) \cong \prod_{\beta \in B} h(S^n)$ by the wedge axiom. In particular, we may look at the element:

$$u'_n := \big(u_n, (\alpha)_{\alpha \in h(S^{n+1})}\big) \in h(Y)$$

obtained under this isomorphism. Now, we have a commutative diagram:

$$\begin{array}{ccccc} h(X) & \xleftarrow{f^*} & h(Y) & \xrightarrow{\simeq} & h(K) \times \prod_{\alpha \in h(S^{n+1})} h(S^{n+1}), \\ & {\scriptstyle (\beta^*)_{\beta \in B}} \nwarrow & \downarrow {\scriptstyle i^*} & \swarrow {\scriptstyle \pi_1} & \\ & & h(K) & & \end{array}$$

so that $f^*(u'_n) = (\beta^*(u_n))_{\beta \in B} = 0$, since each $\beta^*(u_n) = T_{u_n}(\beta) = 0$ by construction. It follows by exactness of 3.8 that we can find an element $u_{n+1} \in h(L)$ such that $j^*(u_{n+1}) = u'_n$, hence in particular $u_{n+1}|_K = u_n$ by definition of u'_n.

It remains to show that $u_{n+1} \in h(L)$ as obtained above is $(n+1)$-universal. For a given $t \geq 0$, we may look at the following triangle:

$$\begin{array}{ccc} \pi_t(K) & \xrightarrow{i_*} & \pi_t(L), \\ {\scriptstyle T_{u_n}} \searrow & & \swarrow {\scriptstyle T_{u_{n+1}}} \\ & h(S^t) & \end{array}$$

where i_* is the induced map on homotopy groups. This diagram commutes, since for every $[g] \in \pi_t(K) = [S^t, K]$ we have that:

$$T_{u_{n+1}}(i_*[g]) = T_{u_{n+1}}([ig]) = g^*(i_* u_{n+1}) = g^*(u_{n+1}|_K) = g^*(u_n) = T_{u_n}([g]).$$

Now, for $t < n$, $T_{u_n} : \pi_t(K) \to h(S^t)$ is an isomorphism by assumption, and $i_* : \pi_t(K) \to \pi_t(L)$ is also an isomorphism since the inclusion $K \xrightarrow{i} L$ is an n-equivalence. Hence, $T_{u_{n+1}} : \pi_t(L) \to h(S^t)$ must also be an isomorphism. For $t = n$, since $T_{u_n} : \pi_n(K) \to h(S^n)$ is a surjection, so is $T_{u_{n+1}} : \pi_n(L) \to h(S^n)$. Thus it remains to check that (1): $T_{u_{n+1}} : \pi_n(L) \to h(S^n)$ is an injection, and (2): $T_{u_{n+1}} : \pi_{n+1}(L) \to h(S^{n+1})$ is a surjection.

(1) Start with an element $[x] \in \ker(T_{u_{n+1}} : \pi_n(L) \to h(S^n))$. Since i is an n-equivalence, the map $i_* : \pi_n(K) \to \pi_n(L)$ is surjective. Hence, we can find an element $[y] \in \pi_n(K)$ with $i_*([y]) = [x]$. By commutativity,

$$T_{u_n}([y]) = T_{u_{n+1}}(i_*[y]) = T_{u_{n+1}}([x]) = 0,$$

hence $[y] \in \ker(T_{u_n} : \pi_n(K) \to h(S^n))$, so that by construction we can find a representative $\beta_0 \in B \subset \pi_n(K)$ with $y \simeq \beta_0$. Now, letting $l_{\beta_0} : S^n \hookrightarrow X = \bigvee_{\beta \in B} S^n$ denote the inclusion into the β_0^{th} factor, we have by definition that $(\bigvee_{\beta \in B} \beta) \circ l_{\beta_0} = \beta_0 : S^n \to K$, so that

$$(\bigvee_{\beta \in B} \beta)_*[l_{\beta_0}] = [\beta_0] = [y]. \tag{3.9}$$

In the notation of 3.6 and 3.7, the situation so far is summarized in the following commutative diagram:

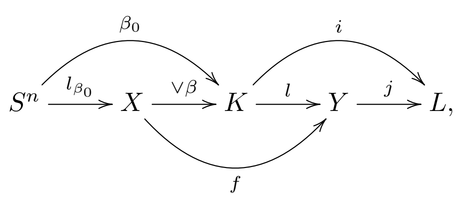

in which the composite $j \circ f$ is null-homotopic since it fits into a cofiber sequence. Thus, by Equation 3.9, we see that:

$$x = i_*([y]) = i_*(\vee\beta)_*[l_{\beta_0}] = [i(\vee\beta)l_{\beta_0}] = [jfl_{\beta_0}] = 0,$$

as needed.

(2) Next, start with an arbitrary element $\alpha_0 \in h(S^{n+1})$, to which we can associate an inclusion map $l_{\alpha_0} : S^{n+1} \hookrightarrow Y = K \vee \bigvee_{\alpha \in h(S^{n+1})} S^{n+1}$ in the α_0^{th} factor. Then, we may look at the composite:

$$S^{n+1} \xrightarrow{l_{\alpha_0}} Y = K \vee \textstyle\bigvee_{\alpha \in h(S^{n+1})} S^{n+1} \xrightarrow{j} L,$$

with $[jl_{\alpha_0}] \in \pi_{n+1}(L)$. As mentioned above, $j^*(u_{n+1}) = u_n'$, so that by construction:

$$T_{u_{n+1}}([jl_{\alpha_0}]) = l_{\alpha_0}^* j^*(u_{n+1}) = l_{\alpha_0}(u_n') = \alpha_0,$$

completing the proof. □

To prove Theorem 2.5, we must construct a space K together with a class $u \in h(K)$ such that the natural transformation T_u is an isomorphism. Such a class u must in particular be universal - and since we are working with CW complexes, Whitehead's lemma suggests that universality may actually suffice. Lemma 3.5 constructs an $(n+1)$-universal class u_{n+1} on a space K_{n+1} from the data of a space K_n with an n-universal class u_n; we now apply the lemma inductively to construct a universal class.

Lemma 3.10. *Let K be a based CW complex and $v \in h(K)$. Then there exists a based CW complex L, a universal class $u \in h(L)$, and an inclusion of based CW complexes $K \subset L$ such that $u|_K = v$.*

Proof. As indicated in Remark 3.4, the condition of being (-1)-universal is vacuous. Hence, v is (-1)-universal. We let $L_{-1} = K$ and $u_{-1} = v$. Applying Lemma 3.5 repeatedly yields an ascending chain of based CW complexes

$$L_{-1} \subset L_0 \subset L_1 \subset L_2 \subset \cdots,$$

together with elements $u_n \in h(L_n)$, such that (i) u_n is n-universal, and (ii) whenever $m \leq n$, $u_n|_{L_m} = u_m$. Let $i_n : L_n \hookrightarrow L_{n+1}$ denote the inclusion. Let L denote the mapping telescope of the above sequence of inclusions; that is,

$$L = \left(\coprod_{n=-1}^{\infty} L_n \times [n-1, n] \right) \Big/ \sim .$$

Here, the equivalence relation $\sim$ is defined by (i) $(x, n) \sim (i_n(x), n)$ for all $x \in L_n$, $n \geq -1$, and (ii) $(\star_n, t) \sim (\star_m, t')$ for any $t \in [n-1, n]$, $t' \in [m-1, m]$, $n, m \geq -1$ (where $\star_n$ denotes the basepoint of L_n). The basepoint of L is the image of $(\star_n, t)$ for any $n \geq -1$ and $t \in [n-1, n]$. Let S_n denote the product $L_n \times [n-1, n]$ appearing above.

We determine precisely when a point $(x, t) \in S_n$ is equivalent to $(x', t') \in S_m$ under the relation $\sim$, if $|n - m| \geq 2$. Since $[n-1, n] \cap [m-1, m] = \emptyset$, it follows from the definition of $\sim$ that $x = \star_n$, $x' = \star_m$. Conversely, $(x, t) \sim (x', t')$ as soon as x and x' are basepoints. We conclude that the images of S_n and S_m in L intersect only at the basepoint $\star$. Next, observe that the image of S_n in L is (based) homeomorphic to $L_n \wedge [n-1, n]^+$. To put these observations to use, define

$$A = \left(\coprod_{n=-1}^{\infty} L_{2n+1} \times [2n, 2n+1] \right) \Big/ \sim,$$

$$B = \left(\coprod_{n=0}^{\infty} L_{2n} \times [2n-1, 2n] \right) \Big/ \sim .$$

There is a slight technicality: since we are aiming to apply the Mayer-Vietoris axiom to the pair (A, B), their interiors should cover L. Therefore, we should 'thicken' these spaces, i.e. replace the interval $[n-1, n]$ with $[n-1-\epsilon, n+\epsilon]$ for some $0 < \epsilon < 1$, where appropriate. Of course, these new spaces deformation retract onto A and B, so we refrain from making this replacement explicit for the sake of notation. We conclude from the discussion above that

$$A \simeq \bigvee_{n \geq -1} L_{2n+1} \wedge [2n, 2n+1]^+,$$
$$B \simeq \bigvee_{n \geq 0} L_{2n} \wedge [2n-1, 2n]^+,$$

and consequently that

$$A \simeq \bigvee_{n \geq -1} L_{2n+1},$$
$$B \simeq \bigvee_{n \geq 0} L_{2n}.$$

Furthermore,

$$\begin{aligned} C := A \cap B &\cong \left(\coprod_{n \geq -1} L_n \times \{n\} \right) \Big/ \sim \\ &\cong \bigvee_{n \geq -1} L_n. \end{aligned}$$

We can now apply the wedge axiom to construct classes $a \in h(A)$, $b \in h(B)$, and $c \in h(C)$ such that $a|_{L_{2n+1}} = u_{2n+1}$ $(n \geq -1)$, $b|_{L_{2n}} = u_{2n}$ $(n \geq 0)$, and $c|_{L_n} = u_n$ $(n \geq -1)$. Since $a|_{L_{2n+1}} = c|_{L_{2n+1}}$ for all $n \geq -1$, we conclude that $a|_C = c$ by the wedge axiom. Similarly, $b|_C = c$, and hence by the Mayer-Vietoris axiom we obtain an element $u \in h(L)$ such that $u|_A = a$ and $u|_B = b$. In particular, $u|_{L_n} = u_n$ for all $n \geq -1$, and consequently $u|_K = v$.

The goal is now to prove that u is universal, which will conclude the proof. Let $n > 0$ be arbitrary, and $0 \leq t < n$. Since $u|_{L_n} = u_n$, the diagram

$$\begin{array}{ccc} \pi_t(L_n) = [S^t, L_n]_* & \xrightarrow{(i_n)_*} & [S^n, L]_* = \pi_n(L) \\ & {\scriptstyle T_{u_n}} \searrow \quad \swarrow {\scriptstyle T_u} & \\ & h(S^n) & \end{array}$$

commutes. Note that Lemma 3.5 constructs L_m from L_n, for $m > n$, by attaching cells of dimension greater than n. It follows that i_n is an n-equivalence, and hence that the horizontal map $(i_n)_*$ is an isomorphism for $t < n$. Since u_n is n-universal, $(T_n)_{S^t}$ is also an isomorphism for $t < n$. Hence, $(T_u)_{S^t}$ is an isomorphism for all $t < n$. Since $n > 0$ was arbitary, so we conclude that u is universal. □

Suppose we start with a based CW complex Y, together with a universal class $u \in h(Y)$. In trying to prove Theorem 2.5, we will want to show that given an arbitrary based CW complex X and an element $w \in h(X)$, there exists a map $G : X \to Y$ such that $G^*u = w$. It will actually be convenient to establish a "relative version" of this construction:

Lemma 3.11. *Let Y be a based CW complex equipped with a universal class $u \in h(Y)$, and let X be a based CW complex, $w \in h(X)$, and $A \subset X$ a based subcomplex. Assume that we have a map $g : A \to Y$ such that $g^*u = w|_A$. Then there exists a map $G : X \to Y$ such that $G|_A = g$ and $G^*u = w$.*

Proof. We begin by considering the space

$$Z = \big((A \wedge I^+) \vee X \vee Y\big) \,/\, \sim,$$

where $(a, 0) \sim a$ and $(a, 1) \sim g(a)$ for any $a \in A \subset X$. Let

$$B = \big((A \wedge [0, 3/4]^+) \vee X\big) \,/\, \sim,$$
$$C = \big((A \wedge [1/4, 1]^+) \vee Y\big) \,/\, \sim,$$

be subcomplexes of X. As we may deformation retract the interval $[0, 3/4]^+$ onto $\{0\}$ (respectively, $[1/4, 1]^+$ onto $\{1\}$), we see that B deformation retracts onto X (respectively, C onto Y). The intersection is given by $B \cap C = (A \wedge [1/4, 3/4]^+)/\sim$, which deformation retracts onto A. As the interiors of B and C cover L, the Mayer-Vietoris axiom yields a class $v \in h(Z)$ such that $v|_X = w$ and $v|_Y = u$.

Now, we invoke Lemma 3.10 to embed Z as a based subcomplex of some CW complex Y' carrying a universal class $u' \in h(Y')$ satisfying $u'|_Z = v$. Although Y' seems quite complicated upon inspection of the proof of Lemma 3.10, the map $j : Y \to Z \hookrightarrow Y'$ is actually a homotopy equivalence. Indeed, as $j^*(u') = u$ by construction, we have a commutative diagram for all $n \geq 0$.

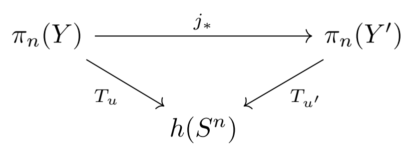

Since the classes u and u' are universal (u by hypothesis, u' by construction), we conclude that the vertical maps are isomorphisms. Hence, j_* is an isomorphism for all $n \geq 0$. Therefore, j is a weak equivalence. Since it is a map of CW complexes, Whitehead's Theorem ensures that j is in fact a homotopy equivalence, as claimed.

Let's now consider the composition $h : A \wedge I^+ \to Z \hookrightarrow Y'$, which is a based homotopy of maps $h_0 : A \to Y'$ and $h_1 : A \to Y'$ (obtained from h by precomposition with the inclusions $A \hookrightarrow A \wedge I^+$ at times 0 and 1, respectively). Since $A \times \{1\}$ is attached to Y in Z via g, we conclude that $h_1 = jg$. Similarly, h_0 is seen to be the composition $A \hookrightarrow X \to Z \hookrightarrow Y'$, which we denote as $g'|_A$ (that is, g' is the composition $X \to Z \hookrightarrow Y'$). So, the existence of h tells us that $g'|_A \simeq jg$.

The map h is a (based) homotopy of maps $A \to Y'$, but we would like to extend it over the entire space X. Fortunately, as A is a subcomplex of X, the inclusion $A \hookrightarrow X$ is a cofibration. We therefore consider the homotopy extension diagram:

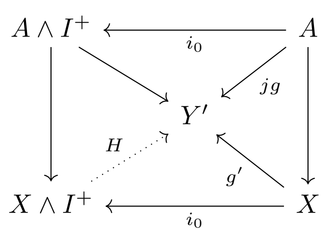

The homotopy extension property gives us the map $H : X \wedge I^+ \to Y'$ above. Inspection of the diagram reveals that $H_0 = g'$, and that $H_1|_A = h_1|_A = jg$. We denote $G' = H_1$. Since the inclusion $A \hookrightarrow X$ is a cofibration, and $j : Y \to Y'$ is a homotopy equivalence, we obtain a map $G : X \to Y$ such that $G|_A = g$ and such that $jG \simeq G'$.

I claim that G is the desired map. Indeed, start with the observation that $G^*u = G^*j^*u' = (jG)^*u' = (G')^*u'$, as $jG \simeq G'$ and $u = j^*u'$ by construction. Next, recall that G' is homotopic to g', via H. Therefore $(G')^*u' = (g')^*u' = (u'|_Z)|_X = v|_X = w$, by construction of u' and v. □

Now we are ready for the proof of Theorem 2.5. The idea is the following: we apply Lemma 3.10 to construct a based CW complex K carrying a universal class $u \in h(K)$, and then leverage Lemma 3.11 to prove that the corresponding natural transformation T_u is in fact an isomorphism.

Proof of Theorem 2.5. Recall that $h(\text{pt}) = \{v\}$ is a one point set. We can apply Lemma 3.10 to pt, $v \in h(\text{pt})$ to obtain a based CW complex K together with a universal class $u \in h(K)$. It now suffices to verify that

$$T_u : [X, K]_* \to h(X)$$

is a bijection for every based CW complex X.

- Surjectivity: Let $w \in h(X)$ be arbitrary. Take $A = \{*\} \subset X$ to be the basepoint, and let $g : A \to K$ be the map to the basepoint of K. Since we automatically have that $w|_A = v = g^*u$, we may invoke Lemma 3.11 to obtain a map $G : X \to K$ such that $G^*u = w$. That is, $T_u(G) = w$, proving surjectivity.
- Injectivity: Suppose $g_0, g_1 : X \to K$ are based maps such that $g_0^*u = g_1^*u$. We set out to prove that $g_0 \simeq g_1$. Let $A = X \wedge \partial I^+$, where $\partial I = \{0, 1, *\}$ is the boundary of I with a disjoint basepoint. Let $g : A \to K$ be the unique map satisfying $g|_{X\times\{0\}} = g_0$ and $g|_{X\times\{1\}} = g_1$. If $p : X \wedge I^+ \to X$ denotes the projection, then we have a commutative diagram.

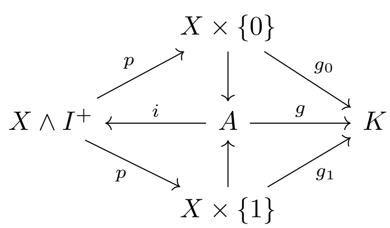

  Now, define $w \in h(X \wedge I^+)$ by $w = p^*g_0^*u$. Inspection of the diagram above yields that $w|_A = g^*u$. Once again applying Lemma 3.11, we construct a map $G : X \wedge I^+ \to K$ such that $G^*u = w$ and, more importantly, $G|_A = g$. We notice that G is a based homotopy of maps $G_0 : X \to K$ and $G_1 : X \to K$. As $G|_A = g$, it follows by construction of g that $G_0 = g_0$ and $G_1 = g_1$. In particular, $g_0 \simeq g_1$ as claimed.

□

4. Connections to stable homotopy theory

Recall the following definition of a spectrum, a central object of study in stable homotopy theory:

Definition 4.1. A *spectrum* E is defined to be a sequence of based spaces $\{E_n\}_{n\geq 0}$, together with *structure maps* $\sigma_n \colon \Sigma E_n \to E_{n+1}$ for each n. Say that a spectrum E is an Ω-*spectrum* if the adjoints of the structure maps $\tilde{\sigma}_n \colon E_n \to \Omega E_{n+1}$ are all weak homotopy equivalences.

Now, suppose we are given a reduced generalized cohomology theory with component functors $h^n : \mathrm{h}\mathbf{CW}_*^{op} \to \mathbf{Set}_*$. By the Brown representability theorem, we can find associated based CW complexes E_n and universal elements $u_n \in h^n(E_n)$ inducing natural isomorphisms for all n:

$$T_{u_n} : [-, E_n] \to h^n(-)$$

Furthermore, via the loop-suspension adjunction and the suspension isomorphism, we obtain a chain of natural isomorphisms:

$$[X, \Omega E_{n+1}] \cong [\Sigma X, E_{n+1}] \cong h^{n+1}(\Sigma X) \cong h^n(X),$$

so that by the Yoneda lemma as in Remarks 2.6 (1) we obtain homotopy equivalences $\tilde{\sigma_n} : E_n \xrightarrow{\simeq} \Omega E_{n+1}$. Thus, the collection $E := \{E_n\}$ together with the adjoints $\sigma_n : \Sigma E_n \to E_{n+1}$ of the maps $\tilde{\sigma}_n$ forms an Ω-spectrum.

Since the functors $h^n := [-, E_n]$ associated to an Ω-spectrum $E = \{E_n\}$ determine a valid reduced generalized cohomology theory, the Brown representability theorem gives us a bijective correspondence between Ω-spectra and reduced generalized cohomology theories given by the assignment:

$$\{E_n\}_{n\geq 0} \mapsto \Big(h^n : \mathrm{Ho}(\mathbf{Top}_*)^{op} \to \mathbf{Ab}, X \mapsto [X, E_n]\Big)_{n\geq 0}.$$

Remark 4.2. (1) It should be noted that, under the proper notion of morphisms of Ω-spectra and cohomology theories, this bijection on objects does not produce an equivalence of categories. This is due to the existence of *hyperphantom maps* between spectra producing the zero map on the associated cohomology theories.
(2) Some of the existing literature requires that a cohomology theory be defined for negative n as well, in which case we may extend the above correspondence via the assignment $X \mapsto [X, \Omega^{-n} E_0]$ for $n < 0$.

Example 4.3. (Eilenberg-Maclane Spectra) Let us apply the Brown representability theorem to reduced singular cohomology with coefficients in a given abelian group G. We get a sequence of spaces $\{K(G, n)\}$, each of which is unique up to homotopy equivalence and fits into a natural isomorphism:

$$\tilde{H}^n(X; G) \cong [X, K(G, n)].$$

In particular, for a given integer n, letting $X = S^q$ yields:

$$\pi_q(K(G, n)) = [S^q, K(G, n)] = \tilde{H}^n(S^q; G) = \begin{cases} G & \text{if } q = n \\ 0 & \text{otherwise.} \end{cases}$$

The spaces $K(G, n)$ are known as the *Eilenberg-Maclane spaces* associated with the group G. By the above discussion, the family of spaces $\{K(G, n)\}_n$ has the structure of an Ω-spectrum, which we call the *Eilenberg-Maclane spectrum* associated to the abelian group G and denote by HG.

5. The ∞-categorical generalization

The proof the Brown representability theorem is, in fact, quite formal. We first observe that Definition 2.1 of a Brown functor $h : \mathrm{h}\mathbf{CW}_*^{\mathrm{op}} \to \mathbf{Set}_*$ can be interpreted categorically. Namely, condition (i) simply states that h sends coproducts in $\mathbf{CW}_*$ to products in $\mathbf{Set}_*$, while (ii) states that given any (homotopy) pushout square

$$\begin{array}{ccc} C & \longrightarrow & C' \\ \downarrow & & \downarrow \\ D & \longrightarrow & D', \end{array}$$

in $\mathbf{CW}_*$, the map $h(D') \to h(D) \times_{h(C)} h(C')$ induced by the universal property of pushouts is surjective. Furthermore, the constructions in the proof generalize quite well. For example, as long as we have a notion of mapping cone and cofiber sequence, the constructions of Lemma 3.1 and Lemma 3.5 will make sense. The mapping telescope used in Lemma 3.10 is nothing but a homotopy colimit in $\mathbf{CW}_*$.

The most important ingredient of the proof, abstractly, was the relationship between the categories $\mathbf{CW}_*$ and $\mathrm{h}\mathbf{CW}_*$. Although our functor h is defined on the latter, many of the arguments of the proof take place in the former; for example, we showed that two maps f, g were equal in $\mathrm{h}\mathbf{CW}_*$ by exhibiting a homotopy between them in $\mathbf{CW}_*$.

This story admits a substantial generalization, established by Lurie in [3], Section 1.4.1. Given any ∞-category $\mathcal{C}$, we can construct its *homotopy category* $h\mathcal{C}$, which is more or less an ordinary category (it is additionally enriched over $\mathrm{Ho}(\mathbf{Top}_*)$). In this section, we explain the generalization of Brown representability to these settings. It would not be possible to give all of the relevant background on ∞-categories here, but the reader should be able to follow the proof without issue after reading Chapter 1 of Lurie [2], and appealing to the nLab as needed[1].

Let's motivate the hypotheses of the theorem. As the classical proof makes frequent use of pushouts, it is not surprising that we'll need to assume that $\mathcal{C}$ is *presentable*, meaning that it is accessible and admits arbitrary small colimits[2] Next, a crucial role in the proof of Theorem 2.5 is played by the spheres $\{S^n\}_{n\geq 0}$. The relevant property of this family is encoded in Whitehead's theorem, which states that a map $f : X \to Y$ of based CW complexes is an equivalence as soon as it induces isomorphisms $\hom_{\mathrm{h}\mathbf{CW}_*}(S^n, X) \to \hom_{\mathrm{h}\mathbf{CW}_*}(S^n, Y)$ for all $n \geq 0$. We'll need to assume that $\mathcal{C}$ contains a collection of objects $\{S_\alpha\}_{\alpha\in I}$ analogous to the collection $\{S^n\}_{n\geq 0}$ in $\mathrm{h}\mathbf{CW}_*$. The first property to formalize is the group structure on $\hom_{\mathrm{h}\mathbf{CW}_*}(S^n, X)$ for each $n \geq 1$, which is induced by the *cogroup* structure on S^n afforded by the fold map $S^n \to S^n \vee S^n$ (collapsing the equator).

Definition 5.1. Let $\mathcal{D}$ be a category admitting finite coproducts. A *cogroup object* of $\mathcal{D}$ is the data of an object $S \in \mathcal{D}$ together with a morphism $S \to S \amalg S$, such that for any $X \in \mathcal{D}$, the induced map

$$\hom_{\mathcal{D}}(S, X) \times \hom_{\mathcal{D}}(S, X) \cong \hom_{\mathcal{D}}(S \amalg S, X) \to \hom_{\mathcal{D}}(S, X)$$

equips $\hom_{\mathcal{D}}(S, X)$ with a group structure.

[1] Alternatively, the reader may keep in mind the slogan that "everything in an ∞-category is done up to coherent homotopy", and believe that the outlined arguments can be formally carried out with this principle in mind.

[2] The notion of accessibility deals with some of the set theoretic difficulties encountered in category theory, and can be safely ignored by the reader not concerned with such things.

Remark 5.2. Let $\emptyset$ denote the initial object in $\mathcal{D}$. Notice that the presence of an identity element in each $\hom_{\mathcal{D}}(S, X)$ is equivalent to the data of a map

$$\mathrm{pt} = \hom_{\mathcal{D}}(\emptyset, X) \to \hom_{\mathcal{D}}(S, X),$$

naturally in X. This is in turn the same data as a map $S \to \emptyset$ in $\mathcal{D}$, which may be thought of as the "counit" of S; the identity in $\hom_{\mathcal{D}}(S, X)$ may then be recovered as the composite $S \to \emptyset \to X$.

We can now make the following (non-standard) definition. Recall that an object S of an ∞-category $\mathcal{C}$ is called *compact* if the functor $\mathrm{Maps}_{\mathcal{C}}(S, -) : \mathcal{C} \to \mathscr{S}$ commutes with all filtered colimits (here and elsewhere, $\mathscr{S}$ denotes the ∞-category of spaces, which may be obtained from the 1-category **CW** by localizing at weak equivalences in the sense of ∞-categories).

Definition 5.3. An object S of a presentable ∞-category $\mathcal{C}$ is called a *generalized sphere* if S is a compact in $\mathcal{C}$ and admits a cogroup structure when viewed in the homotopy category $h\mathcal{C}$.

We begin with the following observation:

Lemma 5.4. *If S is a generalized sphere in $\mathcal{C}$, then so is its suspension ΣS.*

Proof. Recall that the suspension of S is defined as the pushout $\Sigma S := \emptyset \amalg_S \emptyset$, taken with respect to the counit map $S \to \emptyset$ obtained in Remark 5.2. Further notice that the counit map equips any mapping space $\mathrm{Maps}_{\mathcal{C}}(S, X)$ with a basepoint $* = \mathrm{Maps}_{\mathcal{C}}(\emptyset, X) \to \mathrm{Maps}_{\mathcal{C}}(S, X)$, naturally in X, so that we may form the loop space $\Omega\mathrm{Maps}_{\mathcal{C}}(S, X) := * \times_{\mathrm{Maps}_{\mathcal{C}}(S,X)} *$.

To see that ΣS is compact, start with a filtered colimit $\varinjlim_i X_i$ in $\mathcal{C}$. Using that filtered colimits commute with finite limits, we have the following chain of equivalences:

$$\begin{aligned} \mathrm{Maps}_{\mathcal{C}}(\Sigma S, \varinjlim_i X_i) &\simeq \Omega\mathrm{Maps}_{\mathcal{C}}(S, \varinjlim_i X_i) \\ &\simeq \Omega \varinjlim_i \mathrm{Maps}_{\mathcal{C}}(S, X_i) \\ &\simeq \varinjlim_i \mathrm{Maps}_{\mathcal{C}}(\Sigma S, X_i). \end{aligned}$$

It remains to check that ΣS is also a cogroup object. This actually holds whenever S comes equipped with a "counit" map $S \xrightarrow{\epsilon} \emptyset$. Such data enables us to exhibit a comultiplication

$$\Sigma S \to \Sigma S \amalg \Sigma S$$

by identifying the RHS as

$$\Sigma S \amalg \Sigma S = (\emptyset \amalg_S \emptyset) \amalg (\emptyset \amalg_S \emptyset) = \varinjlim \left(\begin{array}{ccccc} \emptyset & \longleftarrow & S & \longrightarrow & \emptyset \\ & & \big\downarrow & & \\ & & \emptyset & & \end{array} \right) = \emptyset \amalg_S \emptyset \amalg_S \emptyset,$$

the LHS as

$$\Sigma S = \emptyset \amalg_S \emptyset = \emptyset \amalg_S S \amalg_S \emptyset,$$

so that we have an evident map

$$\Sigma S = \emptyset \amalg_S S \amalg_S \emptyset \xrightarrow{\mathrm{id}+\epsilon+\mathrm{id}} \emptyset \amalg_S \emptyset \amalg_S \emptyset = \Sigma S \amalg \Sigma S$$

which may be checked to induce a valid group structure on each $\hom_{h\mathcal{C}}(\Sigma S, X)$. $\square$

Call a family of generalized spheres $\{S_\alpha\}_{\alpha\in I}$ in $\mathcal{C}$ a *generating family* if it generates $\mathcal{C}$ under small colimits. Because of Lemma 5.4, we may safely assume that our generating families are closed under taking suspensions. Now, leveraging the fact that ∞-categories are enriched over $\mathscr{S}$, we can formulate a generalization of (and in a sense a corollary to) Whitehead's theorem in our settings:

Lemma 5.5. *Let $\mathcal{C}$ be a presentable ∞-category, and let $\{S_\alpha\}_{\alpha\in I}$ be a generating family of generalized spheres in $\mathcal{C}$. If a morphism $f : X \to Y$ in $\mathcal{C}$ induces isomorphisms $\hom_{h\mathcal{C}}(S_\alpha, X) \to \hom_{h\mathcal{C}}(S_\alpha, Y)$ for all $\alpha \in I$, then f is in fact an equivalence.*

Proof. By a Yoneda argument, showing that the induced map on hom spaces

$$\varphi_Z : \mathrm{Maps}_{\mathcal{C}}(Z, X) \to \mathrm{Maps}_{\mathcal{C}}(Z, Y)$$

is an equivalence for each $Z \in \mathcal{C}$ suffices to establish that f is an equivalence. To this end, let $\mathcal{C}'$ denote the full subcategory of $\mathcal{C}$ consisting of those objects Z for which φ_Z is an equivalence, and let us verify that $\mathcal{C}' = \mathcal{C}$. By the universal property of colimits, we see that $\mathcal{C}'$ is closed under colimits. Since the family $\{S_\alpha\}_{\alpha\in I}$ generates $\mathcal{C}$ under colimits, it therefore suffices to check that φ_{S_α} is an equivalence for all α.

By Whitehead's theorem, the map $\varphi_{S_\alpha} : \mathrm{Maps}_{\mathcal{C}}(S_\alpha, X) \to \mathrm{Maps}_{\mathcal{C}}(S_\alpha, Y)$ is an equivalence iff the induced maps $\pi_n\mathrm{Maps}_{\mathcal{C}}(S_\alpha, X) \to \pi_n\mathrm{Maps}_{\mathcal{C}}(S_\alpha, Y)$ (where the basepoint of either space is obtained from the cogroup structure on S_α) are isomorphisms for every $n \geq 0$. Now, we have that

$$\pi_n\mathrm{Maps}_{\mathcal{C}}(S_\alpha, X) := \pi_0\Omega^n\mathrm{Maps}_{\mathcal{C}}(S_\alpha, X) \simeq \pi_0\mathrm{Maps}_{\mathcal{C}}(\Sigma^n S_\alpha, X) =: \hom_{h\mathcal{C}}(\Sigma^n S_\alpha, X),$$

naturally in X, hence the result follows from the fact that $\{S_\alpha\}_{\alpha\in I}$ was taken to be closed under suspensions, so that by assumption the map

$$\hom_{h\mathcal{C}}(\Sigma^n S_\alpha, X) \xrightarrow{\cong} \hom_{h\mathcal{C}}(\Sigma^n S_\alpha, Y)$$

is an isomorphism. $\square$

We are now in a position to state and prove the following generalization of the Brown representability theorem: parallel to the classical settings, let us call a given $F : h\mathcal{C} \to \mathbf{Set}$ a *Brown functor* if it satisfies the following two properties:

- (i) Given a coproduct $\coprod_i X_i$ in $\mathcal{C}$, the canonical maps $j_i : X_i \hookrightarrow \coprod_i X_i$ induce a bijection

$$F(\amalg_i X_i) \xrightarrow{(F(j_i))_i} \prod_i F(X_i).$$

- (ii) F takes pushouts in $\mathcal{C}$ to weak pullback squares in $\mathbf{Set}$.

Remark 5.6. The relevant colimits are to be taken in the ∞-category $\mathcal{C}$, where they are guaranteed to produce "homotopically correct" answers.

Theorem 5.7. ***(Brown-Lurie Representability Theorem)*** *Let $\mathcal{C}$ be a presentable ∞-category which admits a generating family of generalized spheres $\{S_\alpha\}_{\alpha\in I}$. Then any Brown functor $F : h\mathcal{C}^{op} \to$ **Set** is representable.*

Let us call a pair $(X, \eta \in F(X))$ α-universal if the corresponding map

$$\hom_{\mathrm{h}\mathcal{C}}(S_\alpha, X) \to F(S_\alpha)$$

is bijective, resp. *universal* if it is α-universal for all α in the generating family of generalized spheres. The technical heart of the proof consists of the following analogue of Lemma 3.5 and Lemma 3.10. Before we start, let us make a preliminary comment: the indexing set I of our generating family $\{S_\alpha\}_{\alpha \in I}$ may no longer be countable, and so it is no longer viable to achieve α-universality "one α at a time" as was done in the classical case. Instead, the appropriate strategy will be to take care of "all α's at once". In line with the slogan that "attaching cells may create higher homotopy data", we do so iteratively for $n \geq 0$, imitate the "mapping telescope" argument of Lemma 3.10 (a construction which is particularly natural in the ∞-categorical settings), and conclude by exploiting the compactness of each S_α.

Lemma 5.8. *Given any pair $(X, \eta \in F(X))$ with $X \in \mathcal{C}$, we can find a universal pair $(\tilde{X}, \tilde{\eta} \in F(\tilde{X}))$ together with a map $X \to \tilde{X}$ with respect to which $\tilde{\eta}$ is taken to η.*

Proof. We begin by taking the colimit $X_0 := X \amalg \coprod_{x \in F(S_\alpha)} S_\alpha$, which comes with a canonical map $X \to X_0$. By axiom (i), $F(X_0) \cong F(X) \times \prod_{x \in F(S_\alpha)} F(S_\alpha)$. Let $\eta_0 \in F(X_0)$ denote the class which restricts to η on $F(X)$ and tautologically to each $x \in F(S_\alpha)$. It follows by construction that we have a surjection for each α:

$$\hom_{\mathrm{h}\mathcal{C}}(S_\alpha, X_0) \xrightarrow{(\theta_{\eta_0})_{S_\alpha}} F(S_\alpha). \tag{5.9}$$

Now, we inductively build a sequence of objects of $\mathcal{C}$

$$X_0 \xrightarrow{j_0} X_1 \xrightarrow{j_1} X_2 \xrightarrow{j_2} \ldots \tag{5.10}$$

together with compatible classes $\eta_n \in F(X_n)$ for each $n \geq 1$, in an attempt to "force injectivity" of the analogues of maps 5.9 in the limit. Namely, starting at level n, for each $\alpha \in I$, let $B_\alpha^{(n)}$ denote a system of representatives for elements in $\ker(\hom_{\mathrm{h}\mathcal{C}}(S_\alpha, X_n) \xrightarrow{(\theta_{\eta_n})_{S_\alpha}} F(S_\alpha))$. We then build X_{n+1} as the following (∞-categorical) pushout:

$$\begin{array}{ccc} \coprod_{\alpha, \gamma \in B_\alpha^{(n)}} S_\alpha & \xrightarrow{\amalg \gamma} & X_n \\ \downarrow & & \downarrow j_n \\ \emptyset & \longrightarrow & X_{n+1}. \end{array} \tag{5.11}$$

Since each γ involved corresponds to a class in the kernel of $(\theta_{\eta_n})_{S_\alpha}$, the pair (η_n, v) is a valid element of $F(X_n) \times_{\prod_{\alpha, \gamma \in B_\alpha^{(n)}} F(S_\alpha)} F(\emptyset)$, hence by axiom (ii) we are guaranteed the existence of a class $\eta_{n+1} \in F(X_{n+1})$ which is taken to η_n under the map induced by the canonical morphism $X_n \to X_{n+1}$.

Consider the colimit $\tilde{X} := \varinjlim_n X_n$ with respect to the sequence 5.10 (notice that this ∞-categorical colimit retrieves the mapping telescope from Lemma 3.10 for $\mathcal{C} = \mathscr{S}_*$). We seek to obtain a class $\tilde{\eta} \in F(\tilde{X})$ which restricts to η_n under each of the canonical maps $X_n \to X$. This may be done by considering the following pushout diagram:

$$\begin{array}{ccc} \amalg_n X_n & \longrightarrow & \amalg_n X_{2n} \\ \downarrow & & \downarrow \\ \amalg_n X_{2n+1} & \longrightarrow & \tilde{X} \end{array}$$

where the top horizontal map is given by $\mathrm{id}_{X_0} + j_1 + \mathrm{id}_{X_2} + j_3 + \ldots$, resp. the left vertical map is given by $j_0 + \mathrm{id}_{X_1} + j_2 + \mathrm{id}_{X_3} + \ldots$. By axioms (i) and (ii), we are guaranteed a surjection

$$F(\tilde{X}) \twoheadrightarrow \prod_n F(X_{2n}) \times_{\prod_n F(X_n)} \prod_n F(X_{2n+1})$$

where the pair $((\eta_{2n})_n, (\eta_{2n+1})_n)$ gives a valid element of the RHS by construction, for which we may therefore find a preimage $\tilde{\eta} \in F(\tilde{X})$.

It remains to ensure that the pair $(\tilde{X}, \tilde{\eta} \in F(\tilde{X}))$ is indeed universal. Given $\alpha \in I$, consider the associated map

$$\hom_{\mathrm{h}\mathcal{C}}(S_\alpha, \tilde{X}) \xrightarrow{(\theta_{\tilde{\eta}})_{S_\alpha}} F(S_\alpha).$$

By construction, if we write $l_0 : X_0 \to \tilde{X}$ for the canonical inclusion, we have that $F(l_0)(\tilde{\eta}) = \eta_0$, and it follows that we have a factorization:

$$\begin{array}{ccc} \hom_{\mathrm{h}\mathcal{C}}(S_\alpha, X_0) & \xrightarrow{(l_0)_*} & \hom_{\mathrm{h}\mathcal{C}}(S_\alpha, \tilde{X}) \\ & {\scriptstyle (\theta_{\eta_0})_{S_\alpha}} \searrow \quad \swarrow {\scriptstyle (\theta_{\tilde{\eta}})_{S_\alpha}} & \\ & F(S_\alpha). & \end{array}$$

where the left diagonal map is surjective, hence so is $(\theta_{\tilde{\eta}})_{S_\alpha}$.

Lastly, we check injectivity. Since $(\theta_{\tilde{\eta}})_{S_\alpha}$ is a group homomorphism, it suffices to show that it has trivial kernel, where we recall from Remark 5.2 that the identity element in $\hom_{\mathrm{h}\mathcal{C}}(S_\alpha, \tilde{X})$ is given by the class of the composite $S_\alpha \to \emptyset \to \tilde{X}$. So start with an element of $\ker((\theta_{\tilde{\eta}})_{S_\alpha})$, represented by a map $\gamma : S_\alpha \to \tilde{X} = \varinjlim_n X_n$. Since S_α is compact, γ factors through a map $\gamma_n : S_\alpha \to X_n$ for some large enough n, with $\gamma_n \in B_\alpha^{(n)}$ by definition. By the construction of X_{n+1} from X_n in 5.11, we are therefore guaranteed the existence of a commutative square

$$\begin{array}{ccc} S_\alpha & \xrightarrow{\gamma_n} & X_n \\ \downarrow & & \downarrow{\scriptstyle j_n} \\ \emptyset & \longrightarrow & X_{n+1}, \end{array}$$

witnessing that $j_n\gamma_n$ is homotopic to the composite $S_\alpha \to \emptyset \to X_{n+1}$, hence that the class $[\gamma]$ equals the identity element inside $\hom_{\mathrm{h}\mathcal{C}}(S_\alpha, \tilde{X})$, as needed. □

Now, let us begin with a universal pair $(X, \eta \in F(X))$ obtained by applying Lemma 5.8 to the pair $(\emptyset, v \in F(\emptyset))$ ($F(\emptyset)$ is a one-point set for the same formal reason as in the classical settings). We set out to prove that the induced map

$$\hom_{\mathrm{h}\mathcal{C}}(T, X) \xrightarrow{(\theta_\eta)_T} F(T)$$

is a bijection for any $T \in \mathrm{h}\mathcal{C}$.

- Surjectivity: Begin with an element in the target $t \in F(T)$. Axiom (i) guarantees that we may pick a preimage $\delta \in F(X \amalg T)$ for the pair (η, t) under the canonical map $F(X \amalg T) \to F(X) \times F(T)$. Then, we may further extend the pair $(X \amalg T, \delta \in F(X \amalg T))$ into a universal pair $(Z, u \in F(Z))$, equipped with a map $X \amalg T \to Z$ with respect to which u is taken to δ.

  Now, consider the composite $X \xrightarrow{l_X} X \amalg T \to Z$, with respect to which u is taken to η by construction. We claim that this map is an equivalence, essentially thanks to Lemma 5.5. Indeed, for any α, we may contemplate the following commutative triangle:

$$\begin{array}{ccc} \hom_{h\mathcal{C}}(S_\alpha, X) & \longrightarrow & \hom_{h\mathcal{C}}(S_\alpha, Z) \\ & {\scriptstyle \simeq}\searrow \quad \swarrow{\scriptstyle \simeq} & \\ & F(S_\alpha) & \end{array}$$

  where the diagonal maps are both bijections, hence so is the horizontal map - hence Lemma 5.5 applies to give that the map $X \to Z$ is an equivalence. We may therefore form the "wrong-way" composite $T \to Z \xrightarrow{\simeq} X$ in the following diagram:

$$\begin{array}{ccccc} X & & & & \\ & {\scriptstyle l_X}\searrow & & \overset{\simeq}{\searrow} & \\ & & X \amalg T & \longrightarrow & Z \\ & {\scriptstyle l_T}\nearrow & & \nearrow & \\ T & & & & \end{array}$$

  By construction, the class $\eta \in F(X)$ is sent to $t \in F(T)$ under this composite, establishing surjectivity.
- Injectivity Suppose two maps $f, g : T \to X$ produce the same class $t \in F(T)$ under $(\theta_\eta)_T$. We shall exhibit a homotopy $f \simeq g$, which in our settings is simply a commutative triangle in the ambient ∞-category $\mathcal{C}$. We start by leveraging the Mayer-Vietoris axiom with respect to the pushout square

$$\begin{array}{ccc} T \amalg T & \xrightarrow{f+g} & X \\ {\scriptstyle \nabla}\downarrow & & \downarrow \\ T & \longrightarrow & W \end{array} \tag{5.12}$$

  in order to obtain a preimage $\delta \in F(W)$ for the pair (η, t) under the canonical map $F(W) \to F(X) \times_{F(T) \times F(T)} F(T)$. We may further extend the pair $(W, \delta \in F(W))$ to a universal pair $(\tilde{W}, \tilde{\delta} \in F(\tilde{W}))$, equipped with a map $W \to \tilde{W}$ with respect to which $\tilde{\delta}$ is taken to δ.

  Now, consider the composite $h : X \to W \to \tilde{W}$, where the map $X \to W$ is the vertical map in 5.12. In a similar way to what was done above, we may fall back on Lemma 5.5 to establish that h must in fact be an equivalence. We claim that the composite maps hf and hg are homotopic. Indeed, writing $H : T \to W \to \tilde{W}$ for the composite starting with the horizontal map in 5.12, we see that we have a commutative triangle in $\mathcal{C}$

$$\begin{array}{ccc} T \amalg T & \xrightarrow{hf+hg} & \tilde{W}, \\ {\scriptstyle \nabla}\big\downarrow & \nearrow_{H} & \\ T & & \end{array}$$

corresponding to the data of a homotopy between hf and hg, since the fold map post-composes with either inclusion $T \to T \amalg T$ to produce the identity on T up to homotopy. Since h is an equivalence, it follows that the original maps f and g are also homotopic, completing the proof of Theorem 5.7.

References

[1] Brown, Edgar. Cohomology Theories. The Annals of Mathematics. Vol. 75. 1962. 467-484. (On p. 1)

[2] Lurie, Jacob. Higher Topos Theory. Princeton University Press. 2009. (On pp. 1 and 13)

[3] Lurie, Jacob. Higher Algebra. http://www.math.harvard.edu/~lurie/papers/HA.pdf (On pp. 1 and 12)

[4] Kochman, Stanley O. Bordism, Stable Homotopy and Adams Spectral Sequence. Fields Institute Monographs, American Mathematical Society. 1996. (On p. 1)