

Homotopy Type Theory in Lean

Floris van Doorn¹, Jakob von Raumer², and Ulrik Buchholtz³

¹ Carnegie Mellon University, Pittsburgh, USA
fpvdoorn@gmail.com

² University of Nottingham, United Kingdom
jakob@von-raumer.de

³ TU Darmstadt, Germany
ulrikbuchholtz@gmail.com

Abstract. We discuss the homotopy type theory library in the Lean proof assistant. The library is especially geared toward synthetic homotopy theory. Of particular interest is the use of just a few primitive notions of higher inductive types, namely quotients and truncations, and the use of cubical methods.

Keywords: Homotopy Type Theory · Formalized Mathematics · Lean · Proof Assistants

1 Introduction

Homotopy type theory (HoTT) refers to the homotopical interpretation of Martin-Löf’s dependent type theory [3,22], which grew out of the groupoid model of [13]. In the standard interpretation, every type-theoretical construct corresponds to a homotopy-invariant construction on spaces. An important example is the identity type, which corresponds to the path space construction.

Just like extensional type theory can be interpreted in a variety of categories, for instance elementary toposes, it is expected that homotopy type theory has homotopy-coherent interpretations in higher toposes. Conversely, the interpretation has inspired new type-theoretic ideas such as higher inductive types (HITs) and Voevodsky’s univalence axiom. (See the HoTT book [21] for more about HoTT.)

Most previous formalizations of HoTT used proof assistants that were not originally designed with the homotopy interpretation in mind. In Coq we have both Voevodsky et al.’s *UniMath* project [23] and the HoTT library [4]. In Agda, there is another substantial HoTT library [5]. The former library eschews the use of HITs by instead using Voevodsky’s resizing axiom. Common for all of these libraries is that certain tricks are used to accommodate HoTT: resizing is implemented bluntly in UniMath using the inconsistent principle *type-in-type*, while HITs are implemented in the other libraries using “Licata’s trick” [15]. There is also an impressive experimental proof assistant implementing cubical type theory [7] which is designed with the homotopy interpretation in mind, but it lacks many features that make a proof assistant convenient to use, and the library is so far rudimentary.

Contributions In this paper, we report on a new library⁴ for HoTT in the proof assistant Lean [18]. Lean is open source and implements dependent type theory. It is designed to have a small kernel, with many features built outside the kernel. We describe Lean in greater detail in Sect. 2. The `cloc` tool⁵ reports the library as having 30 400 lines of specification and proof and 3 600 lines of comments. Thus, our library is roughly the same size as the Coq HoTT library, which has 29 800 lines of specification and proof. Our library includes many theorems from synthetic homotopy theory and a large algebraic hierarchy. We describe the library in more detail in Sect. 3. In the library we heavily use cubical techniques for higher path algebra, see Sect. 4. We also have a novel approach to implement HITs, which amounts to having two simple built-in HITs and reducing everything else to those, as described in Sect. 5.

2 The Lean Proof Assistant

Lean [18] is an interactive theorem prover which is mainly developed at Microsoft Research and Carnegie Mellon University. The project was started in 2013 by Leonardo de Moura and has since gained the attention of academics as well as hands-on users. Lean is an open-source program released under the Apache License 2.0 and welcomes additions to its code and mathematical libraries.

In its short history, Lean has undergone several major changes. The second version (Lean 2) supports two kernel modes. The standard mode is for proof irrelevant reasoning, in which `Prop`, the bottom universe, contains types whose objects are considered to be judgmentally equal. Since this is incompatible with homotopy type theory, a second HoTT mode was added, where proof irrelevance is not present. In 2016, the third major version of Lean (Lean 3) was released [17]. In this version, many components of Lean have been rewritten. Of note, the unification procedure has been restricted, since the full higher-order unification which is available in Lean 2 can lead to timeouts and error messages that are unrelated to the actual mistakes. Due to certain design decisions, such as proof erasure in the virtual machine and a function definition package which requires axiom K [11], the homotopy type theory mode is currently not supported in Lean 3. This has led to the situation that the homotopy type theory library is kept in the still maintained but not further developed Lean 2. In the future we hope that we will find a way to support a version of homotopy type theory in Lean 3 or a fork thereof.

The HoTT kernel of Lean 2 provides the following primitive notions:

- *Type universes* `Type.{u} : Type.{u + 1}` for each universe level $u \in \mathbb{N}$. In Lean, this chain of universes is non-cumulative, and all universes are predicative.
- *Function types* `A → B : Type.{max u v}` for types `A : Type.{u}` and `B : Type.{v}` as well as *dependent function types* `Π a, B a : Type.{max u`

⁴ Available as part of: <https://github.com/leanprover/lean2>

⁵ <https://github.com/AlDanial/cloc>

- $v\}$ for each type $A : \mathbf{Type}. \{u\}$ and type family $B : A \rightarrow \mathbf{Type}. \{v\}$. These come with the usual β and η rules.
- *inductive types* and *inductive type families*, as proposed by Peter Dybjer [10]. Every inductive definition adds its constructors and dependent recursors to the environment. Pattern matching is *not* part of the kernel
- two kinds of *higher inductive types*: n -truncation and (typal) quotients (cf. Sect. 5).

Outside the kernel, Lean’s elaborator uses backtracking search to infer implicit information. It does the following simultaneously.

- The elaborator fills in *implicit arguments*, which can be inferred from the context, such as the type of the term to be constructed and the given explicit arguments. Users mark implicit arguments with curly braces. For example, the type of equality is $\mathbf{eq} : \Pi\{A : \mathbf{Type}\}, A \rightarrow A \rightarrow \mathbf{Type}$, which allows the user to write $\mathbf{eq} \ a_1 \ a_2$ or $a_1 = a_2$ instead of $\mathbf{@eq} \ A \ a_1 \ a_2$. The symbol $\mathbf{@}$ allows the user to fill in implicit arguments explicitly. The elaborator supports both first-order unification and higher-order unification.
- We can mark functions as *coercions*, which are then “silently” applied when needed. For example, we have equivalences $f : A \simeq B$, which is a structure consisting of a function $A \rightarrow B$ with a proof that the function is an equivalence. The map $(A \simeq B) \rightarrow (A \rightarrow B)$ is marked as a coercion. This means that we can write $f \ a$ for $f : A \simeq B$ and $a : A$, and the coercion is inserted automatically.
- Lean was designed with *type classes* in mind, which can provide canonical inhabitants of certain types. This is especially useful for algebraic structures (see Sect. 3.3) and for type properties like truncatedness and connectedness. Type class instances can refer to other type classes, so that we can chain them together. This makes it possible for Lean to automatically infer why types are n -truncated if our reasoning requires this, for example when we are eliminating out of a truncated type. For example we show that the type of functors between categories C and D is equivalent to an iterated sigma type.

$$\begin{aligned}
& (\Sigma (F_0 : C \rightarrow D) (F_1 : \Pi \{a \ b\}, \mathbf{hom} \ a \ b \rightarrow \mathbf{hom} (F_0 \ a) (F_0 \ b)), \\
& (\Pi (a), F_1 (ID \ a) = ID (F_0 \ a)) \times \\
& (\Pi \{a \ b \ c\} (g : \mathbf{hom} \ b \ c) (f : \mathbf{hom} \ a \ b), \\
& \quad F_1 (g \circ f) = F_1 \ g \circ F_1 \ f)) \simeq \mathbf{functor} \ C \ D
\end{aligned}$$

Note the use of coercions here: $F_0 : C \rightarrow D$ really means a function from the objects of C to the objects of D . From this equivalence, Lean’s type class inference can automatically infer that $\mathbf{functor} \ C \ D$ is a set if the objects of D form a set. Type class inference will repeatedly apply the rules when sigma-types and pi-types are sets, and use the facts that hom-sets are sets and that equalities in sets are sets (in total 20 rules are applied for this example).

- Instead of giving constructions by explicit terms, we can also make use of Lean’s *tactics*, which give us an alternative way to construct terms step by

step. This is especially useful if the proof term is large, or if the elaboration relies heavily on higher-order unification.

- We can define custom syntax, including syntax with binding. In the following example we declare two custom notations.

```
infix · := concat
notation 'Σ' binders ` , ` r:(scoped P, sigma P) := r
```

The first line allows us to write $p \cdot q$ for path concatenation `concat p q`. The second line allows us to write $\Sigma x, P x$ instead of `sigma P`. This notation can also be chained: $\Sigma (A : \text{Type}) (a : A), a = a$ means `sigma (λ(A : Type), sigma (λ(a : A), a = a))`.

2.1 Consistency of HoTT Lean

Voevodsky’s model of univalence in simplicial sets [14] covers the type theory with empty, unit, disjoint sums, π , σ , identity, and W -types and one univalent universe à la Tarski closed under these type formers. The model validates the β and η rules for function types.

The cubical type theory of [8] interprets Martin-Löf type theory using Andrew Swan’s construction of the identity types. (The cubical path types of this model do not satisfy the computation rule for identity types.) It has been checked that the corresponding model in cubical sets based on de Morgan algebras models two HITs, namely suspension and propositional truncation. (The model even satisfies the computation rules for the path constructors.) The technique used also covers pushouts, so by the reduction of n -truncation to pushouts [19], the model covers all n -truncations. We believe this model also covers all the ordinary inductive families supported by Lean, but this has not been checked in detail.

Mark Bickford’s formalization of the cubical model⁶ covers a whole hierarchy of universes like we have in the Lean kernel. It additionally verifies some novel type constructors such as a higher dimensional intersection type.

These models provide us with high confidence that the logic implemented by the Lean HoTT kernel is consistent. Furthermore, the kernel is very small compared to other kernels implementing dependent type theory. The kernel does not contain pattern matching, a termination checker, fixpoint operations or module management. This increases the confidence that the kernel implements the logic correctly. Furthermore, the only thing we do outside the kernel to extend the logic is to posit the univalence axiom; we do not use type-in-type or Licata’s trick or anything else which might introduce inconsistencies.

3 The Structure of the Library

In this section we describe the overall structure of the homotopy type theory library and we highlight some examples.

⁶ <http://www.nuprl.org/wip/Mathematics/cubical!type!theory/>

The library contains a markdown file in each folder to describe the contents of the files in that folder. For readers familiar with [21], the library includes a file⁷ `book.md` that describes where in the library the various parts of the book are formalized.

Figures 1-4 contain graphs of the files in various parts of the library; the edges denote the dependencies of the files. Each folder contains a file `default` which only contains imports of various files in the folder and which is imported if the user imports the folder. There are also three additional folders: `types` (see Subsect. 3.2), `cubical`, related to the cubical methods discussed in Sect. 4; and `hit`, related to higher inductive types as discussed in Sect. 5. There are also some files in the root folder which we do not describe here.

There is a separate `Spectral` repository,⁸ the goal of which is to formalize the Serre Spectral Sequence, and which will be merged into the Lean-HoTT library in the future. Some examples below are located in this repository.

3.1 The initial part of the library

Figure 1 illustrates the files of the initial part of the library. These files are imported by default when opening a Lean file. The very first file, `datatypes`, defines the basic datatypes, such as `unit`, `empty`, `eq`, `prod`, `sum`, `sigma`, `bool`, `nat`. Higher up, the `path` file develops the basic properties of the identity type (also called equality or identification type) in HoTT. This includes the basic properties of homotopies, transport and the low-dimensional ∞ -groupoid structure of types.

In the rest of the files we define equivalences, posit the univalence axiom and derive function extensionality from univalence (in `equiv`, `ua` and `funext`, respectively). However, in order to be able to track which definitions only depend on function extensionality and not univalence, via the `print axioms` command, we also add function extensionality directly as an axiom.

Lastly, we develop `n`-truncated types, initialize the primitive HITs, prove that types with decidable equality are sets [12] and define the basic notions of pointed types (in `trunc`, `hit`, `hedberg` and `pointed`, respectively).

3.2 Facts about Types

The files in subdirectory `types` develop in more detail the properties and constructions related to individual types and type formers. For types like `sum`, `sigma` and `pi` we characterize the equality in that type, define the functorial action and show that the functorial action preserves equivalences. In `univ` we prove properties of type universes, such as the object classifier property. Of particular importance is the file `pointed`, which contains properties of pointed types, maps, equivalences and homotopies, which contains over 2 000 lines (also counting the corresponding file in the Spectral repository).

⁷ <https://github.com/leanprover/lean2/blob/master/hott/book.md>

⁸ It has 7700 lines of code and 1400 lines of comments. It is available at <https://github.com/cmu-phil/Spectral>

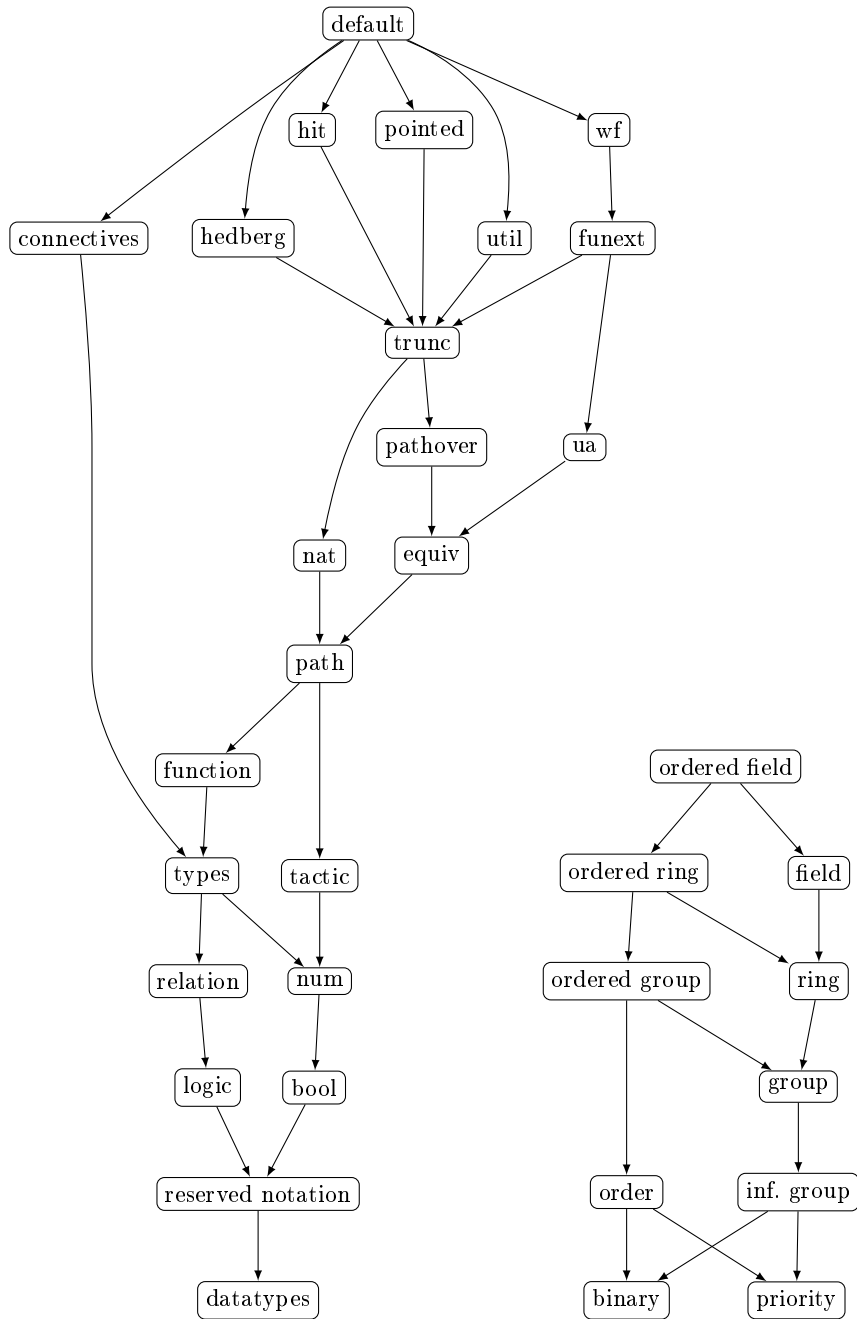


Fig. 1. The initial part of the library

Fig. 2. The algebraic hierarchy

3.3 The Algebraic Hierarchy

The algebraic hierarchy, all in the `algebra` subdirectory, is structured as seen in figure 2. That figure does not contain files that depend on the category theory sublibrary. The algebraic hierarchy defines common algebraic structures, starting with small structures, like semigroups and partial orders, and extending them to groups, rings, all the way up to discrete linear ordered fields. (Discrete means that the order is decidable.)

We combine the “partially bundled” approach with the “fully bundled” approach in the algebraic hierarchy, similar to how algebraic structures in the Coq library are defined [20]. The partially bundled approach means that given a type `A` we define what it means that `A` has a group-structure or ring-structure. This is used for concrete structures, and we use type classes to infer these inhabitants. For example, we prove that \mathbb{N} forms a decidable linear ordered semiring, and mark this as an instance. If we want to show that for `n m k : ℕ` we have $(n * m) * k = n * (m * k)$, we can use `mul.assoc`, the theorem that multiplication in any semigroup is associative. Then type class inference will try to show that \mathbb{N} is a semigroup, and it will use the instance that every decidable linear ordered semiring is a semigroup. We use Lean’s `extend` syntax to easily define new algebraic structures. For example, the following code defines a structure `ab_group` of abelian groups, which consists of the fields of both `group` and `comm_monoid`. Also, the instances `ab_group A → group A` and `ab_group A → comm_monoid A` are automatically generated.

```
structure ab_group [class] (A : Type)
  extends group A, comm_monoid A
```

We use the fully bundled approach when doing group theory and other algebra. A bundled structure is a type together with a structure on that type. For example, this is the definition of a bundled group:

```
structure Group := (carrier : Type) (struct : group carrier)
```

We define `Group.carrier` to be a coercion. We make `Group.struct` an instance, which means that if we have to synthesize a term of type `group (Group.carrier G)`, Lean will automatically find this instance. We use the bundled structures for group theory. For example, if `G H : Group` then we define the product `group G ×g H`. We use `×g` for the product of two groups to disambiguate it from other products, like the product of two types, two pointed types or two truncated types (type class inference does not work well to disambiguate here, since all these structures coerce to types).

If we go back to the example $(n * m) * k = n * (m * k)$ on \mathbb{N} , we also interpret the multiplication symbol on \mathbb{N} using type class inference. In this case, Lean will try to find an instance of `has_mul ℕ`, where `has_mul` is a type class stating that the type has a multiplication. Lean can find this instance since we have a general instance `semigroup A → has_mul A`. However, since we want to also have additive semigroups, we have a different notion of additive semigroups, `add_semigroup`, with corresponding instance `add_semigroup A → has_add A`.

To minimize overhead, we can define additive structures as the multiplicative counterpart, and then prove theorems about additive structures by using the corresponding theorem for multiplicative structures. We do have to manually define the instances for additive structures. Here is an example for semigroups:

```

definition add_semigroup [class] : Type → Type := semigroup
definition has_add_of_add_semigroup [instance] (A : Type)
  [s : add_semigroup A] : has_add A :=
  has_add.mk (@semigroup.mul A s)
definition add.assoc {A : Type} [s : add_semigroup A] (a b c :
  A)
  : (a + b) + c = a + (b + c) :=
  @mul.assoc A s a b c

```

This approach has advantages and disadvantages. An advantage is that theorem names are different for additive structures and multiplicative structures, so we can write `add.assoc` for associativity of addition and `mul.assoc` for associativity of multiplication. Furthermore, we can easily define a ring by extending an additive abelian group and a multiplicative monoid (plus distributivity).

A disadvantage is that operations that are traditionally not written using `+` or `*`, such as concatenating two lists, do not fall in either category. Also, in our formalization we make a distinction between additive and multiplicative groups. Since we define additive groups as multiplicative groups, we can still apply theorems about multiplicative groups to additive groups, but some care is needed when doing this: if one applies a theorem about multiplicative groups with assumption $n * k = 1$ to an additive group, the new subgoal becomes $n * k = 1$, even though in an additive group this really means $n + k = 0$.

All the algebraic structures we mentioned so far (not including `has_mul` and `has_add`) are assumed to be sets, i.e., 0-truncated. We also have variants of some of these structures which are not assumed to be sets. For example, we have `inf_group` and `inf_ab_group`, which are (abelian) groups without the assumption that they are sets, but without higher coherences. This is useful for, e.g., loop spaces or pointed maps into loop spaces, since those types are not groups, but will become groups (the homotopy and cohomology groups) after applying set-truncation.

3.4 Homotopy Theory

The homotopy theory part of the library is organized as shown in figure 3. Almost all results in Chapter 8 of the HoTT book have been formalized in Lean. In particular it contains various results about connectedness, a version of the Freudenthal suspension theorem, the complex and quaternionic Hopf fibration [6] and the long exact sequence of homotopy groups. Together these results show:

```

definition πnSn (n : ℕ) : πg[n+1] (S* (n+1)) ≃g gℤ
definition π3S2 : πg[3] (S* 2) ≃g gℤ

```


This is to say that the n -th homotopy group of the n -sphere (for $n \geq 1$) and the 3rd homotopy group of the 2-sphere are group isomorphic to the integers. Of note here is the notation $\pi g[n] A$ which denotes the n -th homotopy group of A , as a group. In contrast, we also have the operation $\pi[n] A$ which is the n -th homotopy group of A as a pointed type, which is also defined for $n = 0$. Originally, we defined `ghomotopy_group` : $\mathbb{N} \rightarrow \text{Type} \rightarrow \text{Group}$ where `ghomotopy_group n A` is the $(n+1)$ -st homotopy group of A and we had notation $\pi g[n+1] A$ for this. However, this requires the user to write the third homotopy group as $\pi g[2+1]$. To remedy this, we changed the definition of `ghomotopy_group` to have type $\Pi(n : \mathbb{N}) [H : \text{is_succ } n], \text{Type} \rightarrow \text{Group}$, where H is a proof that n is a successor of a natural number, and which is synthesized using type class inference.

We also prove Whitehead’s principle for truncated types and the Seifert-van Kampen theorem, and we define the Eilenberg-MacLane spaces and show that they are unique. Furthermore, we define operations on types of homotopy theoretic significance, such as cofibers, joins, and wedge and smash products, and prove various properties about them, such as the associativity of the join and smash products and the fact that the suspension and smash product have right adjoints, respectively loop spaces and pointed maps.

3.5 Category Theory

It seems a constant across many libraries of formalized mathematics that the development of category theory takes up a substantial fraction of the files, and our library is the same way, as can be seen in figure 4. Highlights include the Yoneda lemma and the Rezk completion. [1]

As an example from this part of the library, consider this excerpt which formalizes the fact that the Yoneda embedding preserves existing limits:

```

definition yoneda_embedding (C : Precategory) : C  $\Rightarrow$  cset  $\hat{\sim}$  c
  Cop

variables {C D : Precategory}
definition preserves_existing_limits [class] (G : C  $\Rightarrow$  D) :=
   $\Pi$ (I : Precategory) (F : I  $\Rightarrow$  C)
    [H : has_terminal_object (cone F)],
    is_terminal (cone_obj_compose G (terminal_object (cone F)))

theorem preserves_existing_limits_yoneda_embedding
  (C : Precategory)
  : preserves_existing_limits (yoneda_embedding C)

```

4 Path Algebra and Cubical Methods

The core innovation in homotopy type theory is its new interpretation of equality. In contrast to proof irrelevant Martin-Löf type theory, we need to be careful

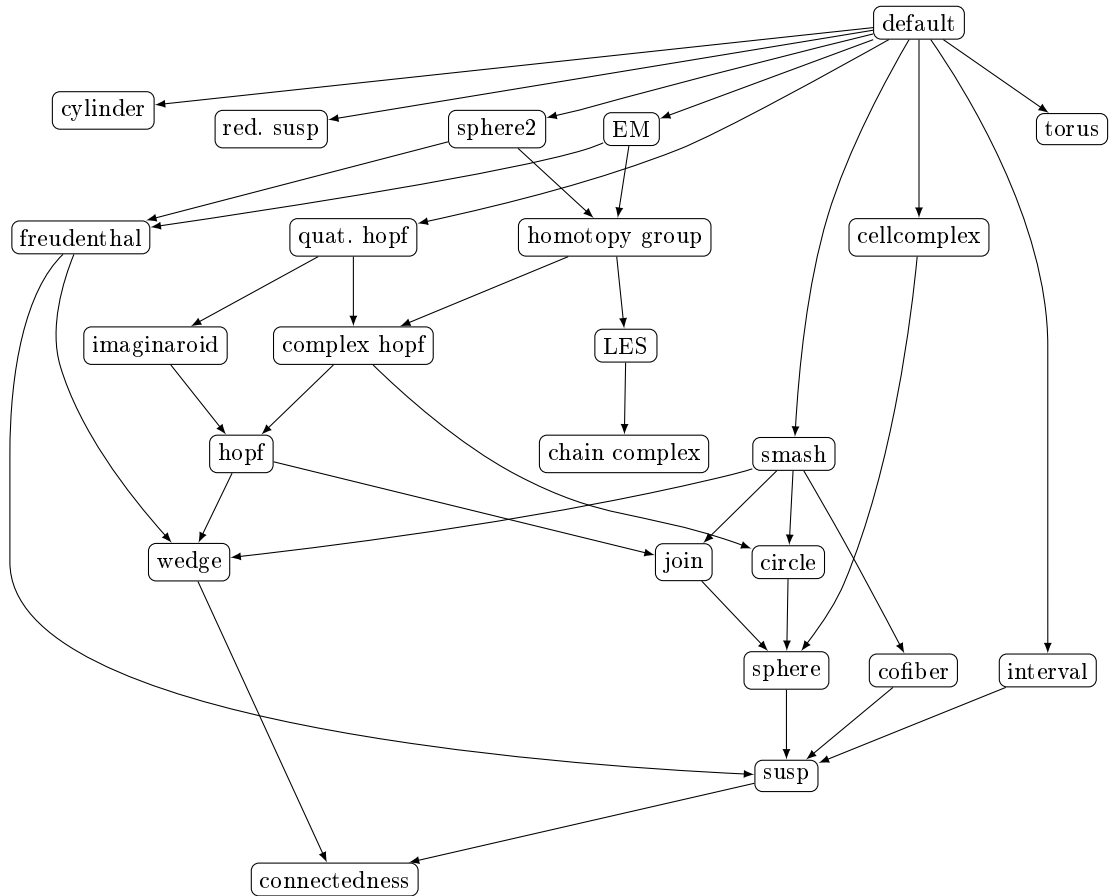


Fig. 3. The homotopy theory part of the library




Fig. 4. The category theory part of the library (c = constructions, f = functor, l = limits)

about choosing well-behaved equality proofs in the library since we might need to prove lemmas about these proof objects themselves. We want to maintain brevity using tactics and equational rewriting while making sure that the generated proofs do not become unwieldy.

After defining equality on a type A in the library’s prelude as an inductive type family over two objects of A which is generated by the reflexivity witness $\text{refl} : \prod(x : A), x = x$, we can provide operations and proofs for the basic *higher groupoid structure* of these “equality paths”: Concatenation $p \cdot q$ and inversion p^{-1} of paths as well as proofs about associativity and cancellation. These are constructed using the dependent recursor of equality which we call *path induction* and which, for each $a : A$, provides a function $\prod(b : A) (p : a = b), P b p$ given the reflexivity case $P a (\text{refl } a)$. Likewise, we can prove the functoriality of functions with respect to equality: For a function $f : A \rightarrow B$ and $p : a = a'$ we define $\text{ap } f p : f a = f a'$ by induction on p . Using an equality $p : a = a'$ in a type A to compare elements of two fibers in a type family C over A , we define the *transport* of an element $x : C a$ along p as $p \triangleright x : C a'$.

For higher paths and dependent paths, we follow what Dan Licata calls the “cubical approach” [16]. The basic notion is that of *pathover*, or a “path over a path”, which compares elements $x : C a$ and $y : C a'$ in different fibers of a type family over some path $p : a = a'$ in the base type. We define the type of pathovers above a base point $a : A$ and $x : C a$ to be the type family $\text{pathover } C x : \prod\{a' : A\}, a = a' \rightarrow C a' \rightarrow \text{Type}$ which is inductively generated by

```
idpo : x =[refl a] x
```

where $x =[p] x'$ is notation for $\text{pathover } _ x p x'$. This definition allows us to define a version $\text{apd } f p : f a =[p] f a'$ of ap for dependent functions $f : \prod(a : A), C a$. It is also used by Lean to express the dependent eliminators for higher inductive types (c.f. Sect. 5). To work with pathovers we provide a variety of operations and lemmas, analogous to the higher groupoid structure of paths. Pathovers correspond to equalities in a sigma type.

For higher paths in a type, we use squares and squareovers. Just like paths were defined as an inductive type family indexed by their endpoints we define the squares in a certain type A as the type family indexed by four corners and four paths between those corners, which is generated by some identity square with refl on all its sides. Squares arise naturally when you need to prove a pathover in an equality type, which is often required when proving equalities involving higher inductive types.

Squareovers are dependent squares over a square. It takes as arguments a square in the base type and four pathovers over the sides of this square. These correspond to squares in a sigma type. We also have a library of cubes three-dimensional equalities. We could generalize these to cubeovers, though we didn’t need those yet.

5 Higher Inductive Types

One novel idea in homotopy type theory is the introduction of *higher inductive types* or HITs [21, Chapter 6]. Higher inductive types are a generalization of inductive types. With inductive types you can specify which terms or points are freely added to that type. In contrast, when defining a HIT, you can specify not only the points in that type, but also paths and higher paths. For example, the circle S^1 is a HIT with one point constructor and one path constructor:⁹

```
HIT circle : Type :=
| base : S1
| loop : base = base
```

This means that the circle is generated by one point and one path `loop : base = base`. There will be more loops in the circle, such as `refl base` and `loop · loop` and `loop-1`, which are all different. Higher inductive types have elimination principles analogous to those of ordinary inductive types.

The most commonly used proof assistants which have HoTT support (such as Coq and Agda) do not support HITs natively. Just adding HITs as constants is not satisfactory, because then the computation rules are not judgmental equalities. Instead, users of Coq use “Dan Licata’s trick” [15]. The idea is that to define a higher inductive type, one first defines a private inductive type inside a module with only the point constructors, and then adds the path constructors as axioms. One then defines the desired induction principle using the induction principle of the private inductive type and adds the computation rules of this induction principle on paths as additional axioms. Then the user closes the module, and the result is that only the data of the higher inductive type are accessible, while the induction principle of the private inductive type is hidden. This ensures that the computation rules are judgmentally true for point constructors (but not for path constructors), but a disadvantage is that inside the module inconsistent axioms were assumed, and one needs to trust that the code in these modules does not introduce an inconsistency in the system. In Agda the rewriting feature is used so that users can extend the kernel with judgmental rewrite rules, though there are no checks for any rewrite rule declared in this way.

In Lean we follow an approach similar to Agda rewriting feature, by building in judgmental rewrite rules. However, we only extend the kernel with the rewrite rules for two “trusted” higher inductive types, namely the `n`-truncation and the typical quotient (quotient for short). The quotient is parameterized by a type `A` and a family of types `R : A → A → Type`. So “typal” (the adjective of “type”) means that we quotient by a family of types and not a family of mere propositions. The quotient is the following HIT:

```
HIT quotient (A : Type) (R : A → A → Type) : Type :=
| i : A → quotient A R
| e : Π{x y : A}, R x y → i x = i y
```

⁹ Although we use syntax inspired by the Lean syntax for inductive types, this is not valid syntax in Lean.

For the n -truncation and the quotient, we add the type formation rule, point and path constructors, and induction principle as constants/axioms.¹⁰ Then we add the judgmental computation rules for the point constructors to the Lean kernel; the Lean kernel is extensible in such a way that certain new computation rules can be added to it. After that, we add the computation rules on paths as axioms. As remarked in Subsect. 2.1, we know that the resulting type theory is consistent, because n -truncations and typal quotients can be reduced to pushouts, and type theory with univalent universes closed under pushouts is modeled by [8].

Given these two HITs, we define all other HITs in the Lean HoTT library using just these two. Some reductions are simple, for example the homotopy pushout of $f : A \rightarrow B$ and $g : A \rightarrow C$ is the quotient on type $B + C$ with the edges R defined as an inductive family with constructor $\Pi(a : A), R(\text{inl}(f a))(\text{inr}(g a))$. Proving the usual induction principle for the pushout is then trivial. Given the pushout, we have defined the other usual HITs: the suspension, circle, join, smash, wedge, cofiber, mapping cylinder and spheres. In particular, we define `circle` as `sphere 1`, which is `susp (susp empty)`. We can then *prove* the usual induction principle for this type, and it satisfies the computation rules on the point constructors judgmentally

We can also define HITs with 2-path constructors using quotients. This uses a method similar to the hubs-and-spokes method described in [21, Sect. 6.7]. From the elimination principle of the circle it follows that for any path $p : x = x$ in type A we can define a map $f : S^1 \rightarrow A$ with $\text{ap } f \text{ loop} = p$ by circle induction. Then we can prove the equivalence

$$(p = \text{refl } x) \simeq \Sigma(x_0 : A), \Pi(z : S^1), f z = x_0$$

This equivalence informally states that filling in a loop is the same as adding a new point x_0 , the *hub*, and *spokes* $f z = x_0$ for every $z : S^1$, similar to the spokes in a wheel. This means that in a higher inductive type, we can replace a 2-path constructor $p = \text{refl } x$ by a new point constructor $x_0 : A$ and a family of 1-path constructors $\Pi(z : S^1), f z = x_0$.

However, this does not quite define 2-HITs in terms of the quotient, since this family of path constructors refers to other path constructors (via the definition of f), which is not allowed in quotients. For this reason, we construct 2-HITs using two nested quotients. We first define a quotient with only the 1-paths and the hubs, and then use another quotient to add the spokes.

For a formal treatment of this, we need the following inductive family, which are the paths in a graph:

```

inductive path {A} (R : A → A → Type) : A → A → Type :=
| of_rel  : Π{a a' : A}, R a a' → path R a a'
| of_path : Π{a a' : A}, a = a' → path R a a'
| symm   : Π{a a' : A}, path R a a' → path R a' a
| trans  : Π{a a' a''}, path R a a' → path R a' a'' → path R a a''

```

¹⁰ For the n -truncation we treat the fact that the new type is n -truncated as a “path-constructor.” In [21, Sect. 7.3] it is explained that the fact that a type is n -truncated can be reduced to (recursive) path constructors.

A *specification for a (nonrecursive) 2-HIT* consists of a type A and two families $R : A \rightarrow A \rightarrow \mathbf{Type}$ and $Q : \prod\{a \ a' : A\}, \text{path } R \ a \ a' \rightarrow \text{path } R \ a \ a' \rightarrow \mathbf{Type}$. Using this, we define the 2-HIT $\text{two_quotient } A \ R \ Q$ with constructors

```
HIT two_quotient A R Q : Type :=
| i_0 : A → two_quotient A R Q
| i_1 : ∏{a a' : A}, R a a' → i_0 a = i_0 a'
| i_2 : ∏{a a' : A} {r r' : path R a a'}, Q r r' →
      extend i_1 r = extend i_1 r'
```

where $\text{extend } i_1 \ r$ is the action of i_1 on paths in R , e.g. $\text{extend } i_1 \ (\text{trans } r_1 \ r_2) := \text{extend } i_1 \ r_1 \cdot \text{extend } i_1 \ r_2$. We first define a special case where the 2-path constructor has only reflexivities on the right hand side. We call this $\text{simple_two_quotient } A \ R \ Q'$, where Q' has type $\prod(a : A), \text{path } R \ a \ a \rightarrow \mathbf{Type}$ and where

```
i_2' : ∏{a} {r : path R a a'}, Q r → extend i_1 r = refl (i_0 a)
```

As mentioned before, we define $\text{simple_two_quotient } A \ R \ Q$ in two steps. We first define a type X with only the 1-path constructors and the hubs:

```
X := quotient A R + ∑(a : A) (r : path R a a), Q' r
```

We then define $\text{simple_two_quotient } A \ R \ Q' := \text{quotient } X \ R'$ where

```
inductive R' : X → X → Type :=
| mk : ∏{a : A} (r : path R a a) (q : Q' r) (x : S1),
      R' (f q x) (inr (a,q))
```

with $f \ q : S^1 \rightarrow X$ defined by induction so that $\text{ap } (f \ q) \ (\text{loop}) = \text{extend } (\text{inl} \circ e) \ r$ for $q : Q' \ r$.

We now prove the expected (dependent) induction principle, (nondependent) recursion principle, and computation rules for this two-quotient. The only computation rule which we did not manage to prove is the computation rule of the induction principle on 2-paths. However, this rule is not necessary to determine the type up to equivalence.

We then define the general version, $\text{two_quotient } A \ R \ Q$, to be equal to $\text{simple_two_quotient } A \ R \ Q'$ where:

```
inductive Q' : ∏{a : A}, path R a a → Type :=
| q_0 : ∏{a a' : A} {r r' : path R a a'},
      Q r r' → Q' (trans r (symm r'))
```

We then show that $\text{two_quotient } A \ R \ Q$ and $\text{trunc } n \ (\text{two_quotient } A \ R \ Q)$ have the right elimination principles and computation rules. It (perhaps surprisingly) requires quite some work to show that the correct computation rules of the truncated version follow from the untruncated version.

This allows us to define all nonrecursive HITs with point, 1-path and 2-path constructors. For example, we define the torus $T^2 := \text{two_quotient } \text{unit} \ R \ Q$. Here $R \ \star \ \star = \text{bool}$, which gives two path constructors p and q from the

basepoint to itself. Q is generated by the constructor $q_0 : Q$ (`trans [ff] [tt]`) (`trans [ff] [tt]`) where `[b]` is notation for `of_rel b`. This gives a path $p \cdot q = q \cdot p$. We also define the *groupoid quotient*: For a groupoid G we define its quotient as `trunc 1 (two_quotient G (@hom G) Q)` where:

```

inductive Q :=
| q_0 :  $\prod(a\ b\ c)$  (g : hom b c) (f : hom a b),
      Q (g o f) (trans f g)

```

If G is just a group (considered as a groupoid with a single object), then the groupoid quotient of G is exactly the Eilenberg-MacLane space $K\ G\ 1$.

We have also defined the propositional truncation just using quotients in Lean [9]. An extension of this construction to n -truncations has been given on paper [19]. If we formalize this generalization in Lean, it is possible to remove n -truncations as a primitive HIT in Lean.

6 Conclusion

We have described the HoTT library for the Lean proof assistant, which formalizes many results in HoTT, including higher inductive types, synthetic homotopy theory and category theory. It has a large library of pointed types, and uses cubical methods for reasoning about higher paths. In the future, we hope to make a HoTT mode for Lean 3, possibly using a version of cubical type theory [8,2].

Acknowledgments We wish to thank the members of the HoTT group at Carnegie Mellon University for many fruitful discussions and Lean hacking sessions, and in particular Steve Awodey and Jeremy Avigad who have been very supportive of our work. Additionally, we deeply appreciate all the times Leonardo de Moura fixed an issue in the Lean kernel to accommodate our library. Lastly, we want to thank all contributors to the HoTT library and the Spectral repository, most notably Egbert Rijke and Mike Shulman.

The first and second authors gratefully acknowledge the support of the Air Force Office of Scientific Research through MURI grant FA9550-15-1-0053. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the AFOSR.

References

1. Ahrens, B., Kapulkin, K., Shulman, M.: Univalent categories and the Rezk completion. *Mathematical Structures in Computer Science* 25(5), 1010–1039 (2015)
2. Angiuli, C., Harper, R., Wilson, T.: Computational higher-dimensional type theory. In: *POPL '17: Proceedings of the 44th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM (2017)
3. Awodey, S., Warren, M.A.: Homotopy theoretic models of identity types. *Math. Proc. Cambridge Philos. Soc.* 146(1), 45–55 (2009)

4. Bauer, A., Gross, J., LeFanu Lumsdaine, P., Shulman, M., Sozeau, M., Spitters, B.: The HoTT Library: A formalization of homotopy type theory in Coq. ArXiv e-prints (Oct 2016)
5. Brunerie, G., Hou (Favonia), K.B., Cavallo, E., Finster, E., Cockx, J., Sattler, C., Jeris, C., Shulman, M., et al.: Homotopy Type Theory in Agda (2017), <https://github.com/HoTT/HoTT-Agda>, code library
6. Buchholtz, U., Rijke, E.: The Cayley-Dickson construction in Homotopy Type Theory. ArXiv e-prints (Oct 2016)
7. Cohen, C., Coquand, T., Huber, S., Mörtberg, A.: Cubical type theory, <https://github.com/mortberg/cubicaltt>, code library
8. Cohen, C., Coquand, T., Huber, S., Mörtberg, A.: Cubical type theory: a constructive interpretation of the univalence axiom. In: 21st International Conference on Types for Proofs and Programs (TYPES 2015). LIPIcs. Leibniz Int. Proc. Inform., Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2016), to appear
9. van Doorn, F.: Constructing the propositional truncation using non-recursive hits. In: Proceedings of the 5th ACM SIGPLAN Conference on Certified Programs and Proofs. pp. 122–129. ACM (2016)
10. Dybjer, P.: Inductive families. Formal aspects of computing 6(4), 440–465 (1994)
11. Goguen, H., McBride, C., McKinna, J.: Eliminating dependent pattern matching. Algebra, Meaning, and Computation pp. 521–540 (2006)
12. Hedberg, M.: A coherence theorem for Martin-Löf’s type theory. J. Funct. Programming 8(4), 413–436 (1998)
13. Hofmann, M., Streicher, T.: The groupoid interpretation of type theory. In: Twenty-five years of constructive type theory (Venice, 1995), Oxford Logic Guides, vol. 36, pp. 83–111. Oxford Univ. Press, New York (1998)
14. Kapulkin, C., Lumsdaine, P.L.: The simplicial model of univalent foundations (after voevodsky) (2012), preprint
15. Licata, D.: Running circles around (in) your proof assistant; or, quotients that compute. blog post (Apr 2011), <http://homotopytypetheory.org/2011/04/23/running-circles-around-in-your-proof-assistant/>
16. Licata, D., Brunerie, G.: A cubical approach to synthetic homotopy theory. In: Proceedings of the 2015 30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). pp. 92–103. LICS ’15, IEEE Computer Society, Washington, DC, USA (2015)
17. de Moura, L., Ebner, G., Roesch, J., Ullrich, S.: The Lean theorem prover. slides (Jan 2017), https://leanprover.github.io/presentations/20170116_P0PL
18. de Moura, L., Kong, S., Avigad, J., van Doorn, F., von Raumer, J.: The Lean theorem prover (system description). In: International Conference on Automated Deduction. pp. 378–388. Springer (2015)
19. Rijke, E.: The join construction. ArXiv e-prints (Jan 2017)
20. Spitters, B., van der Weegen, E.: Developing the algebraic hierarchy with type classes in coq. In: International Conference on Interactive Theorem Proving. pp. 490–493. Springer (2010)
21. The Univalent Foundations Program: Homotopy Type Theory: Univalent Foundations of Mathematics. <http://homotopytypetheory.org/book>, Institute for Advanced Study (2013)
22. Voevodsky, V.: A very short note on the homotopy λ -calculus (2006), http://www.math.ias.edu/~vladimir/Site3/Univalent_Foundations_files/Hlambda_short_current.pdf, online
23. Voevodsky, V., Mörtberg, A., Ahrens, B., Lelay, C., Pannila, T., Matthes, R.: UniMath: *Univalent Mathematics* (2017), <https://github.com/UniMath>, code library