

Internalizing Representation Independence with Univalence

CARLO ANGIULI, Carnegie Mellon University, USA

EVAN CAVALLO, Carnegie Mellon University, USA

ANDERS MÖRTBERG, Stockholm University, Sweden

MAX ZEUNER, Stockholm University, Sweden

In their usual form, representation independence metatheorems provide an external guarantee that two implementations of an abstract interface are interchangeable when they are related by an operation-preserving correspondence. If our programming language is dependently-typed, however, we would like to appeal to such invariance results within the language itself, in order to obtain correctness theorems for complex implementations by transferring them from simpler, related implementations. Recent work in proof assistants has shown that Voevodsky’s univalence principle allows transferring theorems between isomorphic types, but many instances of representation independence in programming involve non-isomorphic representations.

In this paper, we develop techniques for establishing internal relational representation independence results in dependent type theory, by using higher inductive types to simultaneously quotient two related implementation types by a heterogeneous correspondence between them. The correspondence becomes an isomorphism between the quotiented types, thereby allowing us to obtain an equality of implementations by univalence. We illustrate our techniques by considering applications to matrices, queues, and finite multisets. Our results are all formalized in Cubical Agda, a recent extension of Agda which supports univalence and higher inductive types in a computationally well-behaved way.

Additional Key Words and Phrases: Representation Independence, Univalence, Higher Inductive Types, Proof Assistants, Cubical Type Theory

1 INTRODUCTION

There are countless representation independence results in the theory of programming languages, but common among them is the principle that two implementations of an abstract interface are contextually equivalent provided one can exhibit an operation-preserving correspondence between them [Mitchell 1986]. Early representation independence theorems were developed in the context of System F, as “free theorems” obtained through parametricity metatheorems [Reynolds 1983; Wadler 1989]; researchers have since scaled these results to realistic module systems such as those found in OCaml [Leroy 1995] and Standard ML [Crary 2017].

We consider the problem of representation independence for a programming language with full-spectrum dependent types. In this setting, *metatheoretic* invariance results leave much to be desired: ideally, we would like to appeal to these results *within* the language itself to obtain correctness theorems at deep discounts, by establishing a correspondence between two implementations and transferring correctness theorems from the simpler one to the more complex one.

Proof reuse and transfer. This problem is closely related to one studied extensively in the context of the Coq proof assistant [The Coq Development Team 2020], namely, the ability to transfer proofs between representations of mathematical structures [Cohen et al. 2013; Magaud 2003; Magaud and

Authors’ addresses: Carlo Angiuli, Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, 15213, USA, cangiuli@cs.cmu.edu; Evan Cavallo, Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, 15213, USA, ecavallo@cs.cmu.edu; Anders Mörtberg, Department of Mathematics, Stockholm University, Kräftriket Hus 6, Stockholm, 10691, Sweden, anders.mortberg@math.su.se; Max Zeuner, Department of Mathematics, Stockholm University, Kräftriket Hus 6, Stockholm, 10691, Sweden, zeuner@math.su.se.

2018. 2475-1421/2018/1-ART1 \$15.00

<https://doi.org/>

[Bertot 2002; Tabareau et al. 2018]. A simple motivating example is that of the unary natural numbers, which are well-suited to proofs for their natural induction principle, but too computationally inefficient to dispatch large arithmetic subgoals [Chyzak et al. 2014]; binary numbers and machine integers are much more efficient, but are difficult to reason about.

In the Coq Effective Algebra Library (CoqEAL), Cohen et al. [2013] developed a framework for rewriting proof goals from *proof-oriented* data structures and algorithms to *computation-oriented* structures and algorithms. In their framework, users parameterize their development over the relevant types, prove a refinement relation between the two instantiations, and then apply a free theorem obtained via a syntactic translation of the Calculus of Inductive Constructions (CIC) into itself. The final step uses the fact that the CIC is sufficiently expressive to define its own parametricity model, albeit externally [Anand and Morrisett 2017; Bernardy et al. 2012; Keller and Lassel 2012]. However, CoqEAL’s transfer framework does not support dependently-typed goals.

Recently, Tabareau et al. [2018] combined techniques from CoqEAL with Voevodsky’s univalence axiom, which asserts that isomorphic types are equal (and thence, interchangeable) [Voevodsky 2010a]. Their *univalent parametricity* approach transfers theorems, including dependently-typed ones, between isomorphic types without appealing to a meta-level translation. However, their framework is limited to isomorphisms, which does not cover all cases of mathematical interest: for example, CoqEAL handles sparse representations of matrices, which are often not unique, and are therefore not isomorphic to standard representations.

Isomorphisms are not enough. Restricting our attention to isomorphisms is particularly untenable in the context of programming. Consider a standard `Queue A` interface, consisting of a type `Q`, an `empty` queue, and functions `enqueue` : $A \rightarrow Q \rightarrow Q$ and `dequeue` : $Q \rightarrow \text{Maybe}(Q \times A)$. If we intend to prove our code correct, we may reach for a naive implementation $Q = \text{List } A$, which `enqueues` on the front and `dequeues` from the back (writing `last` for the function sending `[]` to `nothing` and `xs ++ [x]` to `just (xs , x)`):

```
ListQueue : (A : Type) → Queue A
ListQueue A = queue (List A) [] _::_ last
```

If we intend to run our code on serious workloads, however, we may reach instead for batched queues ($Q = \text{List } A \times \text{List } A$), which `enqueue` on the front of the first list and `dequeue` from the front of the second list [Okasaki 1999, Section 5.2]:

```
BatchedQueue : (A : Type) → Queue A
BatchedQueue A =
  queue (List A × List A) ([], [])
  (λ x (xs , ys) → fastcheck (x :: xs , ys))
  (λ { _ , [] } → nothing ; (xs , x :: ys) → just (fastcheck (xs , ys) , x))
where
  fastcheck : {A : Type} → List A × List A → List A × List A
  fastcheck (xs , ys) = if isEmpty ys then ([], reverse xs) else (xs , ys)
```

Whereas `ListQueue A .dequeue` takes linear time, `BatchedQueue A .dequeue` is amortized constant-time (being constant-time except when `ys` is empty, in which case one must `reverse xs`). To understand the batched representation, we can observe that each `BatchedQueue` has the same extensional behavior as exactly one `ListQueue`, computed as follows:

```
appendReverse : {A : Type} → BatchedQueue A .Q → ListQueue A .Q
appendReverse (xs , ys) = xs ++ reverse ys
```

This correspondence is *structure-preserving*—it preserves `empty` and commutes with `enqueue` and `dequeue`—and therefore, by representation independence, `ListQueues` and `BatchedQueues` are contextually equivalent. We are tantalizingly close to obtaining the same result internally by univalence, or rather, by the *structure identity principle* (SIP), a consequence of univalence which asserts that isomorphic *structured* types (types equipped with operations) are equal, provided that the isomorphism is structure-preserving [Univalent Foundations Program 2013, Section 9.8]. The gap is simple: `appendReverse` is not an isomorphism, as it sends both $([], [1, 0])$ and $([0], [1])$ to $[0, 1]$.

We arrive at the core of our problem: representation independence metatheorems are unsatisfactory, because we want to leverage representation independence to prove internal correctness theorems; CoqEAL-style approaches do not allow us to replace `ListQueues` with `BatchedQueues` in dependently-typed goals; and univalence does not apply, because our types are not isomorphic.

Contribution. In this paper, we develop techniques for establishing *internal* relational representation independence results in dependent type theory, without appealing to parametricity translations. Taking inspiration from Homotopy Type Theory/Univalent Foundations (HoTT/UF), we use *higher inductive types* (HITs) to simultaneously quotient two related implementation types by a heterogeneous correspondence between them, improving that correspondence to an isomorphism; we then apply univalence (and in particular, the SIP) to obtain an equality of implementations. In the case of `Queues`, our technique leaves `ListQueues` untouched while equating two `BatchedQueues` if and only if `appendReverse` maps them to the same list.

Our results are all formalized in Cubical Agda [Vezzosi et al. 2019], an extension of Agda [The Agda Development Team 2020] which supports univalence and HITs in a computationally well-behaved way using ideas from Cubical Type Theory [Angiuli et al. 2018; Cohen et al. 2018].

From the perspective of proof reuse, our work bridges the gap between CoqEAL and univalent parametricity, because it applies to both dependently-typed goals and non-isomorphism correspondences. In addition, our techniques require less engineering work than these systems, because we rely on built-in features of Cubical Agda rather than typeclass instances for univalence or a Coq plugin to generate parametricity translations. However, we stress that our theorems hold in any dependent type theory with univalence and HITs, features which can be added axiomatically to Coq [Bauer et al. 2017] and Lean [van Doorn et al. 2017], among other languages.

Outline. In Section 2, we review some features of Cubical Agda that we use in the remainder of the paper, including univalence and HITs. In Section 3, we state and prove a version of the SIP, which states that structure-preserving isomorphisms yield equal structured types; we also develop reflection-based automation for automatically deriving the appropriate notion of structure-preservation for a given structure. In Section 4, we demonstrate how to use the SIP in concert with HITs to obtain representation independence results for matrices and queues. In Section 5, we generalize these techniques by introducing the notion of a *quasi-equivalence relation* (QER), and present the main technical result of the paper: any structured QER can be improved to a structured isomorphism between quotients. To illustrate the technique, we apply our theorem to two implementations of finite multisets. Finally, in Section 6, we conclude with future directions and a discussion of related work.

2 PROGRAMMING IN CUBICAL TYPE THEORY

In this section, we review some basic features of Cubical Agda, a recent extension of Agda. All code in this paper is formalized in Agda 2.6.2, which is in development at the time of writing, and

our main results have been integrated into the agda/cubical library.¹ Experts on Cubical Type Theory or HoTT/UF can safely skim this section; readers who wish to learn more about Cubical Agda are encouraged to consult [Vezzosi et al. \[2019\]](#) for a description of the system, or [Cohen et al. \[2018\]](#) and [Coquand et al. \[2018\]](#) for a detailed exposition of its core type theory.

2.1 Equalities as paths

Most languages with full-spectrum dependent types, including Agda [[The Agda Development Team 2020](#)], Coq [[The Coq Development Team 2020](#)], Idris [[Brady 2013](#)], and Lean [[de Moura et al. 2015](#)], have two notions of equality. *Definitional equality* is a purely syntactic notion implemented in typechecking algorithms, which use computation to silently discharge many trivial obligations. Equations that cannot be discharged in this way are instead mediated by an *equality type* (or *propositional equality*) whose proofs are often provided by the user, and appeals to which are marked explicitly in terms and types.

In those languages, the equality type is defined as an inductive family $\text{Eq } A \ x \ y$ generated by a *reflexivity* constructor $\text{refl } x : \text{Eq } A \ x \ x$, following [Martin-Löf \[1975\]](#). One can easily show that equality types satisfy many mathematical properties of equality (e.g., symmetry, transitivity, congruence) but it famously lacks several others, including *function extensionality*: the property that pointwise-equal functions are equal as functions [[Boulier et al. 2017](#)].

Principles such as function extensionality, univalence [[Voevodsky 2010a](#)], and excluded middle can be added to type theory as axioms. However, axioms in type theory lack computational content, causing proofs using axioms to become “stuck,” and thus less amenable to definitional equality. More importantly, type theories with stuck terms cannot serve as programming languages because they lack *canonicity*, the property that all closed terms of natural number type compute numerals.

Building on ideas from HoTT/UF, Cubical Type Theory is a computationally well-behaved extension to dependent type theory whose formulation of equality types, called *path types*, resolve many longstanding issues with Eq types, including the failure of function extensionality, propositional extensionality ([Section 2.2](#)), and the inability to define quotients ([Section 2.3](#)). In Cubical Type Theory, paths are defined as maps out of an *interval* type \mathbb{I} which contains two elements $i0 : \mathbb{I}$ and $i1 : \mathbb{I}$ that are computationally distinct but logically equivalent, in the sense that no function can map them to unequal objects. Therefore, functions $f : \mathbb{I} \rightarrow A$ serve as evidence that $f(i0)$ and $f(i1)$ are equal in A , just as paths in topological spaces can be represented by continuous functions out of the real unit interval $[0, 1] \subset \mathbb{R}$. Iterated equality proofs in A are thus functions $\mathbb{I}^n \rightarrow A$ which correspond topologically to squares, cubes, etc. in A , leading to the term *cubical*.

In Cubical Agda, the interval \mathbb{I} is a primitive type which, in addition to the elements $i0 : \mathbb{I}$ and $i1 : \mathbb{I}$, is equipped with three operations—*minimum* ($_ \wedge _ : \mathbb{I} \rightarrow \mathbb{I} \rightarrow \mathbb{I}$), *maximum* ($_ \vee _ : \mathbb{I} \rightarrow \mathbb{I} \rightarrow \mathbb{I}$), and *reversal* ($\sim _ : \mathbb{I} \rightarrow \mathbb{I}$)—which satisfy the laws of a *De Morgan algebra*, i.e., a bounded distributive lattice ($i0, i1, _ \wedge _, _ \vee _$) with a De Morgan involution $\sim _$. Other formulations of Cubical Type Theory requiring less structure on \mathbb{I} are also possible [[Angiuli et al. 2019, 2018](#)].

Paths. Path types are a special form of dependent function type $(i : \mathbb{I}) \rightarrow A \ i$ that specify the behavior of their elements on $i0$ and $i1$:

$$\text{PathP} : (A : \mathbb{I} \rightarrow \text{Type } \ell) \rightarrow A \ i0 \rightarrow A \ i1 \rightarrow \text{Type } \ell$$

¹The current development version of Agda is available at <https://github.com/agda/agda>, and the agda/cubical library is available at <https://github.com/agda/cubical>.

Here, $A : \mathbb{I} \rightarrow \text{Type } \ell$ is a function from \mathbb{I} to one of Agda’s universes of types.² Elements of PathP are λ -abstractions $\lambda i \rightarrow t : \text{PathP } A \ t \ [i0 / i] \ t \ [i1 / i]$ where $t : A \ i$ for $i : \mathbb{I}$. We can apply a path $p : \text{PathP } A \ a_0 \ a_1$ to $r : \mathbb{I}$, obtaining $p \ r : A \ r$, and paths satisfy β - and η -rules definitionally, just like ordinary Agda functions. The only difference between $p : \text{PathP } A \ a_0 \ a_1$ and a function $f : (i : \mathbb{I}) \rightarrow A \ i$ is that p is subject to additional definitional equalities $p \ i0 = a_0$ and $p \ i1 = a_1$.

Because the two “endpoints” of a path $a_0 : A \ i0$ and $a_1 : A \ i1$ have different types, PathP types in fact represent *heterogeneous* equalities, or *Paths over Paths* in the terminology of HoTT/UF [Univalent Foundations Program 2013, Sect. 6.2]. We can recover homogeneous (non-dependent) paths in terms of PathP as follows:

```
_≡_ : {A : Type ℓ} → A → A → Type ℓ
_≡_ {A = A} x y = PathP (λ _ → A) x y
```

Here, $\{A = A\}$ tells Agda to bind the implicit argument A (the first A) to a variable (the second A) for use on the right-hand side. Path types allow us to manipulate equality proofs using standard functional programming idioms. For instance, a constant path represents a reflexive equality proof.

```
refl : {x : A} → x ≡ x
refl {x = x} = λ _ → x
```

(Note that `refl` is a function, rather than a constructor.) We can directly apply a function to a path in order to prove that dependent functions respect equality:

```
cong : {B : A → Type} {f : (a : A) → B a} {x y : A} (p : x ≡ y) → PathP (λ i → B (p i)) (f x) (f y)
cong f p i = f (p i)
```

We can in fact define `cong` for $B : A \rightarrow \text{Type } \ell$, but to increase readability in the remainder of the paper, we henceforth suppress most universe levels and arguments such as $\{A : \text{Type}\}$ that are easily inferred by the reader. As paths are just functions it is now trivial to *prove* function extensionality, the property that pointwise equal functions are equal:

```
funExt : {f g : A → B} → ((x : A) → f x ≡ g x) → f ≡ g
funExt p i x = p x i
```

The proofs of function extensionality for dependent and n -ary functions are equally direct. Since `funExt` is *definable* in Cubical Agda, it has computational content: namely, to swap the arguments to p . Functional programmers might recognize this as a special case of the `flip` function.

Transport and composition. One of the key operations of equality in type theory is *transport*, which sends equalities between types to coercions between those types. In Cubical Agda, this principle is a special instance of a primitive called `transp`:

```
transport : A ≡ B → A → B
transport p a = transp (λ i → p i) i0 a
```

Consequences of `transport` include the substitution principle and, using the minimum operation on the interval, the usual induction principle for inductively-defined `Eq` types:

```
subst : (B : A → Type) {x y : A} → x ≡ y → B x → B y
subst B p b = transport (λ i → B (p i)) b

J : {x : A} (P : ∀ y → x ≡ y → Type) (d : P x refl) {y : A} (p : x ≡ y) → P y p
J P d p = transport (λ i → P (p i) (λ j → p (i ∧ j))) d
```

²Agda’s universes are typically called `Set`, but this paper and the `agda/cubical` library use `Type` instead because, as we will discuss in Section 2.2, the term *set* has a technical meaning in HoTT/UF.

Using `J` we can easily reproduce the standard type-theoretic proofs that \equiv is symmetric and transitive. One can also prove these properties directly using Cubical Agda’s primitive *homogeneous composition* operation `hcomp`, which expresses these “groupoid laws” and their higher-dimensional analogues. Although `hcomp` is used extensively in the agda/cubical library, we do not need to use it directly in this paper.

2.2 Univalence

The core innovation of HoTT/UF, and the main motivation behind Cubical Type Theory, is Voevodsky’s *univalence* principle [Univalent Foundations Program 2013; Voevodsky 2010a], which states that any equivalence of types (written \simeq , a higher-dimensional analogue of isomorphism) yields an equality of types, such that *transporting* along the equality applies the equivalence [Licata 2016]:

$$\text{ua} : \{A B : \text{Type}\} \rightarrow A \simeq B \rightarrow A \equiv B$$

$$\text{ua}\beta : \{A B : \text{Type}\} (e : A \simeq B) (x : A) \rightarrow \text{transport} (\text{ua } e) x \equiv (\text{equivFun } e) x$$

Here, `equivFun` : $A \simeq B \rightarrow A \rightarrow B$ returns the function underlying an equivalence. Because `transport` always produces an equivalence, univalence essentially states that $A \simeq B$ is not only a necessary but also a sufficient condition for the existence of an equation $A \equiv B$.

In the past decade, several libraries of formalized mathematics have been developed around an axiomatic formulation of univalence, including the UniMath library in Coq [Voevodsky et al. 2020], and the HoTT libraries in Coq [Bauer et al. 2017], Agda [Brunerie et al. 2018], and Lean [van Doorn et al. 2017]. In contrast, `ua` and `uaβ` are definable in Cubical Agda using the `Glue` types of Cohen et al. [2018], and have computational content. In particular, many instances of `uaβ` hold directly by computation (writing `not≃` for a proof that `not` : `Bool` → `Bool` is an equivalence):

$$\begin{aligned} _ & : \text{transport} (\text{ua } \text{not}\simeq) \text{false} \equiv \text{true} \\ _ & = \text{refl} \end{aligned}$$

Recalling that all constructions in type theory respect the equality type by fiat, the force of univalence is therefore that *all constructions respect equivalence/isomorphism*. For instance, without inspecting the definition of `IsMonoid` : `Type` → `Type`, we can conclude that an equivalence $A \simeq B$ gives rise to a path $A \equiv B$, which in turn gives rise to a path `IsMonoid` $A \equiv$ `IsMonoid` B . In Section 3, we will use the structure identity principle to show that a path between two monoids is exactly a monoid isomorphism in the usual sense.

Sets and propositions. Univalence refutes *uniqueness of identity proofs* (UIP), or Streicher’s axiom K [Streicher 1993], because it produces equality proofs in `Type` that are not equal. For example, we can see via `transport` that `ua not≃` and `refl` are unequal proofs of `Bool` \equiv `Bool`:

$$\begin{aligned} \text{not}\neq\text{refl} & : \neg (\text{ua } \text{not}\simeq \equiv \text{refl}) \\ \text{not}\neq\text{refl } e & = \text{true}\neq\text{false} (\text{transport} (\lambda i \rightarrow \text{transport} (e i) \text{true} \equiv \text{false}) \text{refl}) \end{aligned}$$

In the presence of univalence, it is important to keep track of which types satisfy UIP or related principles expressing the complexity of a type’s equality relation. In the terminology of HoTT/UF, a type satisfying UIP is called an *h-set* (homotopy set, henceforth simply *set*), while a type whose elements are all equal is called an *h-proposition* (henceforth *proposition*):

$$\begin{aligned} \text{isProp} & : \text{Type} \rightarrow \text{Type} \\ \text{isProp } A & = (x y : A) \rightarrow x \equiv y \\ \text{isSet} & : \text{Type} \rightarrow \text{Type} \\ \text{isSet } A & = (x y : A) \rightarrow \text{isProp} (x \equiv y) \end{aligned}$$

The empty and unit types are propositions, and propositions are closed under many type formers. For instance, we can show that a dependent function type valued in propositions is a proposition, using a variation of function extensionality:

$$\begin{aligned} \text{isProp}\Pi &: \{B : A \rightarrow \text{Type}\} (h : (x : A) \rightarrow \text{isProp } (B x)) \rightarrow \text{isProp } ((x : A) \rightarrow B x) \\ \text{isProp}\Pi & h p q i x = h x (p x) (q x) i \end{aligned}$$

Although all proofs of a proposition are equal, propositions cannot in general be erased—for example, we will soon see that the property of being an equivalence is a proposition, but it implies the existence of an inverse function, which has computational content. We will consider a related programming application in [Example 2.1](#).

A set is a type whose equality types are propositions, and hence satisfies UIP. While Cubical Type Theory has many non-sets, including `Type` itself, most of the types used in this paper will be sets. Hedberg’s theorem states that all types with decidable equality are sets [[Hedberg 1998](#)], which includes many datatypes such as `Bool`, `N`, and `Z`. In addition, any proposition is a set, from which it follows that `isProp A` and `isSet A` are propositions.

Finally, a type is *contractible* if it has exactly one element:

$$\begin{aligned} \text{isContr} &: \text{Type} \rightarrow \text{Type} \\ \text{isContr } A &= \Sigma [x \in A] ((y : A) \rightarrow x \equiv y) \end{aligned}$$

(Here, we are using the `agda/cubical` library’s notation for Agda’s built-in Σ -types, which are `records` with constructor `_,_`, projections `fst` and `snd`, and definitional η .) We can characterize propositions as types whose equality types are contractible, just as sets are types whose equality types are propositions. Thus contractible types, propositions, and sets serve as the bottom three layers of an infinite hierarchy of types introduced by Voevodsky, known as *h-levels* [[Voevodsky 2010b](#)] or *n-types* [[Univalent Foundations Program 2013](#)].

Equivalences and isomorphisms. An equivalence $A \simeq B$ is a function $A \rightarrow B$ such that the preimage of every point in B is contractible:

$$\begin{aligned} \text{preim} &: \{A B : \text{Type}\} (f : A \rightarrow B) (y : B) \rightarrow \text{Type} \\ \text{preim } \{A = A\} f y &= \Sigma [x \in A] (f x \equiv y) \end{aligned}$$

$$\begin{aligned} \text{isEquiv} &: \{A B : \text{Type}\} (f : A \rightarrow B) \rightarrow \text{Type} \\ \text{isEquiv } \{B = B\} f &= (y : B) \rightarrow \text{isContr } (\text{preim } f y) \end{aligned}$$

$$\begin{aligned} _ \simeq _ &: (A B : \text{Type}) \rightarrow \text{Type} \\ A \simeq B &= \Sigma [f \in (A \rightarrow B)] (\text{isEquiv } f) \end{aligned}$$

(In the HoTT/UF literature, `preim` is often called the *homotopy fiber*.) Every equivalence has an inverse which can be constructed by sending each $y : B$ to the element of its preimage determined by `isContr`. In practice, we often construct equivalences $A \simeq B$ by first constructing ordinary *isomorphisms* `Iso A B` (quadruples of a function $A \rightarrow B$, its inverse $B \rightarrow A$, and proofs that these functions cancel in each direction), then applying the lemma `isoToEquiv` : `Iso A B` \rightarrow `A` \simeq `B`.

One might wonder, why introduce equivalences at all? The reason is that `isEquiv f` is always a proposition (by `isProp\Pi` and the fact that `isContr` is a proposition), and thus two elements of $A \simeq B$ are equal if and only if their underlying functions are equal (because their `isEquiv` proofs must agree). In contrast, the property of a function $f : A \rightarrow B$ being an isomorphism (i.e., the triple of a function $B \rightarrow A$ and proofs that it is left and right inverse to f) is *not* a proposition [[Univalent Foundations Program 2013](#), Theorem 4.1.3]. A subtle consequence of this fact is that if one replaces \simeq in the statement of univalence with `Iso`, the resulting statement is actually *inconsistent* [[Univalent](#)

Foundations Program 2013, Exercise 4.6]! However, for the purposes of this paper, the reader can safely imagine equivalences as isomorphisms.

In dependent type theory, equality in Σ -types is notoriously difficult to manage, due to the equality of second projections being heterogeneous (usually stated with transports). In Cubical Type Theory, however, the natural heterogeneity of `PathP` types allows us to straightforwardly characterize equality in Σ -types without any transports:

$$\begin{aligned} \Sigma\text{Path}\simeq\text{Path}\Sigma &: \{B : A \rightarrow \text{Type}\} (x\ y : \Sigma A\ B) \\ &\rightarrow (\Sigma [p \in (\text{fst } x \equiv \text{fst } y)] (\text{PathP } (\lambda i \rightarrow B (p\ i)) (\text{snd } x) (\text{snd } y))) \simeq (x \equiv y) \\ \Sigma\text{Path}\simeq\text{Path}\Sigma\ x\ y &= \text{isoToEquiv } (\text{iso } (\lambda \{ (p, q)\ i \rightarrow (p\ i, q\ i)\}) \\ &\quad (\lambda p \rightarrow ((\lambda i \rightarrow \text{fst } (p\ i)), (\lambda i \rightarrow \text{snd } (p\ i)))) \\ &\quad (\lambda _ \rightarrow \text{refl}) (\lambda _ \rightarrow \text{refl})) \end{aligned}$$

(Here, `iso` is the constructor for the `Iso` type.) Therefore, we can always exchange an equality proof in a Σ -type with a pair of equalities, which tends to be much easier than manipulating transports. Finally, by `isoToEquiv` it is easy to show that logically equivalent propositions are equivalent and thus equal by `ua`. This principle, *propositional extensionality*, is another often-assumed axiom in type theory that is provable (and hence has computational content) in Cubical Agda.

$$\begin{aligned} \text{prop}\simeq &: \text{isProp } A \rightarrow \text{isProp } B \rightarrow (A \rightarrow B) \rightarrow (B \rightarrow A) \rightarrow A \simeq B \\ \text{prop}\simeq\ pA\ pB\ fg &= \text{isoToEquiv } (\text{iso } fg (\lambda _ \rightarrow pB\ _) (\lambda _ \rightarrow pA\ _)) \end{aligned}$$

2.3 Higher inductive types

Finally, Cubical Agda natively supports *higher inductive types* (HITs), a generalization of inductive datatypes which allows for constructors of equality type [Cavallo and Harper 2019; Coquand et al. 2018; Lumsdaine and Shulman 2019; Univalent Foundations Program 2013]. In HoTT/UF, HITs provide analogues of topological spaces such as the circle, spheres, and torus; in this paper, we use HITs to take quotients of types by equivalence relations.

Propositional truncation. Our first example of a HIT is *propositional truncation*, which quotients a type by the total relation, yielding a proposition. It has two constructors: `|_|`, which includes elements of `A` as elements of `|| A ||`, and `squash`, which equates any two elements of `|| A ||`.

```
data ||_| (A : Type) : Type where
  |_| : A → || A ||
  squash : (x y : || A ||) → x ≡ y
```

We write functions out of `|| A ||` by pattern-matching, noting that `squash x y i` constructs an element of `|| A ||`. For example, we can define the functorial action of `|| - ||` on a function as follows:

```
map : (A → B) → || A || → || B ||
map f | x | = | f x |
map f (squash x y i) = squash (map f x) (map f y) i
```

In addition to checking that these clauses are well-typed, Cubical Agda must also check that the clause for `squash x y : x ≡ y` is a path between `map f x` and `map f y`, by substituting `i0` and `i1` for `i : I` in the left- and right-hand sides of the definition.

We have already seen with `isEquiv` that propositions can have computational content. In the following example, we use propositional truncation to hide the mathematical content of a type without disturbing its computational content.

Example 2.1 (Cost monad). By pairing an element of \mathbb{N} with the output of a function, we obtain a counter that can be incremented at will but whose value is “hidden” to the equality type:

```
Cost : (A : Type) → Type
Cost A = A ×  $\mathbb{N}$ 

Cost ≡ : (x y : Cost A) → fst x ≡ fst y → x ≡ y
Cost ≡ (x, cx) (y, cy) h i = (h i, squash cx cy i)
```

We then define a monad structure on this type which counts the number of binds as a simple proxy for computation steps, by starting at 0 and incrementing once at each bind ($_ \gg= _$). Using $\text{Cost} \equiv$, it is straightforward to prove that these definitions satisfy the monad laws. (In the definition of $_ \gg= _$, map2 is a binary version of map .)

```
 $\_ \gg= \_$  : Cost A → (A → Cost B) → Cost B
(x, m)  $\gg=$  g with g x
... | (y, n) = (y, map suc (map2  $\_ + \_$  m n))

return : A → Cost A
return x = (x, | 0 |)
```

The computational behavior of this monad is precisely the same as one in which the propositional truncation is omitted. For example, if we compare a naive implementation of Fibonacci to a tail-recursive one, the former requires many more recursive calls, as expected:

<pre>fib : \mathbb{N} → Cost \mathbb{N} fib 0 = return 0 fib 1 = return 1 fib (suc (suc x)) = do y ← fib (suc x) z ← fib x return (y + z) _ : fib 20 ≡ (6765, 21890) _ = refl</pre>	<pre>fibAux : \mathbb{N} → \mathbb{N} → \mathbb{N} → Cost \mathbb{N} fibAux 0 res _ = return res fibAux (suc x) res prev = do r ← fibAux x (res + prev) res return r fibTail : \mathbb{N} → Cost \mathbb{N} fibTail 0 = return 0 fibTail (suc x) = fibAux x 1 0 _ : fibTail 20 ≡ (6765, 19) _ = refl</pre>
--	--

However, because we have truncated the cost, we are able to prove that fib and fibTail are *equal* as functions using $\text{Cost} \equiv$, despite having different runtime behavior: $\text{fibEq} : \text{fib} \equiv \text{fibTail}$.

Set quotients. The main HIT that we use in this paper is the *set quotient*, which quotients a type by an arbitrary relation, yielding a set. It has three constructors: $[_]$, which includes elements of the underlying type, eq/ , which equates all pairs of related elements, and squash/ , which ensures that the resulting type is a set:³

```
data  $\_ / \_$  (A : Type) (R : A → A → Type) : Type where
  [ ] : (a : A) → A / R
  eq/ : (a b : A) → (r : R a b) → [ a ] ≡ [ b ]
  squash/ : isSet (A / R)
```

³If we omitted squash/ , then the unit type quotiented by the total relation would be a HIT generated by $[\text{tt}]$ and a path from $[\text{tt}]$ to itself. But this is the circle type S^1 , which is not a set [Univalent Foundations Program 2013, Section 8.1].

Once again, we can write functions out of A / R by pattern-matching; this amounts to writing a function out of A (the clause for `[_]`) which sends R -related elements of A to equal results (the clause for `eq/`), such that the image of the function is a set (the clause for `squash/`). We therefore obtain the universal property of set quotients, which states that functions from A / R to a set are precisely the functions out of A that respect R :

$$\text{setQuotUniversal} : \{R : A \rightarrow A \rightarrow \text{Type}\} \{B : \text{Type}\} \rightarrow \text{isSet } B \rightarrow \\ (A / R \rightarrow B) \simeq (\Sigma [f \in (A \rightarrow B)] ((a b : A) \rightarrow R a b \rightarrow f a \equiv f b))$$

If we additionally assume that R is a proposition-valued equivalence relation, then we can show (by propositional extensionality) that the set quotients are *effective*, in the sense that if `[a] ≡ [b]` then $R a b$ [Voevodsky 2015].

Example 2.2 (Rational numbers). In Cubical Agda, we can define rational numbers \mathbf{Q} (below, left) as pairs of integers and nonzero natural numbers, set quotiented by equality of cross multiplication. Contrast this with the standard type-theoretic definition of rational numbers \mathbf{Q}' (below, right) as pairs of *coprime* numbers [The Mathematical Components Team 2020, rat]:

$$\begin{array}{ll} _ \sim _ : \mathbb{Z} \times \mathbb{N}_+ \rightarrow \mathbb{Z} \times \mathbb{N}_+ \rightarrow \text{Type} & \mathbf{Q}' : \text{Type} \\ (a, b) \sim (c, d) = a * \mathbb{N}_+ d \equiv c * \mathbb{N}_+ b & \mathbf{Q}' = \Sigma [(a, b) \in \mathbb{Z} \times \mathbb{N}_+] (\text{coprime } (\text{abs } a) b) \\ \mathbf{Q} : \text{Type} & \\ \mathbf{Q} = (\mathbb{Z} \times \mathbb{N}_+) / _ \sim _ & \end{array}$$

In essence, \mathbf{Q}' represents rational numbers by choosing a normal-form representative of each equivalence class; because these normal forms are unique, one avoids the need for quotients or setoids. However, the tradeoff is that all operations on \mathbf{Q}' must maintain the coprimality invariant—which can be prohibitively expensive in practice [Chyzak et al. 2014]—whereas operations on the set quotient \mathbf{Q} are free to return any representative as their result.

Many quotients, such as the untyped λ -calculus modulo β -conversion, cannot be defined in type theory without HITs, as they lack a normal form representation. This can be addressed by passing to *setoids*, which supplant the equality type in favor of an explicitly-given equivalence relation for each type [Barthe et al. 2003]. Users of setoids must manually prove that operations respect these relations, which is laborious and can lead to subtle bugs: it was recently discovered that the algebraic hierarchy in Agda’s standard library incorrectly allows a type of rings to be equipped with one setoid structure used by the addition and a different setoid structure used by the multiplication.⁴

3 THE STRUCTURE IDENTITY PRINCIPLE

The structure identity principle (SIP) is the informal principle that properties of mathematical structures should be invariant under isomorphisms of such structures. Making this precise requires committing to a notion of structure; consequently, different formalizations of the SIP are possible, varying in how structures are represented as well as in their scope and generality. The definition we present is a slight reformulation of one given by Escardó [2019]; ours differs by phrasing the condition in terms of dependent paths, which are particularly convenient in Cubical Type Theory.

3.1 Structures

We begin by defining what we mean by a structure; for us, structures are defined over a carrier type and are equipped with a notion of structure-preserving equivalence. Then, we single out the

⁴<https://lists.chalmers.se/pipermail/agda/2020/012009.html>

structures (with their equivalences) that are well-behaved enough for the SIP to apply, which we call univalent structures. The basic definitions closely follow Escardó's account.

A *structure* is a function $S : \mathbf{Type} \rightarrow \mathbf{Type}$. The type of S -structures is defined as follows.

$$\mathbf{TypeWithStr} S = \Sigma [X \in \mathbf{Type}] (S X)$$

A notion of S -structure-preserving equivalences is a term $\iota : \mathbf{StrEquiv} S$, where

$$\mathbf{StrEquiv} S = (A B : \mathbf{TypeWithStr} S) \rightarrow \mathbf{fst} A \simeq \mathbf{fst} B \rightarrow \mathbf{Type}$$

Given two S -structures $A B : \mathbf{TypeWithStr} S$ and an equivalence between the underlying types $e : \mathbf{fst} A \simeq \mathbf{fst} B$, the type $\iota A B e$ is the type of proofs that e is an S -structure-preserving equivalence between A and B . The type of S -structure-preserving equivalences (henceforth, S -structured equivalences) between A and B is then given by

$$A \simeq [\iota] B = \Sigma [e \in \mathbf{fst} A \simeq \mathbf{fst} B] (\iota A B e)$$

We say that (S, ι) defines a *univalent structure* if we have a term of the following type.

$$\begin{aligned} \mathbf{UnivalentStr} S \iota = \{ & A B : \mathbf{TypeWithStr} S \} (e : \mathbf{fst} A \simeq \mathbf{fst} B) \\ & \rightarrow (\iota A B e) \simeq \mathbf{PathP} (\lambda i \rightarrow S (\mathbf{ua} e i)) (\mathbf{snd} A) (\mathbf{snd} B) \end{aligned}$$

A univalent structure is a pair (S, ι) which satisfies the SIP. Our definition is equivalent to Escardó's *standard notion of structure*, but interacts better with cubical machinery.

THEOREM 3.1 (SIP). *For $S : \mathbf{Type} \rightarrow \mathbf{Type}$ and $\iota : \mathbf{StrEquiv} S$, we have a term*

$$\mathbf{SIP} : \mathbf{UnivalentStr} S \iota \rightarrow (A B : \mathbf{TypeWithStr} S) \rightarrow (A \simeq [\iota] B) \simeq (A \equiv B)$$

PROOF. Suppose $\theta : \mathbf{UnivalentStr} S \iota$, and let $A B : \mathbf{TypeWithStr} S$. By applying $\Sigma \mathbf{Path} \simeq \mathbf{Path} \Sigma$ on the right, we reduce the goal to the following.

$$\Sigma [e \in \mathbf{fst} A \simeq \mathbf{fst} B] (\iota A B e) \simeq \Sigma [p \in \mathbf{fst} A \equiv \mathbf{fst} B] (\mathbf{PathP} (\lambda i \rightarrow S (p i)) (\mathbf{snd} A) (\mathbf{snd} B))$$

The first components are equivalent by univalence, while the input θ is a proof that the second components are equivalent over the first equivalence. It is straightforward to check that any pair of equivalences $e_X : X \simeq X'$ and $e_Y : (x : X) \rightarrow Y x \simeq Y' x$ ($\mathbf{equivFun} e_X x$) gives rise to an equivalence $\Sigma [x \in X] (Y x) \simeq \Sigma [x' \in X'] (Y' x')$; applying this lemma concludes the proof. \square

Writing the underlying map of this equivalence explicitly, we have the following.

$$\begin{aligned} \mathbf{sip} : A \simeq [\iota] B & \rightarrow A \equiv B \\ \mathbf{sip} (e, p) i & = (\mathbf{ua} e i, \mathbf{equivFun} (\theta e) p i) \end{aligned}$$

The **SIP** itself is thus a nearly trivial consequence of the definition of univalent structure. The work, then, is in equipping structures with useful definitions of structured equivalence that satisfy **UnivalentStr**. (Note that there is always a distinctly *useless* definition of structured equivalence for any structure $S : \iota A B e = \mathbf{PathP} (\lambda i \rightarrow S (p i)) (\mathbf{snd} A) (\mathbf{snd} B)$, which is trivially univalent.) In the following, we show how this can be done systematically and with support by automation.

Axioms. A structure can typically be separated into two main components: a *raw structure* consisting of operations on the carrier type, and propositional *axioms* (often in the form of equations) governing the behavior of the raw structural components. As observed by Escardó [2019], the latter can be ignored for the purpose of defining structured equivalences.

Definition 3.2. Let $S : \mathbf{Type} \rightarrow \mathbf{Type}$ be a structure and $\iota : \mathbf{StrEquiv} S$. Suppose we have axioms on S -structures in the form of a predicate $ax : \mathbf{TypeWithStr} S \rightarrow \mathbf{Type}$. We then define

$$\mathbf{AxiomsStr} S ax X = \Sigma [s \in S X] (ax (X, s))$$

$$\mathbf{AxiomsEquivStr} \iota ax (X, s, _) (Y, t, _) e = \iota (X, s) (Y, t) e$$

LEMMA 3.3. *Let S, ι , and ax be as above, and assume that ax is proposition-valued. Given some $\theta : \mathbf{UnivalentStr} S \iota$, we have $\mathbf{UnivalentStr} (\mathbf{AxiomsStr} S ax) (\mathbf{AxiomsEquivStr} \iota ax)$.*

PROOF. Let (X, s, a) and $(Y, t, b) : \mathbf{TypeWithStr} (\mathbf{AxiomsStr} S ax)$ and $e : X \simeq Y$. By θ e and the definition of $\mathbf{AxiomsEquivStr}$, we know $\mathbf{AxiomsEquivStr} \iota (X, s, a) (Y, t, b)$ is equivalent to $\mathbf{PathP} (\lambda i \rightarrow S (\mathbf{ua} e i)) s t$, but want it equivalent to $\mathbf{PathP} (\lambda i \rightarrow \mathbf{AxiomsStr} S (\mathbf{ua} e i)) (s, a) (t, b)$. By $\Sigma \mathbf{Path} \simeq \mathbf{Path} \Sigma$, the latter is equivalent to the following Σ -type:

$$\Sigma [p \in \mathbf{PathP} (\lambda i \rightarrow S (\mathbf{ua} e i)) s t] (\mathbf{PathP} (\lambda i \rightarrow ax (\mathbf{ua} e i, p i)) a b)$$

Because we have assumed $ax (\mathbf{ua} e i, p i)$ is a proposition, the second component is contractible, and so the Σ -type is equivalent to its first component as needed. \square

We can use the SIP to transport proofs of axioms between equivalent raw structures to obtain equivalent structures-with-axioms.

COROLLARY 3.4 (INDUCED STRUCTURES). *Let S, ι and ax be as above, and $\theta : \mathbf{UnivalentStr} S \iota$. Suppose we have $(X, s, a) : \mathbf{TypeWithStr} (\mathbf{AxiomsStr} S ax)$ and $(Y, t) : \mathbf{TypeWithStr} S$ and a structured equivalence $e : (X, s) \simeq [\iota] (Y, t)$. Then there exists $b : ax (Y, t)$ with $(X, s, a) \equiv (Y, t, b)$.*

PROOF. Set $b = \mathbf{subst} ax (\mathbf{sip} e) a$; recall that $\mathbf{sip} e : (X, s) \equiv (Y, t)$, so b has type $ax (Y, t)$. It follows by definition that $e : (X, s, a) \simeq [\mathbf{AxiomsEquivStr} \iota ax] (Y, t, b)$, from which we get the desired path by applying \mathbf{sip} once more. \square

Before we show how to build raw structures, we first illustrate the utility of the above lemmas.

Example 3.5 (Monoids). We say that X supports a *raw* monoid structure when it has a distinguished neutral element ε and a binary operation $_ \cdot _$:

$$\mathbf{RawMonoidStructure} X = X \times (X \rightarrow X \rightarrow X)$$

A raw monoid is simply a type equipped with a raw monoid structure: $\mathbf{RawMonoid} = \mathbf{TypeWithStr} \mathbf{RawMonoidStructure}$. A (fully-cooked) monoid structure is a raw monoid structure whose underlying type is a set, and whose binary operation is unital and associative:

$$\mathbf{MonoidAxioms} : \mathbf{RawMonoid} \rightarrow \mathbf{Type}$$

$$\mathbf{MonoidAxioms} (X, \varepsilon, _ \cdot _) = (\mathbf{isSet} X)$$

$$\times (\forall x y z \rightarrow x \cdot (y \cdot z) \equiv (x \cdot y) \cdot z)$$

$$\times (\forall x \rightarrow (x \cdot \varepsilon \equiv x) \times (\varepsilon \cdot x \equiv x))$$

$$\mathbf{MonoidStructure} X = \Sigma [(\varepsilon, _ \cdot _) \in \mathbf{RawMonoidStructure} X] (\mathbf{MonoidAxioms} (X, \varepsilon, _ \cdot _))$$

The type of monoids is thus $\mathbf{Monoid} = \mathbf{TypeWithStr} \mathbf{MonoidStructure}$. A structured equivalence of monoids is a monoid isomorphism in the usual sense: an equivalence that commutes with the monoid operations.

$$\mathbf{MonoidEquiv} : (M N : \mathbf{Monoid}) \rightarrow \mathbf{fst} M \simeq \mathbf{fst} N \rightarrow \mathbf{Type}$$

$$\mathbf{MonoidEquiv} (_, (\varepsilon_M, _ \cdot _), _) (_, (\varepsilon_N, _ \cdot _), _) (\varphi, _) = (\varphi \varepsilon_M \equiv \varepsilon_N)$$

$$\times (\forall x y \rightarrow \varphi (x \cdot y) \equiv (\varphi x) \cdot (\varphi y))$$

It is now easy to see that the monoid axioms are all proposition-valued: the first axiom requires the carrier type to be a set, which forces the equations to be propositions. In order to apply the SIP to monoids, it remains only to check that `RawMonoidStructure` is a `UnivalentStr`. Proving this by hand is not so easy, but rather than proving it directly, we will build it up from elementary combinators in a way that preserves `UnivalentStr`.

3.2 Building structures

Having dealt with axioms, we now define a collection of combinators from which we can build raw structures, their structured equivalences, and proofs they are univalent. We currently support:

`Structures` $S\ X, T\ X := X \mid A \mid S\ X \times T\ X \mid S\ X \rightarrow T\ X \mid \text{Maybe}\ (S\ X)$

The raw fragment of all structures considered in this paper can be built using the above grammar. (Of course, `Maybe` is just one of many inductive types we could consider.)

Constant and pointed structures. Our two base cases are the constant structure $\lambda\ _ \rightarrow A$ and the pointed structure $\lambda\ X \rightarrow X$. Given $(X, a)\ (Y, a') : \text{TypeWithStr}\ (\lambda\ _ \rightarrow A)$, an equivalence $e : X \simeq Y$ is structured when we have a path $a \equiv a'$; this is trivially univalent. For $(X, x)\ (Y, y) : \text{TypeWithStr}\ (\lambda\ X \rightarrow X)$, $e : X \simeq Y$ is structured when `equivFun` $e\ x \equiv y$, i.e., the chosen elements are related by the equivalence. Univalence of the pointed structure is a consequence of `uaβ`.

Product structures. We define the product structure of two structures equipped with notions of structured equivalence (S_1, ι_1) and (S_2, ι_2) , as follows:

`ProductStructure` $S_1\ S_2\ X = S_1\ X \times S_2\ X$

`ProductEquivStr` $\iota_1\ \iota_2\ (X, s_1, s_2)\ (Y, t_1, t_2)\ e = \iota_1\ (X, s_1)\ (Y, t_1)\ e \times \iota_2\ (X, s_2)\ (Y, t_2)\ e$

The product of `UnivalentStructures` is univalent:

`productUnivalentStr` $:(\iota_1 : \text{StrEquiv}\ S_1)\ (\iota_2 : \text{StrEquiv}\ S_2)$
 $\rightarrow \text{UnivalentStr}\ S_1\ \iota_1 \rightarrow \text{UnivalentStr}\ S_2\ \iota_2$
 $\rightarrow \text{UnivalentStr}\ (\text{ProductStructure}\ S_1\ S_2)\ (\text{ProductEquivStr}\ \iota_1\ \iota_2)$

The corresponding HoTT/UF-proof of Escardó [2019] uses the type-theoretic Yoneda lemma [Rijke 2012, §2.8], but the nice interplay of Σ -types and dependent path types in Cubical Agda makes it simple to construct the desired equivalence in the conclusion directly.

`productUnivalentStr` $_ _ \theta_1\ \theta_2\ e = \text{compEquiv}\ (\Sigma\text{-cong-equiv}\ (\theta_1\ e)\ (\lambda\ _ \rightarrow \theta_2\ e))\ \Sigma\text{Path}\simeq\text{Path}\Sigma$

Here, the lemma `Σ -cong-equiv` with a non-dependent second argument gives an equivalence of products from the equivalences on the projections $\theta_1\ e$ and $\theta_2\ e$.

Function structures. We say that two $(\lambda\ X \rightarrow S\ X \rightarrow T\ X)$ -structures f and g are related over e when they take related S -structures to related T -structures:

`FunctionEquivStr` $:\text{StrEquiv}\ S \rightarrow \text{StrEquiv}\ T \rightarrow \text{StrEquiv}\ (\lambda\ X \rightarrow S\ X \rightarrow T\ X)$

`FunctionEquivStr` $\iota_1\ \iota_2\ (X, f)\ (Y, g)\ e = \forall \{s\ t\} \rightarrow \iota_1\ (X, s)\ (Y, t)\ e \rightarrow \iota_2\ (X, f\ s)\ (Y, g\ t)\ e$

Maybe structures. The definition of structured equivalence for $\lambda\ X \rightarrow \text{Maybe}\ (S\ X)$ proceeds by cases on the pair of structures:

`MaybeEquivStr` $:\text{StrEquiv}\ S \rightarrow \text{StrEquiv}\ (\lambda\ X \rightarrow \text{Maybe}\ (S\ X))$

`MaybeEquivStr` $\iota\ (X, \text{nothing})\ (Y, \text{nothing})\ e = \text{Unit}$

`MaybeEquivStr` $\iota\ (X, \text{nothing})\ (Y, \text{just}\ x)\ e = \perp$

$\text{MaybeEquivStr } \iota (X, \text{just } x) (Y, \text{nothing}) e = \perp$
 $\text{MaybeEquivStr } \iota (X, \text{just } x) (Y, \text{just } y) e = \iota (X, x) (Y, y) e$

Transport structures. Although we have given a single definition for each type former thus far, we note that there is often more than one potentially useful definition of structured equivalence for a given $S : \text{Type} \rightarrow \text{Type}$. For example, structured equivalences for $\lambda X \rightarrow \text{Maybe } X$ can also be defined using the functorial action of `Maybe` as follows.

$\text{MaybeEquivStr}' : \text{StrEquiv } (\lambda X \rightarrow \text{Maybe } X)$
 $\text{MaybeEquivStr}' (X, s) (Y, t) e = \text{map-Maybe } (\text{equivFun } e) s \equiv t$

Indeed, this alternative definition is sometimes more convenient; we will use it (in slightly generalized form) in [Section 4.2](#). It belongs to a class of structured equivalence definitions that arise from a functorial action on equivalences (here, `map-Maybe`) and which are captured by the SIP of [Coquand and Danielsson \[2013\]](#). We call these *transport structures* and define them as follows.

$\text{EquivAction} : (S : \text{Type} \rightarrow \text{Type}) \rightarrow \text{Type}_1$
 $\text{EquivAction } S = \{X Y : \text{Type}\} \rightarrow X \simeq Y \rightarrow S X \simeq S Y$
 $\text{TransportStr} : \{S : \text{Type} \rightarrow \text{Type}\} \rightarrow \text{EquivAction } S \rightarrow \text{Type}_1$
 $\text{TransportStr } \{S\} \alpha = \{X Y : \text{Type}\} (e : X \simeq Y) (s : S X) \rightarrow \text{equivFun } (\alpha e) s \equiv \text{subst } S (\text{ua } e) s$

Rather than a type of proofs that an equivalence is structured, we have an action by the structure on equivalences; this action is “correct” when it agrees with `transport`. ([Coquand and Danielsson](#) require only that $\text{equivFun } (\alpha (\text{idEquiv } X)) s \equiv s$, but univalence implies that this is equivalent to our definition.) Any such action gives rise to a notion of structured equivalence defined as follows.

$\text{EquivAction} \rightarrow \text{StrEquiv} : \text{EquivAction } S \rightarrow \text{StrEquiv } S$
 $\text{EquivAction} \rightarrow \text{StrEquiv } \alpha (X, s) (Y, t) e = \text{equivFun } (\alpha e) s \equiv t$

When we have an element of `TransportStr` α , this notion of structured equivalence is univalent. In fact, `TransportStr` α and `UnivalentStr` ($\text{EquivAction} \rightarrow \text{StrEquiv } \alpha$) are equivalent conditions. Note, however, that an element of `StrEquiv` S does not induce an element of `EquivAction` S in any useful way; thus our primary definition of structure is the more permissive of the two, although transport structures are also closed under the grammar at the head of this section.

Transport structures are particularly convenient when we define structures on function types. If we have an action on the domain, however, we can give a more convenient definition as follows.

$\text{FunctionEquivStr}^+ : \text{EquivAction } S \rightarrow \text{StrEquiv } T \rightarrow \text{StrEquiv } (\lambda X \rightarrow S X \rightarrow T X)$
 $\text{FunctionEquivStr}^+ \alpha \iota (X, f) (Y, g) e = \forall s \rightarrow \iota (X, f s) (Y, g (\text{equivFun } (\alpha e) s)) e$

Example 3.6 (Monoids revisited). By [Lemma 3.3](#), to prove the SIP for monoids it remains only to show that `RawMonoidEquiv` is univalent, which we are now ready to do in a systematic way. First, observe that pointed and binary operation structures $\lambda X \rightarrow X$ and $\lambda X \rightarrow (X \rightarrow X \rightarrow X)$ are univalent with their canonical notions of structured equivalence. This is immediate for pointed structures; for binary operations, we obtain the result by applying `FunctionEquivStr`⁺ twice to the pointed structure. Finally, we take the product of these two structures to see that `RawMonoidEquiv` defines a univalent structure. Putting everything together, we get the desired SIP for monoids:

$\text{MonoidPath} : (M N : \text{Monoid}) \rightarrow (M \simeq [\text{MonoidEquiv}] N) \simeq (M \equiv N)$
 $\text{MonoidPath} = \text{SIP monoidUnivalentStr}$

Let us now illustrate the utility of [Corollary 3.4](#). Given $M : \text{Monoid}$ and $N : \text{RawMonoid}$ and a `RawMonoidEquiv` between N and the underlying `RawMonoid` of M , [Corollary 3.4](#) allows us

to transport the axioms from M to N so that we obtain an induced $N' : \mathbf{Monoid}$ with $M \equiv N'$. Specialized to the case of unary and binary numbers, called \mathbb{N} and \mathbf{Bin} in agda/cubical [Vezzosi et al. 2019, Section 2.1], we could show that these types are equal $\mathbf{Monoids}$ by proving only that $(0, +)$ is a monoid on \mathbb{N} , and defining $+_{\mathbf{Bin}}$ such that the function underlying the $\mathbb{N} \simeq \mathbf{Bin}$ equivalence is a monoid homomorphism. By definition, $\mathbb{N} \rightarrow \mathbf{Bin}$ sends 0 to $\mathbf{bin}0$, so we only have to show that

$$(x \ y : \mathbb{N}) \rightarrow \mathbb{N} \rightarrow \mathbf{Bin} (x + y) \equiv (\mathbb{N} \rightarrow \mathbf{Bin} x) +_{\mathbf{Bin}} (\mathbb{N} \rightarrow \mathbf{Bin} y)$$

which can be readily achieved by \mathbb{N} -induction. Therefore, we can transport the theory of natural numbers from \mathbb{N} to \mathbf{Bin} without doing any proofs by \mathbf{Bin} -induction. This is similar to CoqEAL, in which such morphism lemmas were used to ensure that all proofs could be completed on proof-oriented types while computation-oriented types were only used for programming, thereby achieving a clear *separation of concerns* [Dijkstra 1974].

Example 3.6 suggests that other algebraic structures can be treated similarly; indeed, the same proof strategy applies directly to more complex structures like groups and rings. In fact, these proofs are so uniform that we can automate them.

Automation. Using Agda's reflection mechanism, we have defined tactics for automatically generating definitions of structured equivalence and proofs they are univalent. For instance, the macro `AutoEquivStr` ($\lambda (X : \mathbf{Type}) \rightarrow X \times (X \rightarrow X \rightarrow X)$) generates a definition of structured equivalence for raw monoid structures, while `AutoUnivalentStr` ($\lambda (X : \mathbf{Type}) \rightarrow X \times (X \rightarrow X \rightarrow X)$) produces a proof that the definition is univalent. Our tactics opt for transport structures in the domain of a function type and univalent structures otherwise, but we allow annotations to override the default choice; for example, the following definition of equivalence for a queue data structure will use the transport structure for the codomain of the final function.

```

RawQueueEquivStr =
  AutoEquivStr ( $\lambda (X : \mathbf{Type}) \rightarrow X \times (A \rightarrow X \rightarrow X) \times (X \rightarrow \mathbf{Transp}[\mathbf{Maybe} (X \times A)])$ )

```

Given such a definition, the macro `AutoStructure` removes annotations to produce the actual structure definition, in this case $\lambda (X : \mathbf{Type}) \rightarrow X \times (A \rightarrow X \rightarrow X) \times (X \rightarrow \mathbf{Maybe} (X \times A))$.

4 REPRESENTATION INDEPENDENCE THROUGH THE SIP

Before showing how to generalize the SIP to relational correspondences, we first show how to use the SIP to obtain representation independence results of interest to programmers. First, we consider two representations of matrices: as functions out of a finite set of indices, which are well-suited to proofs, and as vectors, which are well-suited to computations. Then, we revisit the queue example from the introduction, and show how to achieve representation independence through a combination of set quotients and the SIP.

4.1 Matrices

Matrices are a standard example in which dependent types can ensure the well-definedness of operations such as multiplication. We can achieve this in Agda with length-indexed lists, or vectors:

```

data Vec (A : Type) :  $\mathbb{N} \rightarrow \mathbf{Type}$  where
  [] : Vec A zero
  _::_ :  $\forall \{n\} (x : A) (xs : \mathbf{Vec} A n) \rightarrow \mathbf{Vec} A (\mathbf{suc} n)$ 

VecMatrix : (A : Type) (m n :  $\mathbb{N}$ )  $\rightarrow \mathbf{Type}$ 
VecMatrix A m n = Vec (Vec A n) m

```

We can define operations on vectors such as `map`, `replicate`, addition, and multiplication in the usual way. Unfortunately, this type is ill-suited for proofs, because as must perform inductions that match the structure of the functions one has written: it is quite difficult, for example, to prove that matrix addition `addVecMatrix` is commutative. Ideally, one would prefer to reason about matrix addition in a *pointwise* fashion, in which commutativity is trivial. We therefore consider a second implementation better-suited to such reasoning:

```
data Fin : ℕ → Type where
  zero : {n : ℕ} → Fin (suc n)
  suc  : {n : ℕ} (i : Fin n) → Fin (suc n)

FinMatrix : (A : Type) (m n : ℕ) → Type
FinMatrix A m n = Fin m → Fin n → A
```

If we now fix natural numbers m and n and assume \mathbf{G} is an additive abelian group, it is trivial to define matrix addition and prove it commutative by function extensionality:

```
addFinMatrix : FinMatrix G m n → FinMatrix G m n → FinMatrix G m n
addFinMatrix M N i j = M i j + N i j

addFinMatrixComm : (M N : FinMatrix G m n) → addFinMatrix M N ≡ addFinMatrix N M
addFinMatrixComm M N i k l = comm (M k l) (N k l) i
```

Similarly, we can easily define the zero matrix and negation of matrices, and prove that these form an abelian group. In the absence of function extensionality, these results are typically stated using setoids, which complicates the proofs [Wood 2019]. Alternatively, MathComp [The Mathematical Components Team 2020] encodes matrices by the graphs of their indexing functions, which makes proofs simpler without relying on setoids, but requires a lot of additional theory that is trivialized by using functions directly.

Although `FinMatrix` is convenient for proofs, it is less natural for computations than `VecMatrix`. Consider the following integer matrices:

```
M : FinMatrix Z 2 2
M i j = if i == j then 1 else 0

N : FinMatrix Z 2 2
N i j = if i == j then 0 else 1
```

Because `if` does not compute on variables, `addFinMatrix M N ≡ (λ _ _ → 1)` does not hold by `refl`. In a `VecMatrix` representation, however, the analogous equation does follow from computation (i.e., by `refl`). Luckily, we can define functions between the two representations and prove that they form an equivalence using `funExt`, which we can transform into a path by `ua`:

```
FinMatrix ≡ VecMatrix : (A : Type) (m n : ℕ) → FinMatrix A m n ≡ VecMatrix A m n
FinMatrix ≡ VecMatrix A m n = ua (FinMatrix ≃ VecMatrix A m n)
```

Using this equivalence, we can easily obtain the aforementioned equation in `FinMatrix` by transferring the analogous equation in `VecMatrix`:

```
_ : addFinMatrix M N ≡ (λ _ _ → 1)
_ = replaceGoal (FinMatrix ≃ VecMatrix Z 2 2) refl
```

The `replaceGoal` lemma takes an equivalence e and transforms a goal of the form $x \equiv y$ into $e^{-1}(ex) \equiv e^{-1}(ey)$, then discharges the former with a proof of the latter. Here, the latter proof is `refl` as `addFinMatrix M N` and `(λ _ _ → 1)` are mapped to definitionally-equal `VecMatrix`s.

We could in fact `transport` the entire abelian group structure on `FinMatrix` to `VecMatrix` along this equivalence, but the operations obtained this way are very naive: for example, the transported

addition on `VecMatrix` converts both arguments to `FinMatrix`, adds them with `addFinMatrix`, then converts them back to `VecMatrix`. What we instead want is an abelian group structure on `VecMatrix` whose addition is given by `addVecMatrix`. We can achieve this using the SIP.

Using the tactics discussed in [Section 3](#), we can automatically obtain a univalent definition of structured equivalence for abelian groups. By [Corollary 3.4](#), we must only show that the function underlying `FinMatrix` \simeq `VecMatrix` sends `addFinMatrix` to `addVecMatrix`. As `FinMatrix` is well-suited for proofs this poses no difficulties, and so we obtain an induced *and equal* abelian group structure on `VecMatrix` whose addition operation is `addVecMatrix`. Using the equality between abelian groups, we can now transport both proofs and programs between the two representations.

4.2 Queues

Thus far, our examples of the SIP have involved equivalent structured types; now, we turn to our original example of two non-equivalent implementations of `Queues`. Fix a set A . Then a raw queue structure consists of an `empty` queue and `enqueue`/`dequeue` functions:

```
RawQueueStructure : Type → Type
RawQueueStructure X = X × (A → X → X) × (X → Maybe (X × A))
```

We add some axioms specifying how these operations should behave: that Q is a set, that `dequeue` of the `empty` queue is `nothing`, and what happens if we `dequeue` after `enqueue`:

```
dequeueEnqueueAxiom Q (empty, enqueue, dequeue) =
  ∀ a q → dequeue (enqueue a q) ≡ just (returnOrEnq a (dequeue q))
  where
    returnOrEnq : A → Maybe (Q × A) → Q × A
    returnOrEnq a nothing = (empty, a)
    returnOrEnq a (just (q, b)) = (enqueue a q, b)
```

These axioms form a proposition because we have assumed Q is a set. As in [Section 3.2](#), we automate the construction of a univalent raw queue structure, then add our axioms to obtain a univalent `QueueStructure`. The `ListQueue` implementation from [Section 1](#) clearly satisfies our queue axioms and therefore admits a `QueueStructure`. Then every `BatchedQueue` (a pair `List A × List A`) corresponds to a single `ListQueue` computed by `appendReverse` (the function sending (xs, ys) to $(xs ++ reverse\ ys)$), but this function is not an isomorphism: multiple `BatchedQueues` are sent to the same `ListQueue`. Worse yet, `BatchedQueue` is *not even an instance of* `QueueStructure`, as it fails to satisfy `dequeueEnqueueAxiom`: we have `dequeue (enqueue c ([b, a], [])) ≡ just (([], [b, c]), a)` but `just (returnOrEnq c (dequeue ([b, a], []))) ≡ just (([c], [b]), a)`.

We can resolve both of these problems by identifying any two `BatchedQueues` sent to the same list by `appendReverse`, which we can do either with a set quotient, or equivalently with the HIT:

```
data BatchedQueueHIT : Type where
  Q⟦_,_⟧ : List A → List A → BatchedQueueHIT
  tilt : ∀ xs ys a → Q⟦ xs ++ [ a ], ys ⟧ ≡ Q⟦ xs, ys ++ [ a ] ⟧
  squash : isSet BatchedQueueHIT
```

We equip `BatchedQueueHIT` with a `QueueStructure` as follows. The `empty` queue is `Q⟦[], []⟧`, and we can define `enqueue` and `dequeue` functions that respect the `tilt` constructor using helper functions similar to `fastcheck`. The `tilt` constructor allows us to shift elements between the ends of the two lists back and forth, thus ensuring that any two `BatchedQueueHIT`s corresponding to the same `ListQueue` are identified. In particular, `appendReverse` can then be extended to an equivalence `BatchedQueueHIT` \simeq `List A` which is moreover structure-preserving, giving an induced raw

queue structure on [BatchedQueueHIT](#). We may now apply the SIP to transfer the axioms satisfied by [ListQueue](#) to the quotiented [BatchedQueue](#) operations.

5 STRUCTURED EQUIVALENCES FROM STRUCTURED RELATIONS

In the queue case study above, we improve a structured *relation* between implementations (generated by [appendReverse](#)) to a structured *equivalence* by considering a quotient of one of the implementations. In this section, we prove that a generalization of this technique is always applicable as soon as we have what we call a *quasi-equivalence relation* (QER).


5.1 Quasi-equivalence relations

Because we will use set quotients, it is natural (and ultimately necessary) to restrict our attention to proposition-valued relations, i.e., $R : X \rightarrow Y \rightarrow \text{Type}$ such that each $R\ x\ y$ is a proposition.

LEMMA 5.1. *We have the following operations on prop-valued relations:*

- For any prop-valued $R : X \rightarrow Y \rightarrow \text{Type}$, the inverse relation R^{-1} is also prop-valued.
- The identity prop-valued relation on X is the truncated path type: $\text{IdRel } X\ x_0\ x_1 = \parallel x_0 \equiv x_1 \parallel$.
- Given two prop-valued relations $R : X \rightarrow Y \rightarrow \text{Type}$ and $S : Y \rightarrow Z \rightarrow \text{Type}$, we define their prop-valued composite $R \cdot S$ using truncation: $(R \cdot S)\ x\ z = \parallel \Sigma[y \in Y] R\ x\ y \times R\ y\ z \parallel$.
- If $f : X \rightarrow Y$ where Y is a set, then its graph, $(\text{graph } f)\ x\ y = (f\ x \equiv y)$, is a prop-valued relation. In particular, if $e : X \simeq Y$, we define $\text{equivGraph } e = \text{graph } (\text{equivFun } e)$.

Definition 5.2. A prop-valued relation $R : X \rightarrow Y \rightarrow \text{Type}$ is *zigzag-complete* when for any $r_0 : R\ x\ y$, $r_1 : R\ x'\ y$, and $r_2 : R\ x'\ y'$, we have $R\ x\ y'$. This property can be visualized as follows:



Zigzag-complete relations, also known as *quasi-PERs* or *QPERs* [Krishnaswami and Dreyer 2013], are a heterogeneous analogue of partial equivalence relations (PERs). A zigzag-complete relation $R : X \rightarrow Y \rightarrow \text{Type}$ expresses a partial correspondence between elements of X and elements of Y ; intuitively, if x corresponds to y , x' corresponds to y , and x' corresponds to y' , then x should also correspond to y' . Note that a QPER induces a PER on each of the two types: $R \cdot R^{-1}$ on X and $R^{-1} \cdot R$ on Y . We henceforth denote these two relations R^{\leftarrow} and R^{\rightarrow} respectively. In fact, R is zigzag-complete exactly when R^{\leftarrow} and R^{\rightarrow} are PERs. In our setting, we require such a partial correspondence to be *total* by additionally asking for truncated choice functions in either direction.

Definition 5.3. A zigzag-complete relation $R : X \rightarrow Y \rightarrow \text{Type}$ is a *quasi-equivalence relation* (QER) when we have $\parallel \Sigma[y \in Y] R\ x\ y \parallel$ for all $x : X$ and $\parallel \Sigma[x \in X] R\ x\ y \parallel$ for all $y : Y$.

As with QPERs, R is a QER if and only if R^{\leftarrow} and R^{\rightarrow} are equivalence relations. Moreover, the truncated choice functions induce an equivalence between the set quotients X / R^{\leftarrow} and Y / R^{\rightarrow} .

LEMMA 5.4. *Given a QER R , there is an equivalence $e : X / R^{\leftarrow} \simeq Y / R^{\rightarrow}$ such that for every $x : X$ and $y : Y$, we have $\text{equivFun } e\ [x] \equiv [y]$ if and only if $R\ x\ y$.*

PROOF. Consider first the forward map, $X / R^{\leftarrow} \rightarrow Y / R^{\rightarrow}$. Observe that for every $x : X$, we have a map $[] \circ \text{fst} : (\Sigma[y \in Y] R\ x\ y) \rightarrow Y / R^{\rightarrow}$. Moreover, this map is *constant*: given $y\ y' : A$ with $R\ x\ y$ and $R\ x\ y'$, we have $[y] \equiv [y']$ by definition of the quotient. It therefore factors through the propositional truncation [Kraus et al. 2017], giving $f_x : \parallel \Sigma[y \in Y] R\ x\ y \parallel \rightarrow Y / R^{\rightarrow}$ such that $f_x \circ [] \equiv [] \circ \text{fst}$. We precompose with the provided choice function $X \rightarrow \parallel \Sigma[y \in Y] R\ x\ y \parallel$ to

get a map $X \rightarrow Y / R^{\rightarrow}$. To show that this map factors through X / R^{\leftarrow} , it is enough to check that for every $x x' : X$ with $R^{\leftarrow} x x'$ and $y y' : A$ with $R x y$ and $R x' y'$, we have $[y] \equiv [y']$. This follows by composing the relational witnesses and applying `eq/`, using that R is zigzag-complete. Thus we have a map $X / R^{\leftarrow} \rightarrow Y / R^{\rightarrow}$; the candidate inverse is constructed in the same way.

The two inverse conditions are propositions, so it suffices to check that they hold on inputs of the form $[x]$ and $[y]$ respectively, which is a straightforward consequence of the definitions. The bi-implication between `equivFun e [x] ≡ [y]` and $R x y$ is also a straightforward calculation; the forward direction uses effectivity of set quotients by equivalence relations [Voevodsky 2015]. \square

Example 5.5 (Finite multisets). We say that X is a type of *finite multisets (bags) over A* if it has an empty multiset X , an insertion function $A \rightarrow X \rightarrow X$, and a multiplicity function $A \rightarrow X \rightarrow \mathbb{N}$. We can of course impose axioms as well, but for the moment we focus on the raw structure; for simplicity, we will also assume that A has decidable equality. We consider two implementations: lists (`List A`) and *association lists* (`AssocList A = List (A × ℕ)`), where in the latter case each element is tagged with a multiplicity. We define `insert` and `count` for each implementation as follows:

```

addlfeq : (a x : A) → ℕ → ℕ → ℕ
addlfeq a x m n =
  if a == x then m + n else n

module L where
  insert : A → List A → List A
  insert x xs = x :: xs

  count : A → List A → ℕ
  count a [] = 0
  count a (x :: xs) =
    addlfeq a x 1 (count a xs)

module AL where
  insert : A → AssocList A → AssocList A
  insert a [] = (a, 1) :: []
  insert a ((x, n) :: ys) =
    if a == x
    then (x, 1 + n) :: ys
    else (x, n) :: insert a ys

  count : A → AssocList A → ℕ
  count a [] = 0
  count a ((x, n) :: ys) = addlfeq a x n (count a ys)

```

In both cases, the empty list is the empty multiset. Note that *neither* of these implementations satisfies many of the “extensional” laws we expect of multisets. For example, we do not have `insert a (insert a' xs) ≡ insert a' (insert a xs)` in either case.

We now define a QER between `List A` and `AssocList A`, identifying a pair of multisets whenever they assign the same multiplicity to each element of A .

```

R : List A → AssocList A → Type
R xs ys = (a : A) → L.count a xs ≡ AL.count a ys

```

This relation is prop-valued because \mathbb{N} is a set, and is zigzag-complete by transitivity and symmetry of \equiv . We define functions in both directions using the implementations’ `insert` functions.

```

φ : List A → AssocList A
φ [] = []
φ (x :: xs) = AL.insert x (φ xs)

ψ : AssocList A → List A
ψ [] = []
ψ ((x, 0) :: ys) = ψ ys
ψ ((x, suc n) :: ys) = L.insert x (ψ ((x, n) :: ys))

```

Finally, proving $\forall xs \rightarrow R xs (\phi xs)$ and $\forall ys \rightarrow R (\psi ys) ys$ takes a bit more work, but the intuition is that these functions preserve the `count` of each element of A . Note that the two derived equivalence relations R^{\leftarrow} and R^{\rightarrow} also wind up identifying multisets on either side precisely when they assign the same multiplicity to each element of A . By [Lemma 5.4](#), we have an equivalence between `List A` and `AssocList A` after quotienting by these relations on either side.

5.2 Structured relations

We now generalize to *structured* relations and equivalences, defining a condition on notions of structured relation such that any structured QER induces a structured equivalence between quotients. First, as with structured equivalences, a candidate *notion of structured relation* for a structure S assigns, for each relation R on types X and Y and structures $s : S X$, $t : S Y$, a type of witnesses that the relation is structured:

$$\begin{aligned} \text{StrRel} &: (S : \text{Type} \rightarrow \text{Type}) \rightarrow \text{Type}_1 \\ \text{StrRel } S &= \{X Y : \text{Type}\} \rightarrow (X \rightarrow Y \rightarrow \text{Type}) \rightarrow (S X \rightarrow S Y \rightarrow \text{Type}) \end{aligned}$$

(For simplicity, we do not require the input relation to be prop-valued, but our correctness conditions will only constrain behavior on prop-valued relations.) We now identify a correctness condition on notions of structured relations.

Definition 5.6. A candidate $\rho : \text{StrRel } S$ is *suitable* when the following hold:

- (1) *Set- and prop-preservation:* If X is a set, then $S X$ is a set. If R is a prop-valued-relation, then ρR is a prop-valued relation.
- (2) *Symmetry:* For any prop-valued relation R , if $\rho R s t$, then $\rho (R^{-1}) t s$.
- (3) *Transitivity:* For any prop-valued relations R and R' , if $\rho R s t$ and $\rho R' t u$, then $\rho (R \cdot R') t u$.
- (4) *Descent to quotients:* If $R : X \rightarrow X \rightarrow \text{Type}$ is a prop-valued equivalence relation and $\rho R s s$ for some $s : S X$, then there is a *unique* $\bar{s} : S (X / R)$ such that we have some $r : \rho (\text{graph } _) s \bar{s}$; that is, the type $\Sigma[\bar{s} \in S (X / R)] (\rho (\text{graph } _) s \bar{s})$ is contractible.

Note that the condition of being suitable is a proposition. First, we check that these conditions are sufficient for our original motivation: obtaining structured equivalences from structured QERs.

THEOREM 5.7. *Let $\rho : \text{StrRel } S$ be a suitable notion of structured relation, and let structured types $(X, s) : \text{TypeWithStr } S$ and $(Y, t) : \text{TypeWithStr } B$ be given. For any QER $R : X \rightarrow Y \rightarrow \text{Type}$ structured by some $r : \rho R s t$, the following exist:*

- (1) a structure $\bar{s} : S (X / R^{\leftarrow})$ with $\rho (\text{graph } _) s \bar{s}$,
- (2) a structure $\bar{t} : S (Y / R^{\rightarrow})$ with $\rho (\text{graph } _) t \bar{t}$,
- (3) an element of $\rho (\text{equivGraph } e) \bar{s} \bar{t}$, where $e : X / R^{\leftarrow} \simeq Y / R^{\rightarrow}$ is the equivalence obtained by applying [Lemma 5.4](#).

In other words, we have the dotted lines in the following picture, where the inner lines indicate relations and each outer line indicates an element of ρ over the inner relation.

$$\begin{array}{ccc} & r : \rho R s t & \\ s & \xrightarrow{\quad} & t \\ \exists! \downarrow & \begin{array}{ccc} X & \xrightarrow{R} & Y \\ \downarrow \text{graph } _ & & \downarrow \text{graph } _ \\ X / R^{\leftarrow} & \xrightarrow[e \simeq]{} & Y / R^{\rightarrow} \end{array} & \downarrow \exists! \\ & \bar{s} \text{-----} \bar{t} & \end{array}$$

PROOF. By symmetry and transitivity applied with r , we have an element of $\rho R^{\leftarrow} s s$. By descent to quotients, we thus obtain $\bar{s} : S (X / R^{\leftarrow})$ as above. We obtain \bar{t} analogously. From $\rho R s t$, $\rho (\text{graph } _) s \bar{s}$, and $\rho (\text{graph } _) t \bar{t}$, we get a proof of $\rho (\text{graph } _)^{-1} \cdot R \cdot \text{graph } _) \bar{s} \bar{t}$ by

symmetry and transitivity. We arrive at the final condition by checking that the composite relation $(\text{graph } \llbracket _ \rrbracket)^{-1} \cdot R \cdot \text{graph } \llbracket _ \rrbracket$ is equal to $\text{equivGraph } e$. \square

If the restriction of ρ to equivalences is moreover a *univalent* notion of structured equivalence, we obtain a *path* between $(X / R^{\leftarrow}, \bar{s})$ and $(Y / R^{\rightarrow}, \bar{t})$ as a corollary.

Definition 5.8. A notion of structured relation $\rho : \text{StrRel } S$ is *univalent* when it is suitable and its restriction to equivalences, $(\lambda (X, s) (Y, t) e \rightarrow \rho (\text{equivGraph } e) s t) : \text{StrEquiv } S$, is univalent.

As in [Section 3](#), we can show that the collection of suitable and univalent relational structures is closed under various type formers with reasonable definitions of structured relation. (Notably, the constant structure $\lambda _ \rightarrow A$ is only suitable when A is a set.)

Positive structures $P X, Q X := X \mid A \mid P X \times Q X \mid \text{Maybe } (P X)$
Structures $S X, T X := P X \mid S X \times T X \mid P X \rightarrow S X \mid \text{Maybe } (S X)$

One major departure from the situation in [Section 3](#) is that the suitable relation structures are *not* closed under function types in general, but only in restricted cases. We can see the issue by considering descent to quotients for the candidate structure $\lambda X \rightarrow (X \rightarrow X) \rightarrow X$. Given an instance of this structure $f : (X \rightarrow X) \rightarrow X$ and a QER $R : X \rightarrow X \rightarrow \text{Type}$, we have no way of obtaining an induced structure $\bar{f} : (X / R \rightarrow X / R) \rightarrow X / R$: at some point, we would need to extract maps $X \rightarrow X$ from maps $X / R \rightarrow X / R$. Nevertheless, the function structure $S \rightarrow T$ is suitable when S belongs to the restricted class of *positive* relation structures, which we define in two steps as follows.

Definition 5.9. A candidate $\rho : \text{StrRel } S$ *acts on functions* when the following hold:

- (1) For any $f : X_0 \rightarrow X_1$, we have some $S f : S X_0 \rightarrow S X_1$, with $S (\text{IdRel } X) \equiv \text{idfun } (S X)$ for all $X : \text{Type}$.
- (2) For any $f : X_0 \rightarrow X_1, g : Y_0 \rightarrow Y_1$, relations $R_0 : X_0 \rightarrow Y_0 \rightarrow \text{Type}$ and $R_1 : X_1 \rightarrow Y_1 \rightarrow \text{Type}$, and $\alpha : (x : X_0) (y : Y_0) \rightarrow R_0 x y \rightarrow R_1 (f x) (g y)$, we have a map $\rho \alpha : (s : S X_0) (t : S Y_0) \rightarrow \rho R_0 s t \rightarrow \rho R_1 (S f s) (S g t)$.

LEMMA 5.10. *If $\rho : \text{StrRel } S$ is suitable and acts on functions, then any prop-valued equivalence relation $R : X \rightarrow X \rightarrow \text{Type}$ induces a map $(S X) / (\rho R) \rightarrow S (X / R)$.*

PROOF. We have $S \llbracket _ \rrbracket : S X \rightarrow S (X / R)$. The codomain is a set by set-preservation, so we just need to show that $S \llbracket _ \rrbracket x_0 \equiv S \llbracket _ \rrbracket x_1$ whenever $\rho R x_0 x_1$. To show this we use the uniqueness condition for descent to quotients: it suffices to show that both $S \llbracket _ \rrbracket x_0$ and $S \llbracket _ \rrbracket x_1$ satisfy $\rho (\text{graph } \llbracket _ \rrbracket) x_0$. For the first, we start by deriving $\rho R^{\leftarrow} x_0 x_0$ from $\rho R x_0 x_1$ using symmetry and transitivity. We have maps $\text{idFun } X : X \rightarrow X$ to $\llbracket _ \rrbracket : X \rightarrow X / R$ and know that $R^{\leftarrow} y y'$ implies $\text{graph } \llbracket _ \rrbracket y \llbracket y' \rrbracket$ for every $y, y' : X$. The action of ρ on functions thus tells us that $\rho (\text{graph } \llbracket _ \rrbracket) (S (\text{idFun } X) x_0) (S \llbracket _ \rrbracket x_0)$ and therefore $\rho (\text{graph } \llbracket _ \rrbracket) x_0 (S \llbracket _ \rrbracket x_0)$ holds. The proof of $\rho (\text{graph } \llbracket _ \rrbracket) x_0 (S \llbracket _ \rrbracket x_1)$ proceeds similarly. \square

Definition 5.11. Let $\rho : \text{StrRel } S$ be a suitable notion of relational structure that acts on functions. We say ρ is *positive* when the following hold.

- (1) *Reflexivity:* For any $X : \text{Type}$ and $s : S X$, we have $\rho (\text{IdRel } X) s s$.
- (2) *Reverse transitivity:* For any prop-valued relations $R : X \rightarrow Y \rightarrow \text{Type}$ and $R' : Y \rightarrow Z \rightarrow \text{Type}$ and terms $s : S X$ and $t : S Z$, the map $(\rho R \cdot \rho R') s t \rightarrow \rho (R \cdot R') s t$ supplied by transitivity is an equivalence.
- (3) *Quotients:* For any prop-valued equivalence relation R , the map $(S X) / (\rho R) \rightarrow S (X / R)$ defined in [Lemma 5.10](#) is an equivalence.

These conditions are designed to validate the following result.

THEOREM 5.12. *If $\rho_1 : \text{StrRel } S_1$ is a positive notion of relational structure and $\rho_2 : \text{StrRel } S_2$ is a suitable notion of relational structure, then the following is a suitable notion of relational structure.*

FunctionRelStr : $\text{StrRel } (\lambda X \rightarrow S_1 X \rightarrow S_2 X)$

FunctionRelStr $R f g = \forall \{x y\} \rightarrow \rho_1 R x y \rightarrow \rho_2 R (f x) (g y)$

PROOF. The (formalized) proof is rather technical; here, we give just enough detail to point out where the conditions in [Definition 5.11](#) are used. Reflexivity and the quotient condition are both used in the proof that **FunctionRelStr** induces quotient structures. Given an equivalence relation R , we must be able to turn maps $f : S_1 X \rightarrow S_2 X$ with **FunctionRelStr** $R f f$ into maps $S_1 (X / R) \rightarrow S_2 (X / R)$. First, we use reflexivity to construct a map $S_1 X \rightarrow S_2 (X / R)$. Given $s : S_1 X$, we have some $r : \rho_1 (\text{IdRel } X) s s$ by reflexivity. By the action of ρ_1 on functions and reflexivity of R , this implies $r : \rho_1 R s s$. Then by definition of **FunctionRelStr** $R f f$, we have $\rho_2 R (f s) (f s)$. By descent to quotients for ρ_2 , we thereby obtain an element of $S_2 (X / R)$.

The uniqueness condition in descent to quotients allows us to prove that this map $S_1 X \rightarrow S_2 (X / R)$ respects $\rho_1 R$, and thus induces a map $S_1 X / \rho_1 R \rightarrow S_2 (X / R)$. We then apply the quotient condition to derive a map $S_1 (X / R) \rightarrow S_2 (X / R)$. Finally, reverse transitivity in the domain is naturally required when we prove that **FunctionRelStr** is transitive. \square

We note that the definitions of suitability and positivity are somewhat negotiable, and are chosen to make these proofs work. For example, one may instead require suitable (not only positive) structures to be reflexive. In that case we are still able to prove the closure conditions; however, we must also impose a *reverse* reflexivity condition for positive structures.

Example 5.13. To apply [Theorem 5.7](#) to our multiset implementations, we must show that the QER R is structured as a relation between multiset implementations. We generate the notion of structured relation using a relational equivalent of the tactics described in [Section 3](#). (In this case, an annotation is required on the constant \mathbb{N} verifying that it is a set.)

multisetShape $X = X \times (A \rightarrow X \rightarrow X) \times (A \rightarrow X \rightarrow \text{Const}[\mathbb{N}, \text{isSet}\mathbb{N}])$

module $S = \text{RelMacro } (\text{autoRelDesc multisetShape})$

The module S defined here packages all the definitions and results we need to work with a relational structure: $S.\text{structure} : \text{Type} \rightarrow \text{Type}$, definitions $S.\text{equiv} : \text{StrEquiv } S.\text{structure}$ and $S.\text{relation} : \text{StrRel } S.\text{structure}$ of structured equivalences and relations respectively, and proofs they are univalent. Here, showing that R is structured requires the following three conditions:

isStructuredEmp : $R [] []$

isStructuredCount : $\forall x \{xs ys\} \rightarrow R xs ys \rightarrow \text{L.count } x xs \equiv \text{AL.count } x ys$

isStructuredInsert : $\forall x \{xs ys\} \rightarrow R xs ys \rightarrow R (\text{L.insert } x xs) (\text{AL.insert } x ys)$

The first holds by calculation, while the second holds by definition of R . The third boils down to showing that $\text{AL.count } a (\text{AL.insert } x ys) \equiv \text{addIfEq } a x 1 (\text{AL.count } a ys)$ for all $a : A$, which follows from a laborious but straightforward case analysis on ys and the outputs of $_ == _$. Defining $\text{Lstructure} = ([], \text{L.insert}, \text{L.count})$ and $\text{ALstructure} = ([], \text{AL.insert}, \text{AL.count})$, we then have $(\text{isStructuredEmp}, \text{isStructuredCount}, \text{isStructuredInsert}) : S.\text{relation } R \text{Lstructure } \text{ALstructure}$.

It remains only to turn the crank. Applying [Theorem 5.7](#), we immediately obtain multiset structures on $\text{List } A / R^{\leftarrow}$ and $\text{AssocList } A / R^{\rightarrow}$ and a path between the two in $\text{TypeWithStr } S.\text{structure}$. Note that we never need to explicitly show that the multiset operations for $\text{List } A$ and $\text{AssocList } A$ preserve R^{\leftarrow} and R^{\rightarrow} respectively; this is a formal consequence of the proof that R is structured.

We can see our queue example from Section 4.2 as another instance of this result, as the graph of `appendReverse : BatchedQueue A → ListQueue A` defines a QER. (The graph of any function is zigzag-complete.) The quotient of `BatchedQueue A` winds up being equivalent to our hand-rolled `BatchedQueueHIT`, while the quotient of `ListQueue A` is the original type. Once we have obtained the raw quotients, we can check that they satisfy any axioms we like; this usually reduces straightforwardly to showing that the axioms hold on the raw types up to R^{\leftarrow} or R^{\rightarrow} .

6 RELATED AND FUTURE WORK

In this paper, we have shown how to combine univalence and higher inductive types to obtain internal relational representation independence results in dependent type theory, specifically Cubical Agda. Univalence captures internally the principle that isomorphic types are equal; the structure identity principle lifts univalence to structured types and structured isomorphisms; and set quotients (via higher inductive types) allow us to improve structured relations to isomorphisms, in order to obtain representation independence results through the SIP, and validate axioms that a concrete implementation may otherwise fail to satisfy.

Our work lies at the intersection of formalized mathematics and programming language theory; we will navigate the related work moving from the former to the latter.

Proof reuse and transfer. As we discuss in Section 1, an important problem in modern proof assistants is the automatic transfer of programs and proofs across related mathematical structures. In the context of Coq, Magaud [2003]; Magaud and Bertot [2002] considered transferring proofs between isomorphic types by means of an external plugin, but their work did not handle dependent types. Cohen et al. [2013] designed a more general framework for CoqEAL, which supported program and data refinements for arbitrary relations. Their work was implemented using a parametricity plugin and proof search using typeclass instances, but does not handle dependently typed goals, and therefore could only transport programs and not proofs between related structures.

Tabareau et al. [2018, 2019] have recently improved this work by combining parametricity and typeclass based proof search à la CoqEAL with (axiomatic) univalence. This new Coq framework, called *univalent parametricity*, enables the transport of programs and proofs between related structures; however, unlike CoqEAL, its scope is limited to isomorphic types due to its reliance on univalence. We emphasize that while their work relies on an axiom for univalence, it nevertheless achieves code reuse in Coq without losing computational content, by means of a careful setup in which the axiom does not interfere with computationally-relevant parts, and by using typeclass instances to mimic computation rules similar to those of Cubical Type Theory. However, their approach does not achieve all the niceties of Cubical Type Theory, including constructive quotients and functional and propositional extensionality. Ringer et al. [2019] automatically build isomorphisms in Coq between inductive types and transport programs and proofs across these isomorphisms. This is accomplished using very similar techniques to those of Tabareau et al. [2018], but written as an external plugin instead of using typeclass instance search.

Another major contribution of Tabareau et al. [2019] is a solution to the *anticipation problem*, which concerns inferring an interface *a posteriori* and transporting programs and proofs across this inferred interface. This is especially useful in dependent type theory, where concrete types enjoy many definitional equalities that abstract types do not. For example, if one parameterizes proofs by an implementation of \mathbb{N} with `+`, the abstract `+` will not reduce on any input, whereas a concrete implementation will reduce on `0`. It is therefore very useful to allow users to develop libraries using concrete types while still enjoying the benefits of automated program and data refinements. In the setting of representation independence, however, one usually considers programs that *are* parameterized by an interface, and that interface determines the notion of structure-preservation, making

the anticipation problem less central. But even a partial solution to the anticipation problem in our setting would improve the applicability and ease of use of our framework.

Similar frameworks exist for other proof assistants. The Isabelle/HOL code generator uses refinements to make abstract code executable [Haftmann et al. 2013]; a similar refinement framework has been developed in Coq by Delaware et al. [2015] to synthesize abstract datatypes and generate OCaml code. There is also an external tool for Isabelle/HOL called Autoref which uses a parametricity-based translation to achieve CoqEAL-style refinements [Lammich 2013].

The structure identity principle. In the philosophy of mathematics, *structuralism* is an informal principle of representation independence: although mathematical objects may be realized in *ad hoc* ways, reasonable mathematical arguments implicitly access these structures only through interfaces [Awodey 2013; Benacerraf 1965]. As the work on proof transfer illustrates, formalized mathematics must often rely on this principle either implicitly or explicitly [Carette et al. 2014].

Researchers in HoTT/UF noticed very early on that univalence captures a form of structuralism, and have subsequently formalized quite a few versions of the SIP in HoTT/UF. The first type-theoretic formalization and proof of the SIP for algebraic structures was given by Coquand and Danielsson [2013]. As discussed in Section 3.2, their notions of structure are given by maps $S : \text{Type} \rightarrow \text{Type}$ with an action on equivalences. That action induces a notion of structured equivalence, which is then required to agree with paths.

Section 9.8 of the HoTT Book [Univalent Foundations Program 2013] contains a version of the SIP for structured categories; Ahrens and Lumsdaine [2017] develop a variation of this SIP using displayed categories. These principles are phrased in terms of structured *functions* between structured types, splitting the definition of a structure into two components: what it means to equip a type with structure, and what it means for a function to preserve that structure. The SIP then applies to structured functions that are also equivalences. This work is limited to structures on sets, as opposed to general types, due to the difficulty of formulating higher categories in HoTT/UF, but Ahrens et al. [2020] have recently developed a higher generalization.

Our SIP is derived from a version proposed by Escardó [2019]; like that of Coquand and Danielsson [2013], it is phrased in terms of equivalences rather than functions, but follows the categorical SIPs by allowing an arbitrary definition of structured equivalences rather than deriving one from an action on equivalences. Our `UnivalentStr` closely resembles Escardó’s *standard notion of structure*, except for being stated in terms of dependent paths; the two conditions are equivalent.

Our univalent structured relations follow the pattern of the categorical SIPs, with relations substituted for functions: a notion of structured relations is univalent when structured relations that are equivalences correspond to paths. However, our suitability conditions requires that even relations that are merely “equivalence-like” (i.e., are QERs) have some relationship with paths. This increased power comes with a more limited applicability, seen in the lack of a suitable notion of structured relation for general function types.

Internalizing relational parametricity. We have shown that univalence and HITs together recover a rich notion of representation independence, a concept which in programming languages is often obtained through parametricity theorems. We note that there are other “free theorems” which cannot be proven in Cubical Agda; a simple example is the contractibility of $(X : \text{Type}) \rightarrow X \rightarrow X$.

Just as univalence internalizes the fact that type theory respects isomorphism, one can consider dependent type theories with *internalized parametricity* [Abel et al. 2020; Bernardy et al. 2015; Cavallo and Harper 2020; Krishnaswami and Dreyer 2013; Nuyts and Devriese 2018; Nuyts et al. 2017]. Parametricity applies to representation independence in greater generality than univalence: for example, it implies that any computation parameterized by a `Queue` produces the same results for `List A` and `AssocList A`, while we only have such a result for the quotients described in Section 5.

Nevertheless, it is worth noting that internal parametricity is not *stronger* than univalence, but rather incomparable: an action on relations does not imply an action on equivalences.

From the standpoint of formalized mathematics, one drawback of internalized parametricity is that it is more difficult to combine with classical principles than univalence. For example, it is incompatible with excluded middle for h -propositions [Cavallo and Harper 2020, Theorem 10], unlike univalence [Kapulkin and Lumsdaine 2020]. Furthermore, there are not currently any large-scale proof assistants which implement type theories with internal parametricity.

Programming in HoTT/UF. We are not the first to consider applications of univalence and HITs to programming languages. Angiuli et al. [2016] model Darcs-style patch theories as HITs; HoTTSQL [Chu et al. 2017] uses a univalent universe to define the semantics of a query language and prove optimizations correct; and Basold et al. [2017] discuss HITs for types often used in programming, including modular arithmetic, integers and finite sets. These examples predate implementations of Cubical Type Theory, and would likely be significantly easier to formalize in Cubical Agda.

Countless variations of finite (multi)sets have been considered in all the main proof assistants for HoTT/UF; we limit our comparisons to those closest to our own. The finite sets of Basold et al. [2017] have been further studied by Frumin et al. [2018], and are defined by encoding finite subsets of A as the free join-semilattice on A . If one drops idempotency of the union operation, one obtains a HIT equivalent to $\text{List } A / R^{\leftarrow}$. They also discuss a variation called *listed finite sets* which are almost exactly $\text{List } A / R^{\leftarrow}$, except that they have a path constructor equating lists with duplicate elements. This HIT and a *listed finite multisets* version have been formalized in Cubical Agda by Choudhury and Fiore [2019]; we have proven their type equivalent to $\text{List } A / R^{\leftarrow}$ as well as to a direct HIT version of $\text{AssocList } A / R^{\rightarrow}$. Gylterud [2020] defines finite multisets using set quotients of W -types of sets. As an application, Forsberg et al. [2020] define in Cubical Agda an ordinal notation system using a variation of listed finite multisets, which is then proven equivalent to two other definitions of ordinal notation systems, and both programs and proofs are transported between the definitions. These transports are done in an ad-hoc way that can be done more elegantly with the SIP. Prior to HoTT/UF, Danielsson [2012] formalized finite multisets as a setoid of lists modulo “bag equivalence,” which is exactly the relation R that we use to quotient lists in Example 5.5.

Batched queues maintain the invariant that the first list is only empty if the queue is empty. One can encode this invariant in the type of the $Q(_, _)$ constructor, so that *enqueue* and *dequeue* are required to preserve the invariant. This variation of queues is equivalent to *BatchedQueueHIT* and has been formalized independently in Cubical Agda by Gjørup and Vindum [2019].

Abstract types. There are of course countless papers on abstract types and representation independence in programming languages, but we limit our discussion to the few most relevant. Most representation independence theorems apply to structured heterogeneous relations [Mitchell 1986], much like the ones we consider. In the context of System F, Robinson [1994] uses invariance under isomorphism to derive a simple form of representation independence for abstract types.

Extended ML [Kahrs et al. 1997] is an extension of Standard ML whose signatures can include axioms drawn from boolean expressions extended with quantifiers and termination predicates; the technique of *algebraic specification* [Sannella and Tarlecki 1987] is a method for establishing such axioms by successive structure-preserving refinements, much like CoqEAL. Miranda’s *data types with laws* extend datatype constructors with a limited ability to maintain invariants [Turner 1985]; as with the normal-form representations of quotients discussed in Example 2.2, these, unlike HITs, have the drawback of incurring runtime cost to maintain normal forms.

Future work. A natural next direction is to formalize representation independence for more sophisticated data structures, such as self-balancing binary search trees. Our methodology already

applies to these and other examples, but proving the correspondences suitable would be more complex. We can also extend the structure tactics described in Sections 3 and 5, particularly to support inductive types beyond `Maybe`. Note, however, that these tactics apply to the types that appear in *interfaces*—not in *representations*—which are usually quite limited.

Finally, we would like to investigate whether our techniques are practical in Coq using axiomatic univalence and HITs. Although we rely heavily on these features computing automatically in Cubical Agda, Tabareau et al. [2019] overcome similar difficulties with the univalence axiom by cleverly using typeclass instances to mimic computation rules.

ACKNOWLEDGMENTS

We thank Robert Harper and Jonathan Sterling for helpful conversations about this work.

The material in this paper is based upon research supported by the Air Force Office of Scientific Research under MURI grants FA9550-15-1-0053 and FA9550-19-1-0216, and the Swedish Research Council (Vetenskapsrådet) under Grant No. 2019-04545. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government, or any other entity.

REFERENCES

- Andreas Abel, Jesper Cockx, Dominique Devriese, Amin Timany, and Philip Wadler. 2020. Leibniz equality is isomorphic to Martin-Löf identity, parametrically. *Journal of Functional Programming* 30 (2020), e17. <https://doi.org/10.1017/S0956796820000155>
- Benedikt Ahrens and Peter LeFanu Lumsdaine. 2017. Displayed Categories. In *2nd International Conference on Formal Structures for Computation and Deduction (FSCD 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 84)*, Dale Miller (Ed.), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 5:1–5:16. <https://doi.org/10.4230/LIPIcs.FSCD.2017.5>
- Benedikt Ahrens, Paige Randall North, Michael Shulman, and Dimitris Tsementzis. 2020. A Higher Structure Identity Principle. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (Saarbrücken, Germany) (LICS '20)*. Association for Computing Machinery, New York, NY, USA, 53–66. <https://doi.org/10.1145/3373718.3394755>
- Abhishek Anand and Greg Morrisett. 2017. *Revisiting Parametricity: Inductives and Uniformity of Propositions*. arXiv:1705.01163 [cs.LO] <http://arxiv.org/abs/1705.01163> Preprint.
- Carlo Angiuli, Guillaume Brunerie, Thierry Coquand, Kuen-Bang Hou (Favonia), Robert Harper, and Daniel R. Licata. 2019. Syntax and Models of Cartesian Cubical Type Theory. (February 2019). <https://github.com/dlicata335/cart-cube> Preprint.
- Carlo Angiuli, Kuen-Bang Hou (Favonia), and Robert Harper. 2018. Cartesian Cubical Computational Type Theory: Constructive Reasoning with Paths and Equalities. In *27th EACSL Annual Conference on Computer Science Logic (CSL 2018) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 119)*, Dan Ghica and Achim Jung (Eds.), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 6:1–6:17. <https://doi.org/10.4230/LIPIcs.CSL.2018.6>
- Carlo Angiuli, Edward Morehouse, Daniel R. Licata, and Robert Harper. 2016. Homotopical patch theory. *Journal of Functional Programming* 26 (2016). <https://doi.org/10.1017/S0956796816000198> Special issue dedicated to ICFP 2014.
- Steve Awodey. 2013. Structuralism, Invariance, and Univalence. *Philosophia Mathematica* 22, 1 (10 2013), 1–11. <https://doi.org/10.1093/phimat/nkt030>
- Gilles Barthe, Venanzio Capretta, and Olivier Pons. 2003. Setoids in type theory. *Journal of Functional Programming* 13, 2 (2003), 261–293. <https://doi.org/10.1017/S0956796802004501>
- Henning Basold, Herman Geuvers, and Niels van der Weide. 2017. Higher Inductive Types in Programming. *Journal of Universal Computer Science* 23, 1 (January 2017), 63–88. <https://doi.org/10.3217/jucs-023-01-0063>
- Andrej Bauer, Jason Gross, Peter LeFanu Lumsdaine, Michael Shulman, Matthieu Sozeau, and Bas Spitters. 2017. The HoTT Library: A Formalization of Homotopy Type Theory in Coq. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs (Paris, France) (CPP 2017)*. ACM, New York, NY, USA, 164–172. <https://doi.org/10.1145/3018610.3018615>
- Paul Benacerraf. 1965. What Numbers Could not Be. *The Philosophical Review* 74, 1 (1965), 47–73. <https://doi.org/10.2307/2183530>

- Jean-Philippe Bernardy, Thierry Coquand, and Guilhem Moulin. 2015. A Presheaf Model of Parametric Type Theory. *Electronic Notes in Theoretical Computer Science* 319 (2015), 67–82. <https://doi.org/10.1016/j.entcs.2015.12.006> 31st Conference on the Mathematical Foundations of Programming Semantics.
- Jean-Philippe Bernardy, Patrik Jansson, and Ross Paterson. 2012. Proofs for Free: Parametricity for Dependent Types. *Journal of Functional Programming* 22, 2 (March 2012), 107–152. <https://doi.org/10.1017/S0956796812000056>
- Simon Boulrier, Pierre-Marie Pédro, and Nicolas Tabareau. 2017. The next 700 Syntactical Models of Type Theory. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs (Paris, France) (CPP 2017)*. Association for Computing Machinery, New York, NY, USA, 182–194. <https://doi.org/10.1145/3018610.3018620>
- Edwin Brady. 2013. Idris, a general-purpose dependently typed programming language: Design and implementation. *Journal of Functional Programming* 23, 5 (2013), 552–593. <https://doi.org/10.1017/S095679681300018X>
- Guillaume Brunerie, Kuen-Bang Hou (Favonia), Evan Cavallo, Tim Baumann, Eric Finster, Jesper Cockx, Christian Sattler, Chris Jeris, Michael Shulman, et al. 2018. Homotopy Type Theory in Agda. <https://github.com/HoTT/HoTT-Agda>
- Jacques Carette, William M. Farmer, and Michael Kohlhase. 2014. Realms: A Structure for Consolidating Knowledge about Mathematical Theories. In *Intelligent Computer Mathematics*, Stephen M. Watt, James H. Davenport, Alan P. Sexton, Petr Sojka, and Josef Urban (Eds.). Springer International Publishing, Cham, 252–266. https://doi.org/10.1007/978-3-319-08434-3_19
- Evan Cavallo and Robert Harper. 2019. Higher Inductive Types in Cubical Computational Type Theory. *Proceedings of the ACM on Programming Languages* 3, POPL, Article 1 (January 2019), 27 pages. <https://doi.org/10.1145/3290314>
- Evan Cavallo and Robert Harper. 2020. Internal Parametricity for Cubical Type Theory. In *28th EACSL Annual Conference on Computer Science Logic (CSL 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 152)*, Maribel Fernández and Anca Muscholl (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 13:1–13:17. <https://doi.org/10.4230/LIPIcs.CSL.2020.13>
- Vikraman Choudhury and Marcelo Fiore. 2019. The finite-multiset construction in HoTT. <https://hott.github.io/HoTT-2019/conf-slides/Choudhury.pdf> Talk at the International Conference on Homotopy Type Theory (HoTT 2019).
- Shumo Chu, Konstantin Weitz, Alvin Cheung, and Dan Suciu. 2017. HoTTSQL: Proving Query Rewrites with Univalent SQL Semantics. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (Barcelona, Spain) (PLDI 2017)*. ACM, New York, NY, USA, 510–524. <https://doi.org/10.1145/3062341.3062348>
- Frédéric Chyzak, Assia Mahboubi, Thomas Sibut-Pinote, and Enrico Tassi. 2014. A Computer-Algebra-Based Formal Proof of the Irrationality of $\zeta(3)$. In *Interactive Theorem Proving (Vienna, Austria) (ITP 2014)*, Gerwin Klein and Ruben Gamboa (Eds.). Springer International Publishing, Cham, 160–176. https://doi.org/10.1007/978-3-319-08970-6_11
- Cyril Cohen, Thierry Coquand, Simon Huber, and Anders Mörtberg. 2018. Cubical Type Theory: A Constructive Interpretation of the Univalence Axiom. In *21st International Conference on Types for Proofs and Programs (TYPES 2015) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 69)*, Tarmo Uustalu (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 5:1–5:34. <https://doi.org/10.4230/LIPIcs.TYPES.2015.5>
- Cyril Cohen, Maxime Dénès, and Anders Mörtberg. 2013. Refinements for Free!. In *Certified Programs and Proofs (CPP 2013)*, Georges Gonthier and Michael Norrish (Eds.). Springer International Publishing, Cham, 147–162. https://doi.org/10.1007/978-3-319-03545-1_10
- Thierry Coquand and Nils Anders Danielsson. 2013. Isomorphism is equality. *Indagationes Mathematicae* 24, 4 (2013), 1105–1120. <https://doi.org/10.1016/j.indag.2013.09.002> In memory of N.G. (Dick) de Bruijn (1918–2012).
- Thierry Coquand, Simon Huber, and Anders Mörtberg. 2018. On Higher Inductive Types in Cubical Type Theory. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (Oxford, United Kingdom) (LICS 2018)*. ACM, New York, NY, USA, 255–264. <https://doi.org/10.1145/3209108.3209197>
- Karl Cray. 2017. Modules, Abstraction, and Parametric Polymorphism. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (Paris, France) (POPL 2017)*. ACM, New York, NY, USA, 100–113. <https://doi.org/10.1145/3009837.3009892>
- Nils Anders Danielsson. 2012. Bag Equivalence via a Proof-Relevant Membership Relation. In *Interactive Theorem Proving (Princeton, NJ, USA) (ITP 2012)*, Lennart Beringer and Amy Felty (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 149–165. https://doi.org/10.1007/978-3-642-32347-8_11
- Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. 2015. The Lean Theorem Prover (System Description). In *Automated Deduction – CADE-25 (Berlin, Germany)*, Amy P. Felty and Aart Middeldorp (Eds.). Springer International Publishing, Cham, 378–388. https://doi.org/10.1007/978-3-319-21401-6_26
- Benjamin Delaware, Clément Pit-Claudel, Jason Gross, and Adam Chlipala. 2015. Fiat: Deductive Synthesis of Abstract Data Types in a Proof Assistant. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (Mumbai, India) (POPL 2015)*. Association for Computing Machinery, New York, NY, USA, 689–700. <https://doi.org/10.1145/2676726.2677006>
- Edsger W. Dijkstra. 1974. On the role of scientific thought. (August 1974). <https://www.cs.utexas.edu/users/EWD/transcriptions/EWD04xx/EWD447.html>

- Martín Hötzel Escardó. 2019. Introduction to univalent foundations of mathematics with Agda. <https://www.cs.bham.ac.uk/~mhe/HoTT-UF-in-Agda-Lecture-Notes/index.html>
- Fredrik Nordvall Forsberg, Chuangjie Xu, and Neil Ghani. 2020. Three Equivalent Ordinal Notation Systems in Cubical Agda. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs* (New Orleans, LA, USA) (CPP 2020). Association for Computing Machinery, New York, NY, USA, 172–185. <https://doi.org/10.1145/3372885.3373835>
- Dan Frumin, Herman Geuvers, Léon Gondelman, and Niels van der Weide. 2018. Finite Sets in Homotopy Type Theory. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs* (Los Angeles, CA, USA) (CPP 2018). Association for Computing Machinery, New York, NY, USA, 201–214. <https://doi.org/10.1145/3167085>
- Emil Gjørup and Simon Friis Vindum. 2019. Case Study: BatchedQueue. <https://github.com/limehloh/CubicalAgda-BatchedQueue>
- Håkon Robbestad Gylterud. 2020. Multisets in type theory. *Mathematical Proceedings of the Cambridge Philosophical Society* 169, 1 (2020), 1–18. <https://doi.org/10.1017/S0305004119000045>
- Florian Haftmann, Alexander Krauss, Ondřej Kunčar, and Tobias Nipkow. 2013. Data Refinement in Isabelle/HOL. In *Interactive Theorem Proving* (Rennes, France) (ITP 2013), Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 100–115. https://doi.org/10.1007/978-3-642-39634-2_10
- Michael Hedberg. 1998. A coherence theorem for Martin-Löf’s type theory. *Journal of Functional Programming* 8, 4 (1998), 413–436. <https://doi.org/10.1017/S0956796898003153>
- Stefan Kahrs, Donald Sannella, and Andrzej Tarlecki. 1997. The definition of Extended ML: A gentle introduction. *Theoretical Computer Science* 173, 2 (1997), 445–484. [https://doi.org/10.1016/S0304-3975\(96\)00163-6](https://doi.org/10.1016/S0304-3975(96)00163-6)
- Chris Kapullkin and Peter LeFanu Lumsdaine. 2020. The law of excluded middle in the simplicial model of type theory. arXiv:2006.13694 [math.LO] <https://arxiv.org/abs/2006.13694> Preprint.
- Chantal Keller and Marc Lasson. 2012. Parametricity in an Impredicative Sort. In *21st EACSL Annual Conference on Computer Science Logic (CSL 2012) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 16)*, Patrick Cégielski and Arnaud Durand (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 381–395. <https://doi.org/10.4230/LIPIcs.CSL.2012.381>
- Nicolai Kraus, Martín Escardó, Thierry Coquand, and Thorsten Altenkirch. 2017. Notions of Anonymous Existence in Martin-Löf Type Theory. *Logical Methods in Computer Science* 13, 1 (2017). [https://doi.org/10.23638/LMCS-13\(1:15\)2017](https://doi.org/10.23638/LMCS-13(1:15)2017)
- Neelakantan R. Krishnaswami and Derek Dreyer. 2013. Internalizing Relational Parametricity in the Extensional Calculus of Constructions. In *Computer Science Logic 2013 (CSL 2013) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 23)*, Simona Ronchi Della Rocca (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 432–451. <https://doi.org/10.4230/LIPIcs.CSL.2013.432>
- Peter Lammich. 2013. Automatic Data Refinement. In *Interactive Theorem Proving* (Rennes, France) (ITP 2013), Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 84–99. https://doi.org/10.1007/978-3-642-39634-2_9
- Xavier Leroy. 1995. Applicative Functors and Fully Transparent Higher-Order Modules. In *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California) (POPL 1995). ACM, New York, NY, USA, 142–153. <https://doi.org/10.1145/199448.199476>
- Dan Licata. 2016. *Weak univalence with “beta” implies full univalence.* <https://groups.google.com/d/msg/homotopytypetheory/j2KBIvDw53s/YTDK4D0NFQAJ> Email to Homotopy Type Theory mailing list.
- Peter LeFanu Lumsdaine and Michael Shulman. 2019. Semantics of higher inductive types. *Mathematical Proceedings of the Cambridge Philosophical Society* (2019). <https://doi.org/10.1017/S030500411900015X>
- Nicolas Magaud. 2003. Changing Data Representation within the Coq System. In *Theorem Proving in Higher Order Logics* (Rome, Italy) (TPHOLs 2003), David Basin and Burkhart Wolff (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 87–102. https://doi.org/10.1007/10930755_6
- Nicolas Magaud and Yves Bertot. 2002. Changing Data Structures in Type Theory: A Study of Natural Numbers. In *Types for Proofs and Programs (TYPES 2000) (Lecture Notes in Computer Science, Vol. 2277)*, Paul Callaghan, Zhaohui Luo, James McKinna, and Robert Pollack (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 181–196. https://doi.org/10.1007/3-540-45842-5_12
- Per Martin-Löf. 1975. An Intuitionistic Theory of Types: Predicative Part. In *Logic Colloquium ’73*, H. E. Rose and J. C. Shepherdson (Eds.). Studies in Logic and the Foundations of Mathematics, Vol. 80. North-Holland, 73–118. [https://doi.org/10.1016/S0049-237X\(08\)71945-1](https://doi.org/10.1016/S0049-237X(08)71945-1)
- John C. Mitchell. 1986. Representation Independence and Data Abstraction. In *Proceedings of the 13th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages* (St. Petersburg Beach, Florida) (POPL ’86). Association for Computing Machinery, New York, NY, USA, 263–276. <https://doi.org/10.1145/512644.512669>
- Andreas Nuyts and Dominique Devriese. 2018. Degrees of Relatedness: A Unified Framework for Parametricity, Irrelevance, Ad Hoc Polymorphism, Intersections, Unions and Algebra in Dependent Type Theory. In *Proceedings of the 33rd Annual*

- ACM/IEEE Symposium on Logic in Computer Science* (Oxford, United Kingdom) (*LICS 2018*). ACM, New York, NY, USA, 779–788. <https://doi.org/10.1145/3209108.3209119>
- Andreas Nuyts, Andrea Vezzosi, and Dominique Devriese. 2017. Parametric Quantifiers for Dependent Type Theory. *Proceedings of the ACM on Programming Languages* 1, ICFP, Article 32 (January 2017), 29 pages. <https://doi.org/10.1145/3110276>
- Chris Okasaki. 1999. *Purely functional data structures*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511530104>
- John C. Reynolds. 1983. Types, Abstraction and Parametric Polymorphism. In *Information Processing '83: Proceedings of the IFIP 9th World Computer Congress*, R. E. A. Mason (Ed.). North-Holland, 513–523.
- Egbert Rijke. 2012. *Homotopy Type Theory*. Master's thesis. University of Ljubljana. <http://hottheory.files.wordpress.com/2012/08/hott2.pdf>
- Talia Ringer, Nathaniel Yazdani, John Leo, and Dan Grossman. 2019. Ornaments for Proof Reuse in Coq. In *10th International Conference on Interactive Theorem Proving (ITP 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 141)*, John Harrison, John O'Leary, and Andrew Tolmach (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 26:1–26:19. <https://doi.org/10.4230/LIPIcs.ITP.2019.26>
- Edmund Robinson. 1994. Parametricity as isomorphism. *Theoretical Computer Science* 136, 1 (1994), 163–181. [https://doi.org/10.1016/0304-3975\(94\)00126-4](https://doi.org/10.1016/0304-3975(94)00126-4)
- Donald Sannella and Andrzej Tarlecki. 1987. On observational equivalence and algebraic specification. *J. Comput. System Sci.* 34, 2 (1987), 150–178. [https://doi.org/10.1016/0022-0000\(87\)90023-7](https://doi.org/10.1016/0022-0000(87)90023-7)
- Thomas Streicher. 1993. *Investigations Into Intensional Type Theory*. Habilitation thesis. Ludwig-Maximilians-Universität München. <https://www2.mathematik.tu-darmstadt.de/~streicher/HabilStreicher.pdf>
- Nicolas Tabareau, Éric Tanter, and Matthieu Sozeau. 2018. Equivalences for Free: Univalent Parametricity for Effective Transport. *Proceedings of the ACM on Programming Languages* 2, ICFP (September 2018), 92:1–92:29. <https://doi.org/10.1145/3236787>
- Nicolas Tabareau, Éric Tanter, and Matthieu Sozeau. 2019. *The Marriage of Univalence and Parametricity*. arXiv:1909.05027 [cs.PL] <https://arxiv.org/abs/1909.05027> Preprint.
- The Agda Development Team. 2020. The Agda Programming Language. <http://wiki.portal.chalmers.se/agda/pmwiki.php>
- The Coq Development Team. 2020. The Coq Proof Assistant. <https://www.coq.inria.fr>
- The Mathematical Components Team. 2020. The Mathematical Components library. <https://github.com/math-comp/math-comp>
- D. A. Turner. 1985. Miranda: A non-strict functional language with polymorphic types. In *Functional Programming Languages and Computer Architecture (FPCA 1985)*, Jean-Pierre Jouannaud (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–16.
- The Univalent Foundations Program. 2013. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Self-published. <https://homotopytypetheory.org/book/>
- Floris van Doorn, Jakob von Raumer, and Ulrik Buchholtz. 2017. Homotopy Type Theory in Lean. In *Interactive Theorem Proving (Brasília, Brazil) (ITP 2017)*, Mauricio Ayala-Rincón and César A. Muñoz (Eds.). Springer, Cham, 479–495. https://doi.org/10.1007/978-3-319-66107-0_30
- Andrea Vezzosi, Anders Mörtberg, and Andreas Abel. 2019. Cubical Agda: A Dependently Typed Programming Language with Univalence and Higher Inductive Types. *Proceedings of the ACM on Programming Languages* 3, ICFP, Article 87 (August 2019), 29 pages. <https://doi.org/10.1145/3341691>
- Vladimir Voevodsky. 2010a. The equivalence axiom and univalent models of type theory. http://www.math.ias.edu/vladimir/files/CMU_talk.pdf Notes from a talk at Carnegie Mellon University.
- Vladimir Voevodsky. 2010b. Univalent Foundations. https://www.math.ias.edu/vladimir/sites/math.ias.edu.vladimir/files/Bonn_talk.pdf Notes from a talk in Bonn.
- Vladimir Voevodsky. 2015. An experimental library of formalized Mathematics based on the univalent foundations. *Mathematical Structures in Computer Science* 25, 5 (2015), 1278–1294. <https://doi.org/10.1017/S0960129514000577>
- Vladimir Voevodsky, Benedikt Ahrens, Daniel Grayson, et al. 2020. UniMath – a computer-checked library of univalent mathematics. <https://github.com/UniMath/UniMath>
- Philip Wadler. 1989. Theorems for Free!. In *Proceedings of the Fourth International Conference on Functional Programming Languages and Computer Architecture* (Imperial College, London, United Kingdom) (*FPCA '89*). Association for Computing Machinery, New York, NY, USA, 347–359. <https://doi.org/10.1145/99370.99404>
- James Wood. 2019. Vectors and Matrices in Agda. Blog post at <https://personal.cis.strath.ac.uk/james.wood.100/blog/html/VecMat.html>.