Proof Repair and (Co)algebra Equivalences

E.G. PLANTE III, University of Illinois at Urbana-Champaign, USA
ERIC PAUL, University of Illinois at Urbana-Champaign, USA
COSMO VIOLA, University of Illinois at Urbana-Champaign, USA
G.A. KAVVOS, University of Bristol, United Kingdom
TALIA RINGER, University of Illinois at Urbana-Champaign, USA

Maintaining formal proofs in response to breaking changes in their dependencies is a major challenge of proof engineering. *Proof repair* is a kind of proof automation that lessens the burden of this challenge by automatically adapting proofs in response to those changes. While proof repair is well understood through an algorithmic lens, it has, until now, been poorly understood through a theoretical lens. We address this gap by giving the first foundational treatment of proof repair, focusing in particular on type-theoretic proof repair in response to changes in datatypes. Our key insight is that the core component of a proof repair algorithm—a "configuration"—is exactly an algebra isomorphism. We formalize and prove this insight, answering open questions from prior work. We then take this one step further and show that, by taking the dual of this representation, we obtain proof repair over coinductive types. All of our results are formalized in Agda.

CCS Concepts: • **Theory of computation → Type theory**; • **Software and its engineering → Software verification**.

Additional Key Words and Phrases: Proof Repair, Type Theory, Category Theory

1 Introduction

Formal proof developments in proof assistants like Agda, Rocq, and Lean [3, 38, 40] have grown larger as the technology has grown more accessible. One timeless challenge in large formal proof developments is the burden of change [47]. Proof repair has emerged as a promising tool to cope with change by automatically fixing formal proofs in response to breaking changes.

This paper is about a class of tools for coping with changes in datatypes in proof developments that we refer to as *type-theoretic proof repair* [46, 48, 58]: given an old and a new version of a datatype, these tools automatically rewrite the programs and proofs stated over the old version into programs and proofs stated over the new one. Type-theoretic proof repair leverages the underlying type theory of a proof assistant to establish relationships between datatypes across which programs and proofs can be repaired. This is, of course, not the only kind of proof repair—other repair tools exist that, for example, repair proofs concerning higher-order imperative code [25], or leverage non-type-theoretic methods like large language models [34] (see Section 5 for more).

The objective of this paper is to provide a mathematical foundation for type-theoretic proof repair tools that clarifies the scope of changes contemporary tools can handle and the extensions future repair tools can support. To this end, we develop the connection between proof repair and the categorical notion of algebras on endofunctors, from which results concerning the limits, scope, and extensions of repair follow directly. Our treatment characterizes the objects at the heart of repair tools; the syntactic transformations the tools perform over terms are beyond our present scope.

The principal theme of our work is that **proof repair is an algebra isomorphism**. We bring this conclusion to light by relating prior proof repair work to the type-theoretic understanding of datatypes; the goal is to connect these two fields and leverage type theory to obtain formal results about proof repair tools. Once established, the connection between algebra and repair gives

Authors' Contact Information: E.G. Plante III, University of Illinois at Urbana-Champaign, USA; Eric Paul, University of Illinois at Urbana-Champaign, USA; Cosmo Viola, University of Illinois at Urbana-Champaign, USA; G.A. Kavvos, University of Bristol, United Kingdom; Talia Ringer, University of Illinois at Urbana-Champaign, USA.

a mathematical explanation for why repair works, settles questions left open in prior work, and relaxes assumptions repair was thought to require. It also points beyond current tools, to repair for types no existing tool can handle. All of our results and examples are marked with a blue checkmark linking to the corresponding Agda formalization.

Contributions. Concretely, our contributions are:

- We give a formal account of proof repair within type theory, which lays the groundwork for its connection to algebras (Section 2). In doing so, we formally establish results about the fundamental constructs of repair tools that prior work had only conjectured informally.
- We demonstrate that proof repair can be understood through the notion of initial algebras on a suitable functor (Section 3). This new viewpoint clarifies the class of relationships between types that configuration-based repair captures, characterizing them exactly in terms of universal properties. We recast a case study from the original work on proof repair in terms of algebras, and formalize the connection to proof repair algorithms. Our main result asserts that configurations are algebra isomorphisms.
- We find that repair does not require univalence, as was previously thought, and so applies to type theories that are inconsistent with univalence. This opens the door for new implementations of repair in proof assistants like Lean (Section 3.4).
- With our new categorical perspective on proof repair, we develop the foundations for repair algorithms that handle a new class of types: coinductive types (Section 4).

Motivation

We introduce proof repair using a classic case study ✓. Suppose that a proof engineer has previously developed a verified library using the `Nat` datatype. They want to change the library to use a more efficient binary representation, `Bin`:

```
data Nat : Type where
  zero : Nat
  succ : Nat -> Nat
```

```
data Pos : Type where
 one : Pos
 pb0 : Pos -> Pos
 pb1 : Pos -> Pos

data Bin : Type where
 bzero : Bin
 bpos  : Pos -> Bin
```

The `Bin` datatype has a constructor for zero, `bzero`, and a constructor for positive numbers, `bpos`, which according to the `Pos` constructors is either 1, $2n$, or $2n + 1$ for some positive number n, itself represented by a term of type `Pos`. Proof repair is the study of algorithms that replace one of these representations with the other. To understand what such an algorithm must accomplish, consider the following definitions:

```
isZeroN : Nat -> Bool
isZeroN zero   = true
isZeroN (succ _) = false
```

```
zUnit : (n : Nat) -> n + zero = n
zUnit zero   = refl
zUnit (succ m) = ap succ (zUnit m)
```

While making a datatype change in a verified library, a proof engineer may encounter many such programs and proofs, requiring broad changes to their development. A simple way of achieving that is to compose each old term with a conversion function `toNat : Bin -> Nat`:

```
isZeroB-nat : Bin -> Bool
isZeroB-nat b = isZeroN (toNat b)
```

```
zUnitB-nat : (b : Bin) -> toNat b + zero = toNat b
zUnitB-nat b = zUnit (toNat b)
```

However, this output is unsatisfying because the transferred terms still route every operation back through `Nat` via `toNat`, so the old representation remains a dependency of the library.

Repair algorithms aim to do better. When given a *configuration* (Section 2), i.e., a decomposition of a type equivalence into a correspondence between the constructors and eliminators of `Nat` and `Bin`, the repair algorithm specializes the old code to the new datatype, producing native definitions over `Bin` that never mention `Nat`:

```
isZeroB : Bin -> Bool
isZeroB bzero   = true
isZeroB (bpos _) = false
```

```
zUnitB : (b : Bin) -> b + bzero = b
zUnitB bzero   = refl
zUnitB (bpos p) = ap bpos (pUnit p)
```

This algorithmic view of proof repair is well understood [48]. Before developing the algebra perspective, we give a formal account of proof repair, both to connect it to our abstract development and to establish results about the mechanics of repair.

Type Theory, Notation, and Conventions

We work within an intensional Martin-Löf type theory with dependent products, dependent sums, and identity types, using Agda as a proxy for this type theory. Our development uses the homotopy-theoretic reading of the identity type [56], but assumes no univalence principles. The only axioms we assume on the inductive side are function extensionality together with its computation rule on happly and the corresponding uniqueness rule (every path of functions is the extensionality of its happly). For the coinductive development (Section 4) we work in Cubical Agda [57], whose path type supports the copattern-style reasoning that coinductive types demand; there, function extensionality is a theorem, and again we make no use of univalence (or of other cubical-specific axioms beyond the path type itself).

We write $=$ for the identity (path) type, $p \bullet q : a = c$ for the composition of paths $p : a = b$ and $q : b = c$, $\mathrm{tpt}\, P\, p : P\, a \to P\, b$ for transport along $p : a = b$, and $=_{\mathrm{def}}$ for definitions. We refer to a term of type Type as a *carrier*. We write $(a : A) \to B\, a$ for dependent functions and freely use $\to$ for its non-dependent instance, with $g \circ f : A \to C$ the composite of functions $f : A \to B$ and $g : B \to C$. We write $(a : A) \times B\, a$ for dependent sums and freely use $\times$ for its non-dependent instance, with projections fst and snd. We write $\mathbb{1}$ for the unit type and $* : \mathbb{1}$ for its inhabitant. A *type equivalence* $f : A \simeq B$ is a bi-invertible map: a function $f : A \to B$ together with a left inverse and a right inverse. A type is *contractible* when it has a term to which every term is equal; it is a *proposition* when any two of its terms are equal, and a *set* when its identity types are all propositions. We suppress universe levels throughout; the formalization is universe-polymorphic. Finally, when one equation depends on another we display it as a plain equality rather than writing the transport or dependent path explicitly; the formal Agda statements carry the dependency in full.

2 Proof Repair via Configurations

In this section, we internalize the core machinery of repair algorithms, *configurations*, in type theory. The formalization of proof repair serves as the nexus of the connection we draw between repair and (co)algebras, and simultaneously provides a source of the questions the rest of the paper seeks to answer. The results developed here about configurations are primarily confirmations of facts that were thought to be true of configurations, but were not previously made precise or proven in a formal setting. We begin by informally working through the execution of a proof repair algorithm, and then give a formal definition of configurations.

The prior work on tools for type-theoretic repair (PUMPKIN Pi [48, 58]) develops the notion of a configuration to capture a decomposed version of an equivalence $f : A \simeq B$ at the level of the metalanguage. The repair tool then operates on the syntax of the type theory, replacing terms of type A with terms of type B according to this specification. A configuration for the motivating example of Section 1 names each type's constructors by `DepConstr` and its dependent eliminator

by DepElim, whose computation rules are Iota.[1] Together these present Nat and Bin through a common interface that pairs each Nat constructor (resp. eliminator) with the Bin constructor (resp. eliminator) it corresponds to under the equivalence, with constructors numbered by the metalanguage indices 0 and 1. Since Bin has no successor constructor, the configuration pairs succ with the successor function binSucc : Bin -> Bin; the constructors a configuration names need not be the datatype's own.

```
-- DepConstr: constructors, indexed by type and constructor number
DepConstr Nat 0 = zero   DepConstr Bin 0 = bzero
DepConstr Nat 1 = succ   DepConstr Bin 1 = binSucc

-- DepElim: a dependent eliminator for each type
DepElim Nat : (P : Nat -> Type) -> P (DepConstr Nat 0) -> ((n : Nat) -> P n
            -> P (DepConstr Nat 1 n)) -> (n : Nat) -> P n
DepElim Bin : (P : Bin -> Type) -> P (DepConstr Bin 0) -> ((b : Bin) -> P b
            -> P (DepConstr Bin 1 b)) -> (b : Bin) -> P b

-- Iota: the eliminator computes on each constructor
Iota Nat 0 : DepElim Nat P pz ps (DepConstr Nat 0) = pz
Iota Nat 1 : DepElim Nat P pz ps (DepConstr Nat 1 n) = ps n (DepElim Nat P pz ps n)
Iota Bin 0 : DepElim Bin P pz ps (DepConstr Bin 0) = pz
Iota Bin 1 : DepElim Bin P pz ps (DepConstr Bin 1 b) = ps b (DepElim Bin P pz ps b)
```

Proof repair tools then use the decomposition of the source and target types' inductive structure to automatically repair code by replacing the original constructors and eliminators componentwise with the new ones. The repair algorithm replaces every DepElim Nat by DepElim Bin, DepConstr Nat by DepConstr Bin, and Iota Nat by Iota Bin, while the remaining case data is carried over unchanged. For example, the proof zUnit given below would be repaired into zUnitB:

```
zUnit : (n : Nat) -> n + zero = n
zUnit = DepElim Nat (λn -> n + zero = n)
  (Iota Nat 0)
  (λn ih -> Iota Nat 1 • ap succ ih)
```

```
zUnitB : (b : Bin) -> b + bzero = b
zUnitB = DepElim Bin (λb -> b + bzero = b)
  (Iota Bin 0)
  (λb ih -> Iota Bin 1 • ap binSucc ih)
```

The basis of a configuration lies in the constructors of a type, which may take recursive and non-recursive arguments. We specify these constructors by a *signature*, and call an interpretation of a signature on a carrier a *constructor algebra*. Our definitions here are inspired by the work of Dybjer [22] and Chapman et al. [14], both of which proved crucial in formalizing the previously informal notion of configurations.

Definition 2.1 (Constructor Algebra ✓)**.** To describe the argument list of a single constructor as a telescope, we define a *constructor arity* type generated by:

$$\begin{aligned}\mathsf{Done} &: \mathsf{ConstrArity}\\ \mathsf{Nonrec} &: (A : \mathsf{Type}) \to (A \to \mathsf{ConstrArity}) \to \mathsf{ConstrArity}\\ \mathsf{Rec} &: (I : \mathsf{Type}) \to \mathsf{ConstrArity} \to \mathsf{ConstrArity}.\end{aligned}$$

Here Done closes the telescope, Nonrec $A\,k$ records a non-recursive argument of type A on which the remaining arity k may depend, and Rec $I\,cs$ records a batch of recursive arguments, one for each element of I. Non-recursive and recursive arguments may be freely interleaved.

[1] We omit the Eta data that prior work included because our formalization shows it is redundant.

We define the type of *arguments* to a constructor, parametrized by a carrier X standing for the inductive type itself:

$$\begin{aligned}
&\mathsf{Args} : \mathsf{ConstrArity} \to \mathsf{Type} \to \mathsf{Type}\\
&\mathsf{Args}\ \mathsf{Done}\ X =_{\mathrm{def}} \mathbb{1}\\
&\mathsf{Args}\ (\mathsf{Nonrec}\ A\ k)\ X =_{\mathrm{def}} (a : A) \times \mathsf{Args}\ (k\ a)\ X\\
&\mathsf{Args}\ (\mathsf{Rec}\ I\ cs)\ X =_{\mathrm{def}} (I \to X) \times \mathsf{Args}\ cs\ X.
\end{aligned}$$

We call an element $as : \mathsf{Args}\ cs\ X$ an *argument list*: a nested tuple with one entry for each step of the telescope cs, where a non-recursive entry is a value $a : A$ and a recursive entry is a map $r : I \to X$ giving the subterm at each position. In the clauses below we destructure argument lists as (a, as) or (r, as), with as the remaining entries. A *signature* is a type of constructor labels together with an assignment of an arity to each label:

$$\mathsf{Signature} =_{\mathrm{def}} (\mathsf{Op} : \mathsf{Type}) \times (\mathsf{Op} \to \mathsf{ConstrArity})$$

Finally, a *constructor algebra* for a signature on a carrier X interprets every label as a constructor taking its arguments into the carrier:

$$\begin{aligned}
&\mathsf{ConstrAlgebra} : \mathsf{Signature} \to \mathsf{Type} \to \mathsf{Type}\\
&\mathsf{ConstrAlgebra}\ (\mathsf{Op},\ \mathsf{arity})\ X =_{\mathrm{def}} (o : \mathsf{Op}) \to \mathsf{Args}\ (\mathsf{arity}\ o)\ X \to X.
\end{aligned}$$

For a map $g : X \to Y$ and $cs : \mathsf{ConstrArity}$, we define $\mathsf{mapArgs}\ cs\ g : \mathsf{Args}\ cs\ X \to \mathsf{Args}\ cs\ Y$ by postcomposing each family of recursive subterms $r : I \to X$ with g and leaving non-recursive data untouched:

$$\begin{aligned}
&\mathsf{mapArgs}\ \mathsf{Done}\ g\ as =_{\mathrm{def}} as\\
&\mathsf{mapArgs}\ (\mathsf{Nonrec}\ A\ k)\ g\ (a, as) =_{\mathrm{def}} (a,\ \mathsf{mapArgs}\ (k\ a)\ g\ as)\\
&\mathsf{mapArgs}\ (\mathsf{Rec}\ I\ cs)\ g\ (r, as) =_{\mathrm{def}} (g \circ r,\ \mathsf{mapArgs}\ cs\ g\ as).
\end{aligned}$$

This action is functorial: it preserves identities and composition, and it carries an equivalence $X \simeq Y$ on carriers to an equivalence $\mathsf{Args}\ cs\ X \simeq \mathsf{Args}\ cs\ Y$ on arguments.

A configuration captures the shared inductive structure of two types by equipping each with a constructor algebra for a common signature, an eliminator for that algebra, and a β law for the eliminator.

Definition 2.2 (Configuration ✓)**.** Fix a signature $sig = (\mathsf{Op}, \mathsf{arity})$, a carrier C, and a constructor algebra $\alpha : \mathsf{ConstrAlgebra}\ sig\ C$. Relative to a *motive* $P : C \to \mathsf{Type}$, the *inductive hypotheses* at an argument list, $as : \mathsf{Args}\ cs\ C$, provide a proof of P at every recursive child, that is, at $r\ d$ for each recursive entry r of as and each position d. Formally,

$$\begin{aligned}
&\mathsf{IH} : (cs : \mathsf{ConstrArity}) \to (C \to \mathsf{Type}) \to \mathsf{Args}\ cs\ C \to \mathsf{Type}\\
&\mathsf{IH}\ \mathsf{Done}\ P\ as =_{\mathrm{def}} \mathbb{1}\\
&\mathsf{IH}\ (\mathsf{Nonrec}\ A\ k)\ P\ (a, as) =_{\mathrm{def}} \mathsf{IH}\ (k\ a)\ P\ as\\
&\mathsf{IH}\ (\mathsf{Rec}\ I\ cs)\ P\ (r, as) =_{\mathrm{def}} \big((d : I) \to P\ (r\ d)\big) \times \mathsf{IH}\ cs\ P\ as.
\end{aligned}$$

An *induction principle* for α takes a motive P together with, for each label o, a case sending arguments as and their inductive hypotheses to a value of $P\ (\alpha\ o\ as)$, and produces a dependent section:

$$\begin{aligned}
\mathsf{Ind}\ sig\ \alpha =_{\mathrm{def}} (P : C \to \mathsf{Type}) \to\\
\Big((o : \mathsf{Op})\ (as : \mathsf{Args}\ (\mathsf{arity}\ o)\ C) \to \mathsf{IH}\ (\mathsf{arity}\ o)\ P\ as \to P\ (\alpha\ o\ as)\Big) \to\\
(x : C) \to P\ x.
\end{aligned}$$

Any dependent function $f : (x : C) \to P\ x$ yields inductive hypotheses at every argument list, by applying f to each recursive child (see the full definition in the formalization ✓):

$$\mathsf{mkIH} : (cs : \mathsf{ConstrArity}) \to ((x : C) \to P\ x) \to (as : \mathsf{Args}\ cs\ C) \to \mathsf{IH}\ cs\ P\ as.$$

An eliminator $\mathsf{ind} : \mathsf{Ind}\ \mathit{sig}\ \alpha$ satisfies the *β law* when it computes on constructors: for all P, cases, label o, and arguments as,

$$\mathsf{ind}\ P\ \mathit{cases}\ (\alpha\ o\ as)\ =\ \mathit{cases}\ o\ as\ \big(\mathsf{mkIH}\ (\mathsf{arity}\ o)\ (\mathsf{ind}\ P\ \mathit{cases})\ as\big).$$

We write $\mathsf{Beta}\ \mathit{sig}\ \mathsf{ind}$ for the type of proofs that ind satisfies the β law. An *inductive algebra* for sig on C packages a constructor algebra with an induction principle and its β law (referred to as ι in prior proof repair work):

$$\mathsf{IndAlg}\ \mathit{sig}\ C =_{\mathrm{def}} (\alpha : \mathsf{ConstrAlgebra}\ \mathit{sig}\ C) \times (\mathsf{ind} : \mathsf{Ind}\ \mathit{sig}\ \alpha) \times \mathsf{Beta}\ \mathit{sig}\ \mathsf{ind}.$$

Finally, a *configuration* for sig between carriers C and D is a pair of inductive algebras for sig, one on each carrier:

$$\mathsf{Config}\ \mathit{sig}\ C\ D =_{\mathrm{def}} \mathsf{IndAlg}\ \mathit{sig}\ C \times \mathsf{IndAlg}\ \mathit{sig}\ D.$$

This definition internalizes the `DepConstr`/`DepElim`/`Iota` data that PUMPKIN Pi externally maintains; the results that follow are theorems about this reconstruction, and we support its faithfulness by recovering the tool's correctness criteria (Proposition 2.3) and case studies from it.

Though our signatures are non-indexed, our formalization handles the repair of inductive families in a way that is consistent with the proof repair literature. Our formalization includes the case study from Ringer et al. [48] of repairing between $\mathtt{Vec}\ A\ n$ and $\mathtt{Fin}\ n \to A$, where the per-index signature $\mathit{sig}\ n$ has a single non-recursive operation taking the n values ✓. Programs and proofs over vectors repair index-by-index, with the coherence conditions of Proposition 2.3 instantiated at each n. The restriction to non-indexed signatures thus reflects the class of changes contemporary repair tools perform, which does not include repairing types with differing index types.

A configuration can also satisfy correctness criteria [48] ensuring that it behaves coherently with the equivalence it is intended to capture. Our formal definition satisfies these criteria.

Proposition 2.3 (Correctness Conditions ✓)**.** *Fix $(\mathcal{A}_C, \mathcal{A}_D) : \mathsf{Config}\ \mathit{sig}\ C\ D$ with constructor algebras α_C, α_D, eliminators $\mathsf{ind}_C, \mathsf{ind}_D$, and β laws β_C, β_D. Let $\varphi : C \to D$ be the forward map of the equivalence induced by the configuration (Proposition 2.4). The configuration satisfies the three coherence conditions of PUMPKIN Pi [48].*

(1) constr_ok. *Constructors commute with φ: for every label o and arguments $as : \mathsf{Args}\ (\mathsf{arity}\ o)\ C$,*

$$\varphi\ (\alpha_C\ o\ as)\ =\ \alpha_D\ o\ \big(\mathsf{mapArgs}\ (\mathsf{arity}\ o)\ \varphi\ as\big).$$

(The formal statement is a heterogeneous strengthening, relating C-side and D-side arguments across the lifted argument equivalence; the form displayed here is the instance at φ-mapped arguments.)

(2) iota_ok. *Each eliminator computes on its own constructors: this is exactly the β law, β_C for ind_C and β_D for ind_D.*

(3) elim_ok. *The eliminators commute with φ: for every motive $P : D \to \mathsf{Type}$ and cases cases_D, there are transported C-side cases $\varphi^*\,\mathit{cases}_D$ (delegating to cases_D on the φ-mapped arguments and transporting back along* constr_ok*) such that, for every $x : C$,*

$$\mathsf{ind}_C\ (P \circ \varphi)\ (\varphi^*\,\mathit{cases}_D)\ x\ =\ \mathsf{ind}_D\ P\ \mathit{cases}_D\ (\varphi\ x).$$

Proof sketch. All three are straightforwardly derivable from the structure carried by the configuration and function extensionality. □

We apply the equivalence explicitly rather than using univalence and transport as in Ringer et al. [48]. The first condition, constr_ok, states that *φ respects the constructors*, sending a constructed value to the corresponding construction on the other side. The remaining two ask nothing of φ in isolation: iota_ok is just each type computing on its own constructors, and elim_ok is recovered from constr_ok once each type is known to be generated by its constructors in the strong sense the eliminators capture. We will see in Section 3 that this is no accident: a structure-respecting

map between two such canonically generated types must be an isomorphism, and naming this phenomenon precisely is what the algebraic reformulation buys. This is the first hint of our central thesis, that proof repair is an algebra isomorphism.

Additionally, our formalization enables us to prove a conjecture about configurations stated in Viola et al. [58]. The original proof repair work included extra data for an η principle in both the configuration and the correctness conditions; we confirm that this data is not needed, since it is derivable from the rest ✓.

Configurations were developed as an algorithm-friendly decomposition of type equivalences, which in a tool setting are externalized in some automation language. Equipped with our formal definitions, we are able to verify that configurations do indeed correspond to type equivalences.

Proposition 2.4 (Configurations and Inductive Equivalences are Equivalent ✓)**.** *Fix a signature sig and carriers C, D. In the presence of an inductive algebra on C, type equivalences $C \simeq D$ and configurations between C and D are equivalent:*

$$(\mathsf{IndAlg}\ sig\ C\ \times\ (C \simeq D))\ \simeq\ \mathsf{Config}\ sig\ C\ D.$$

Proof sketch. An inductive algebra on C appears on both sides, so, fixing one, it suffices to give an equivalence $(C \simeq D) \simeq \mathsf{IndAlg}\ sig\ D$. The inductive algebra transports directly across the equivalence. Conversely, a pair of inductive algebras determines an equivalence through their folds. Each inductive algebra sends a constructor algebra γ on a type Y to a map $\mathsf{fold}_C\ \gamma : C \to Y$, the eliminator at the constant motive Y, characterized by its β rule

$$\mathsf{fold}_C\ \gamma\ (\alpha_C\ o\ as)\ =\ \gamma\ o\ \big(\mathsf{mapArgs}\ (\mathsf{arity}\ o)\ (\mathsf{fold}_C\ \gamma)\ as\big).$$

Writing $\varphi = \mathsf{fold}_C\ \alpha_D$ and $\psi = \mathsf{fold}_D\ \alpha_C$, the composite $\psi \circ \varphi$ is a fold of α_C into itself; so is the identity, and the fold into a fixed algebra is unique, so $\psi \circ \varphi = \mathsf{id}_C$, and symmetrically $\varphi \circ \psi = \mathsf{id}_D$. These two constructions are mutually inverse, giving the equivalence. □

This is not a surprising result, but it confirms the intuition with which configurations were designed and serves as a benchmark upon which the rest of the paper improves. However, this result does not give a satisfying mathematical description of repair, since it simply transports half of the data of a configuration across an equivalence to obtain the other half. We show (Section 3) that P-algebra isomorphisms provide an elegant mathematical answer to the question of what transformations configurations capture, and from this answer we obtain results about possible extensions of repair algorithms and the limits of repair.

In the typical case of repair, a proof engineer has a natural choice of configuration in mind, but, in general, unnatural and useless configurations exist. For example, a configuration could reuse the same type or trivially specify a single constructor and eliminator. It is up to the proof engineer and repair tooling to ensure that configurations execute a useful form of repair. We return to this point in Section 3, where the algebraic reformulation makes this precise.

Configurations are a useful framework for building proof repair algorithms, but they expose a high level of syntactic detail that is a burden when trying to reason about repair in general. The rest of the paper is dedicated to developing an abstract formulation of repair that is more amenable to reasoning about the limits and properties of repair. We proceed by connecting configurations to P-algebras, and then observing the results of this connection in the context of repair algorithms.

3 Algebraic Foundations for Proof Repair

We formulate repair in the language of algebras on polynomial endofunctors [7, 11], which we introduce from scratch for readers unfamiliar with them. We first introduce P-algebras and the algebraic view of inductive datatypes, providing the necessary background. We then establish that configurations are a representation of algebra isomorphisms. Finally, we showcase a purely algebraic form of proof repair.

3.1 Warming up to Algebras

We begin with a tour of P-algebras. Note that we are working within type theory, and so the definitions we use have a type-theoretic flavor and are not the precise definitions one may see in a category theory textbook. We nonetheless refer to these type-theoretic notions by their usual category-theoretic names.

Definition 3.1 (Polynomial Endofunctor ✓)**.** The polynomial endofunctor takes a type A of *shapes* and a family $B : A \to \mathsf{Type}$ assigning to each shape its type of *positions*:

$$P : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to \mathsf{Type} \to \mathsf{Type}$$
$$P\,A\,B\,X \;=_{\text{def}}\; (a : A) \times (B\,a \to X)$$

An element of $P\,A\,B\,X$ is a shape a together with a $B\,a$-indexed family of children drawn from X. The action of $P\,A\,B$ on a function $f : X \to Y$ is postcomposition on the children, which makes $P\,A\,B$ a functor.

$$\mathsf{P\text{-}map} : (X \to Y) \to P\,A\,B\,X \to P\,A\,B\,Y$$
$$\mathsf{P\text{-}map}\,f\,(a, u) \;=_{\text{def}}\; (a, f \circ u)$$

We call a pair $(A : \mathsf{Type},\ B : A \to \mathsf{Type})$ a *polynomial*.

Definition 3.2 (P-algebra ✓)**.** Fix $A : \mathsf{Type}$ and $B : A \to \mathsf{Type}$. For the polynomial endofunctor $P\,A\,B$, a *P-algebra* is a pair (C, σ) of a carrier C and a structure map $\sigma : P\,A\,B\,C \to C$. We define a P-algebra formally as follows:

$$\mathsf{P\text{-}Alg}\,A\,B : \mathsf{Type}$$
$$\mathsf{P\text{-}Alg}\,A\,B =_{\text{def}} (C : \mathsf{Type}) \times (P\,A\,B\,C \to C)$$

Example 3.3 (Nat Algebra ✓)**.** Consider the functor

$$P : \mathsf{Type} \to \mathsf{Type} \qquad P\,X =_{\text{def}} \mathbb{1} + X.$$

A P-algebra is a pair (C, σ) with a structure map $\sigma : \mathbb{1} + C \to C$. By the universal property of the coproduct, such a map is equivalently a point $z : C$ together with an endomap $s : C \to C$, which we bundle as $\sigma = [\,z\,,\ s\,]$. Thus, a P-algebra is exactly a type equipped with an interpretation of a `zero` constructor and a `succ` constructor. The type `Nat` carries the algebra $(\mathtt{Nat}, [\,\mathtt{zero}, \mathtt{succ}\,])$, reading the nullary shape as `zero` and the unary shape as `succ`. The payoff of the algebraic view is that the same functor has many carriers: `Bin` carries $(\mathtt{Bin}, [\,\mathtt{bzero}, \mathtt{binSucc}\,])$, a different interpretation of the same two shapes.

Definition 3.4 (Algebra Homomorphism ✓)**.** Given a polynomial (A, B) and two algebras (C, σ) and (D, τ) for $P\,A\,B$, an algebra homomorphism from (C, σ) to (D, τ) is a function $h : C \to D$ such that the following diagram commutes:

$$\begin{array}{ccc} PC & \xrightarrow{Ph} & PD \\ \downarrow{\scriptstyle \sigma} & & \downarrow{\scriptstyle \tau} \\ C & \xrightarrow{h} & D \end{array}$$

Formally, we define the type of algebra homomorphisms as the function space between the carriers paired with the proof of the above commutative diagram:

$$\mathsf{isAlgHom} : (X\ Y\ : \mathsf{P\text{-}Alg}\,A\,B) \to (\mathsf{Carrier}\,X \to \mathsf{Carrier}\,Y) \to \mathsf{Type}$$
$$\mathsf{isAlgHom}\ (C,\ \sigma)\ (D,\ \tau)\ h =_{\text{def}} h \circ \sigma = \tau \circ \mathsf{P\text{-}map}\,h$$

$$\mathsf{AlgHom} : \mathsf{P\text{-}Alg}\,A\,B \to \mathsf{P\text{-}Alg}\,A\,B \to \mathsf{Type}$$
$$\mathsf{AlgHom}\,X\,Y =_{\text{def}} (h\ :\ \mathsf{Carrier}\,X \to \mathsf{Carrier}\,Y) \times (\mathsf{isAlgHom}\,X\,Y\,h)$$

Algebra homomorphisms compose (AlgHom-∘), and every algebra carries the identity homomorphism id-hom, whose carrier map is the identity function.

Example 3.5 (Nat Homomorphism ✓)**.** Homomorphisms between P-algebras are maps that preserve the constructors of the source algebra. The map `fromNat` : `Nat` → `Bin` is a homomorphism:

$$\begin{array}{ccc} \mathbb{1} + \mathtt{Nat} & \xrightarrow{\mathbb{1} + \mathtt{fromNat}} & \mathbb{1} + \mathtt{Bin} \\ {\scriptstyle [\,\mathtt{zero},\,\mathtt{succ}\,]}\downarrow & & \downarrow{\scriptstyle [\,\mathtt{bzero},\,\mathtt{binSucc}\,]} \\ \mathtt{Nat} & \xrightarrow[\mathtt{fromNat}]{} & \mathtt{Bin} \end{array}$$

This commutation is precisely the constr_ok condition of Proposition 2.3: a homomorphism is a map that respects the constructors.

Since the central claim of this paper is that proof repair is an algebra isomorphism, we next fix precisely what an isomorphism of algebras is.

Definition 3.6 (Algebra Equivalence ✓)**.** Given a polynomial (A, B) and algebras X, Y : P-Alg $A\,B$, an algebra homomorphism f : AlgHom $X\,Y$ is an *algebra equivalence* when it has a right inverse and a left inverse:

$$\begin{aligned} &\mathsf{isAlgEquiv} : (X\ Y : \mathsf{P\text{-}Alg}\ A\ B) \to \mathsf{AlgHom}\ X\ Y \to \mathsf{Type} \\ &\mathsf{isAlgEquiv}\ X\ Y\ f =_{\mathrm{def}} \big((g : \mathsf{AlgHom}\ Y\ X) \times (\mathsf{AlgHom\text{-}}\circ\ f\ g = \mathsf{id\text{-}hom}_X)\big) \\ &\qquad\qquad \times \big((h : \mathsf{AlgHom}\ Y\ X) \times (\mathsf{AlgHom\text{-}}\circ\ h\ f = \mathsf{id\text{-}hom}_Y)\big) \\ &\mathsf{AlgEquiv} : \mathsf{P\text{-}Alg}\ A\ B \to \mathsf{P\text{-}Alg}\ A\ B \to \mathsf{Type} \\ &\mathsf{AlgEquiv}\ X\ Y =_{\mathrm{def}} (f : \mathsf{AlgHom}\ X\ Y) \times \mathsf{isAlgEquiv}\ X\ Y\ f \end{aligned}$$

Forgetting the homomorphism witnesses, the carrier maps of an algebra equivalence form a type equivalence Carrier $X \simeq$ Carrier Y; an algebra equivalence is exactly a type equivalence whose two maps respect the algebra structure. We say *algebra isomorphism* interchangeably.

Definition 3.7 (Homotopy-initial Algebra ✓)**.** Given a polynomial (A, B), an algebra X : P-Alg $A\,B$ is *homotopy-initial* if for every Y : P-Alg $A\,B$ the type of homomorphisms $X \to Y$ is contractible (isContr). Formally,

$$\begin{aligned} &\mathsf{isInitAlg} : \mathsf{P\text{-}Alg}\ A\ B \to \mathsf{Type} \\ &\mathsf{isInitAlg}\ X =_{\mathrm{def}} (Y : \mathsf{P\text{-}Alg}\ A\ B) \to \mathsf{isContr}(\mathsf{AlgHom}\ X\ Y) \end{aligned}$$

For brevity, we may refer to homotopy-initial algebras as initial algebras.

Equipped with our P-algebra toolbox, we turn back to our motivating example and phrase it in terms of algebras. In this example we are implicitly relying on the correspondence between homotopy-initial algebras and W-types established by Awodey, Gambino, and Sojakova [12]. Of course, we are using Agda's `data` declaration instead of W-types, but we are still able to put an algebra structure on the Agda datatypes as they capture the same inductive structure.

Example 3.8 (Motivating Example with Initial Algebras ✓)**.** First, we obtain that `Nat` and `Bin` are initial algebras for the functor $P\,X =_{\mathrm{def}} \mathbb{1} + X$.

```
NatInit : isInitAlg (Nat , [ zero , succ ])
BinInit : isInitAlg (Bin , [ bzero , binSucc ])
```

The unique homomorphism out of the initial algebra is the recursor, and its dependent strengthening is the induction principle. Both are derived from `isInitAlg` alone—the recursor as the underlying function of the unique homomorphism, and the eliminator from the section it determines on the total space.

```
recNat : {D : Type} -> D -> (D -> D) -> Nat -> D
recNat-zero : recNat z s zero  = z
recNat-succ : recNat z s (succ n) = s (recNat z s n)

indNat : (Q : Nat -> Type)  -> Q zero  -> ((n : Nat) -> Q n -> Q (succ n))  -> (n : Nat)  -> Q n
indNat-zero : indNat Q qz qs zero = qz
indNat-succ : indNat Q qz qs (succ n) = qs n (indNat Q qz qs n)
```

With this, we can write the programs from before entirely in terms of universal properties. Here addition is itself defined by the recursor on its first argument, $m + n =_{\text{def}} \texttt{recNat}\ n\ \texttt{succ}\ m$, so the equations of + used below are the computation rules of `recNat`:

```
isZeroN : Nat -> Bool
isZeroN = recNat true (λ _ -> false)

zUnit : (n : Nat) -> n + zero = n
zUnit = indNat (λ n -> n + zero = n)  recNat-zero  (λ n ih -> recNat-succ • ap succ ih)
```

Because `Bin` is initial for the same functor, it carries the same universal property, the same recursor and eliminator, now out of `Bin`. Repair is then a simple instantiation of the universal property at the new initial algebra: `recNat`/`indNat` become `recBin`/`indBin` and `zero`/`succ` become `bzero`/`binSucc`, while the programs' defining data is untouched.

```
recBin : {D : Type} -> D -> (D -> D) -> Bin -> D
indBin : (Q : Bin -> Type)  -> Q bzero  -> ((b : Bin) -> Q b -> Q (binSucc b))  -> (b : Bin) -> Q b

isZeroB : Bin -> Bool
isZeroB = recBin true (λ _ -> false)

zUnitB : (b : Bin) -> b + bzero = b
zUnitB = indBin  (λ b -> b + bzero = b)  recBin-zero (λ b ih -> recBin-succ • ap binSucc ih)
```

The algebraic perspective on proof repair is epitomized by the preceding example; the two very different representations of natural numbers are initial for the same functor and hence must have a unique isomorphism between them. This is our thesis emerging: `Nat` and `Bin` are repaired precisely through this algebra isomorphism. We spend the rest of this section cementing this idea and establishing that our slogan (proof repair is an algebra isomorphism) is general and concrete.

3.2 Connecting Configurations and Algebras

Type equivalences are an unsatisfactory answer to the question of what configurations capture. Indeed, because type equivalences do not capture the inductive structure of configurations, our proof of Proposition 2.4 assumed an inductive algebra to generate a configuration from an equivalence. We use P-algebras to put our understanding of configurations on firmer ground. The main idea of this section is that configurations are a syntax-sensitive method for capturing a shared inductive structure on types, which can be represented by algebra isomorphisms. We begin to develop this idea by coaxing configurations into something less burdened by syntax.

Every configuration has a canonical form that we call a *W-type configuration*. These configurations emulate W-types by capturing the inductive structure of types with a single constructor that takes a non-recursive and a recursive argument.

Definition 3.9 (W-type Configurations ✓)**.** We define a one-constructor *W-type signature*: a single operation whose arity supplies a shape $a : A$ and then a $B\,a$-indexed family of recursive children.

$$\begin{aligned}
&\mathsf{WTypeArity} : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to \mathbb{1} \to \mathsf{ConstrArity}\\
&\mathsf{WTypeArity}\,A\,B\,_ =_{\mathrm{def}} \mathsf{Nonrec}\,A\,(\lambda a \to \mathsf{Rec}\,(B\,a)\,\mathsf{Done})\\
&\\
&\mathsf{WTypeSignature} : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to \mathsf{Signature}\\
&\mathsf{WTypeSignature}\,_\,_\,.\mathsf{Signature.Op} =_{\mathrm{def}} \mathbb{1}\\
&\mathsf{WTypeSignature}\,A\,B\,.\mathsf{Signature.arity} =_{\mathrm{def}} \mathsf{WTypeArity}\,A\,B\\
&\\
&\mathsf{WTypeIndAlg} : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to (C : \mathsf{Type}) \to \mathsf{Type}\\
&\mathsf{WTypeIndAlg}\,A\,B\,C =_{\mathrm{def}} \mathsf{IndAlg}\,(\mathsf{WTypeSignature}\,A\,B)\,C
\end{aligned}$$

A *W-type configuration* is the special configuration where the signature is a W-type signature, collapsing all per-constructor data into one constructor with one non-recursive argument and one recursive argument:

$$\begin{aligned}
&\mathsf{WTypeConfig} : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to (C\,D : \mathsf{Type}) \to \mathsf{Type}\\
&\mathsf{WTypeConfig}\,A\,B\,C\,D =_{\mathrm{def}} \mathsf{Config}\,(\mathsf{WTypeSignature}\,A\,B)\,C\,D
\end{aligned}$$

Whereas configurations are useful for capturing the syntactic detail of particular representations of inductive types, W-type configurations collapse that detail, which makes them useful for reasoning about the underlying structure configurations capture. Indeed, we prove every configuration is represented by a corresponding W-type configuration.

Lemma 3.10 (Configurations Faithfully Embed into W-type Configurations ✓)**.** *Every configuration over an arbitrary signature sig determines a W-type configuration on the same carriers. We read off the polynomial whose shapes pair a constructor label with its non-recursive data (an argument list at the carrier $\mathbb{1}$) and whose positions (*recPos*) are the recursive arguments occurring in that shape:*

$$A \;=_{def}\; (c : \mathsf{Op}) \times \mathsf{Args}(\mathsf{arity}\,c)\;\mathbb{1}, \qquad B\,(c, s) \;=_{def}\; \mathsf{recPos}(\mathsf{arity}\,c)\,s,$$

and the configuration transports to a W-type configuration for this (A, B):

$$\mathsf{Config}\,sig\,C\,D \to (A : \mathsf{Type}) \times ((B : (A \to \mathsf{Type})) \times (\mathsf{WTypeConfig}\,A\,B\,C\,D))$$

The embedding is faithful: writing $\varphi, \psi : C \leftrightarrows D$ for the forward and backward maps of the repair equivalence induced by the configuration (Proposition 2.4), and φ_W, ψ_W for those induced by its W-type configuration, the maps are equal:

$$\varphi_W = \varphi \qquad \textit{and} \qquad \psi_W = \psi.$$

Thus, arbitrary configurations translate to W-type configurations with the repair maps preserved. This canonicalization is essential, as it lets us transfer type-theoretic results about W-types to repair. In particular, by showing that configurations over arbitrary signatures carry the same repair data as W-type configurations, we connect repair to the study of W-types by Awodey et al. [12]. Because nothing about repair is lost by passing to the canonical form, we state our main theorem over W-type configurations.

Theorem 3.11 (W-type Configurations are Unique Algebra Equivalences ✓)**.** *For any polynomial (A, B) and carriers C, D, a W-type configuration is exactly an algebra equivalence between algebras on C and D whose source is initial:*

$$\begin{aligned}
\mathsf{WTypeConfig}\,A\,B\,C\,D \;\simeq\; &(\sigma_C : P\,A\,B\,C \to C)\\
&\times (\sigma_D : P\,A\,B\,D \to D)\\
&\times \mathsf{AlgEquiv}\,(C, \sigma_C)\,(D, \sigma_D)\\
&\times \mathsf{isInitAlg}\,(C, \sigma_C)
\end{aligned}$$

Proof sketch. The carrier-level equivalence $\mathsf{WTypeIndAlg}\,A\,B\,X \simeq \mathsf{InitAlgOn}\,A\,B\,X$ (with InitAlgOn as in Definition 3.13) is given by the homotopy-initiality and W-type correspondence of Awodey et al. [12], which we then use to get an equivalence between the corresponding pairs. We then construct an equivalence between a pair of initial algebras and a single initial algebra together with an algebra equivalence to the other. Composing these equivalences yields the desired conclusion. That the equivalence is unique ✓ is given by the universal property and shown by Awodey et al. [12]. □

This result gives a bridge between repair and algebras that justifies our slogan that repair is an algebra isomorphism. Once configurations are translated into W-type configurations the theorem is unsurprising, but it is valuable for the conceptual connection it cements between repair and P-algebras. Indeed, the algebraic perspective clarifies properties of repair that the purely syntactic view offered by configurations (Section 2) does not. What is remarkable about this correspondence is that it shows how configurations realize algebra isomorphisms as syntactic transformations that remove old dependencies (Proposition 3.15), so that repair carries out a mathematically defined transformation without ever referring to the old type.

We now know that a configuration packages precisely the data of a universal construction on two carriers, namely the property of both carriers being initial for a common functor. For a proof engineer, the question of when repair between two datatypes is natural can now be posed precisely and formally as whether there is a natural functor for which both types present an initial algebra. Given a positive answer, the rest of repair is determined; the transformation the algorithm executes is exactly the algebra isomorphism that initiality forces to exist uniquely.

Additionally, we can state that configurations can handle exactly those datatype changes representable by an algebra isomorphism. This class is broad in practice, since the initial algebra of a functor pins down the data structure the programmer has in mind, so changes of representation that preserve that structure fall within its reach. This characterization has force only relative to a fixed polynomial, since any equivalence is an algebra isomorphism for a constant polynomial functor.

Example 3.12 (Motivating Example with Algebra Equivalence ✓)**.** Because `Nat` and `Bin` are initial for the same functor, initiality hands us a homomorphism each way and forces their two composites to be endomorphisms of an initial algebra, hence the identities. This gives the unique algebra isomorphism `toNat/fromNat : Nat ≃ Bin`. Now the swap of initial algebras from the previous example is the action of this isomorphism. The new recursor `recBin z s` precomposed with `fromNat` is again a homomorphism, so by initiality of `Nat` it must coincide with `recNat z s`; the swap is forced, not chosen. Each old program is thus equal to the repaired program precomposed with the isomorphism:

```
recNatSwap : recNat z s = recBin z s ∘ fromNat

isZeroCoh : (n : Nat) -> isZeroN n = isZeroB (fromNat n)
isZeroCoh n = happly recNatSwap n
```

The same story plays out one level up for proofs: the dependent eliminators commute with the isomorphism (elim_ok), so the repaired proof transports back to the original along the round-trip witness `toNatFromNat`. Writing `zUnitB-nat` for the composed `Bin`-side proof from the introduction (ranging over `toNat b`, so the predicate stays on the `Nat` side), we have

```
zUnitB-nat : (b : Bin) -> toNat b + zero = toNat b
zUnitCoh : (n : Nat) -> tpt (λ m -> m + zero = m) (toNatFromNat n) (zUnitB-nat (fromNat n)) =
     zUnit n
```

Repair across `Nat` and `Bin` is, in this precise sense, the unique algebra isomorphism between them: the proof engineer's swap of constructors and eliminators is exactly its action on terms, and

the coherence conditions of Proposition 2.3 are precisely the witnesses that this action sends each original term to its repair. However, the naive use of an algebra isomorphism is not repair since it does not rid the code of the old dependency. Shortly, we characterize the ability to get rid of old datatypes in terms of P-algebras.

3.3 Algebraic Repair

Now, we can formulate repair entirely in P-algebraic terms.

Definition 3.13 (Algebra Configuration ✓)**.** Given a polynomial (A, B) and carriers C, D, an *initial algebra structure* on C packages an algebra with a proof of initiality:

$$\mathsf{InitAlgOn}\, A\, B\, C \;=_{\mathrm{def}}\; (\sigma : P\, A\, B\, C \to C) \times \mathsf{isInitAlg}\, (C, \sigma).$$

An *algebra configuration* between C and D is a pair of initial algebra structures, one on each carrier:

$$\mathsf{AlgConfig}\, A\, B\, C\, D \;=_{\mathrm{def}}\; \mathsf{InitAlgOn}\, A\, B\, C \;\times\; \mathsf{InitAlgOn}\, A\, B\, D.$$

An algebra configuration is deliberately the algebraic analogue of a configuration (Definition 2.2): where a configuration carries an inductive algebra on each carrier, an algebra configuration carries an initial P-algebra on each carrier.

An algebra configuration determines an algebra isomorphism between its two carriers by the initiality of each side, which we refer to as the repair map. Forgetting the homomorphism witnesses and keeping the underlying carrier maps recovers exactly the type equivalence $C \simeq D$ of Proposition 2.4; the algebraic statement is the upgrade, an equivalence whose maps preserve the inductive structure rather than a bare bijection of carriers.

This algebraic reformulation is more than a change of vocabulary. The algebraic picture replaces the syntax-heavy bookkeeping of a configuration with a single mathematical object, a pair of initial algebras for one polynomial functor. This puts the foundations of proof repair on well-trodden mathematical ground, and clarifies what repair computes. It equally clarifies the proof engineer's task. The engineer needs only the common algebraic structure shared by the source and target; the repair, together with its correctness, is then determined. The payoff of the slogan that proof repair is an algebra isomorphism is that it names the goal of a repair algorithm precisely and grounds it in a mathematical fact independent of any particular implementation.

The coherence conditions are subsumed by this viewpoint. In the configuration setting of Proposition 2.3, the three conditions are properties one must establish of the induced map. From an algebra configuration they fall out of initiality, and our development derives all three from the two initial algebras alone ✓. Commutation with the constructors (constr_ok) is given by the homomorphism square and the remaining two rest on the source being canonically generated by its constructors, which for an initial algebra is exactly its dependent eliminator and the eliminator's computation rule. This is the precise content of the foreshadowing in Section 2.

Proposition 3.14 (Algebra Config iff Configuration ✓)**.** *Over carriers C, D, write*

$$\mathsf{SigConfig}\, C\, D =_{def} (sig : \mathsf{Signature}) \times \mathsf{Config}\, sig\, C\, D$$

$$\mathsf{PAlgConfig}\, C\, D =_{def} (A : \mathsf{Type}) \times (B : A \to \mathsf{Type}) \times \mathsf{AlgConfig}\, A\, B\, C\, D$$

for a configuration and algebra configuration over a signature. There are maps in both directions,

$$\mathsf{SigConfig}\, C\, D \;\leftrightarrow\; \mathsf{PAlgConfig}\, C\, D,$$

that do not form an equivalence because the forward map always produces a W-type signature and the map back recovers the polynomial (A, B) only up to equivalence.

Proof sketch. Both maps are the carrier-level bridge of Awodey et al. [12] applied on each carrier, which identifies a W-type inductive algebra with a homotopy-initial algebra. Restricting to a W-type signature upgrades the maps to the type equivalence of Theorem 3.11, with both round-trips proved. □

In our example and other repair case studies, the repaired program is always the *same* program, with the universal property instantiated at the target and the rest of the data untouched. The algebraic framework lets us formulate this as a general statement.

Proposition 3.15 (Repair by Instantiation ✓)**.** *Given a polynomial* (A, B)*, fix an algebra configuration between* C *and* D *with repair map* φ*, and* $(Y, \gamma) : \mathsf{P\text{-}Alg}\, A\, B$*. Writing* $\mathsf{fold}_C\, \gamma : C \to Y$ *and* $\mathsf{fold}_D\, \gamma : D \to Y$ *for the unique homomorphisms out of each initial algebra,*

$$\mathsf{fold}_C\, \gamma \;=\; \mathsf{fold}_D\, \gamma \circ \varphi.$$

Proof. Both sides are algebra homomorphisms $(C, \sigma_C) \to (Y, \gamma)$, and the hom-space out of an initial algebra is contractible. □

A program written against the universal property consists of defining data γ together with an instantiation of fold at one initial algebra. By Proposition 3.15, repair is re-instantiation: replace fold_C by fold_D, keeping γ. The repaired term $\mathsf{fold}_D\, \gamma$ is built from D's universal property and γ alone, so it mentions the source type C exactly when γ does. When the program is written generically, its defining data is carrier-free (like the recursor cases `true` and $\lambda_ \to$ `false`), and the repaired term then mentions C nowhere. When γ does mention C, those occurrences are further uses of C's algebra structure, repairable by another instance of the same swap. Whether a term mentions C is a property of its syntax, which the type theory cannot express internally.

Proposition 3.15 folds into a fixed algebra (Y, γ), covering recursion but not induction, whose motive $P : C \to \mathsf{Type}$ mentions the carrier. It applies anyway, since induction is a fold into the total-space algebra on $(x : C) \times P\, x$, with the eliminator its first-projection section ✓. Because the motive mentions C, the target is transported along the repair map (unique by Theorem 3.11) rather than shared. Projecting the resulting equation of pairs onto the proof component is exactly `elim_ok` (Proposition 2.3), which is why that condition carries transported cases φ^**cases*. The transported motive still mentions C, but only through constructors and folds, themselves repairable by further swaps; iterating over motive and cases yields the native D-side proof, as in `zUnitB` of Section 2.

We view Proposition 3.15 as stating that repair exploits a kind of *representation independence* [9] for algebras in type theory. Programs and proofs written over a specific datatype are independent of the specific implementation of that datatype and are only sensitive to the underlying algebra, which is the key idea behind repair. This is why repair tools can operate by syntactic instantiation, supplying the programs' defining data to the same universal property at a different carrier. Thus, the algebraic perspective additionally gives a mathematical account of why repair can eliminate the source type as a dependency.

3.4 Discussion: Repair without Univalence

While our results have made real use of the homotopy interpretation of the intensional identity type, nowhere have we relied upon univalence: the entire inductive development is carried out in intensional Martin-Löf type theory with function extensionality as the only axiom (Section 1). This is notable because prior proof repair work was largely framed within a univalent context, where the correctness conditions were thought to require univalence [48]. Our development thus broadens the class of type theories in which this kind of proof repair is known to be justified. For example, our results apply to settings where univalence is inconsistent, such as Lean, whose type theory has uniqueness of identity proofs [2].

4 Dualizing Repair: Coinductive Types and Final Coalgebras

With the algebraic perspective on proof repair in hand, we are able to dualize our development to coinductive types and final coalgebras. Mechanized developments model infinite data and processes via coinductive types, which face the same burden of representation change that motivated repair

for inductive types. Yet no existing repair tool addresses them. In this section, we work within Cubical Agda [57], which supports a notion of equality that is amenable to coinductive types and bisimulation.

Coinductive types are defined in terms of destructors, so programs interact with coinductive data principally by observation. Moreover, coinductive types enable reasoning about infinite objects through the principle of coinduction [32]. A classic infinite object is a stream, an infinite sequence over a given type:

$$(a_0, a_1, a_2, \ldots) : \mathsf{Stream}\, A \qquad A : \mathsf{Type} \qquad a_i : A.$$

The coinductive type of streams is defined as follows:

```
record Stream (A : Type) : Type where
  coinductive
  field
    hd : A
    tl : Stream A
```

We define programs that construct streams by specifying the observations that can be made of the stream—its head and its tail—a style known as copattern matching [5]. Here is the program that adjoins an element to the beginning of a stream:

```
cons : A -> Stream A -> Stream A
hd (cons a s) = a
tl (cons a s) = s
```

Using the principle of coinduction, we may verify that any stream can be decomposed into `cons` applied to its head and tail. In Cubical Agda even this proof is built by copattern matching, a point we return to at the end of Section 4.1:

```
streamComp : (s : Stream A) -> cons (hd s) (tl s) = s
hd (streamComp s i) = hd s
tl (streamComp s i) = tl s
```

Just like inductive types, coinductive types that capture the same data structure can have different definitions. For example, a library may also define the type of streams with a single destructor that packages the head and tail, along with its own `cons'`:

```
record Stream' (A : Type) : Type where
  coinductive
  field
    obs : A × Stream' A
```

```
cons' : A -> Stream' A -> Stream' A
obs (cons' a s) = (a , s)
```

Our goal in this section is to lay the foundations for proof repair algorithms that can handle changes between coinductive types, such as the change from `Stream` to `Stream'`. We will work backward, utilizing the coalgebraic perspective on coinductive types to develop repair in terms of universal properties, and then extract a notion of configurations for coinductive types that exposes syntactic detail for an implementation.

4.1 Warming up to Coalgebras

We begin by dualizing the algebraic setup in Section 3 to coalgebras. We internalize the definitions given in standard sources [6, 29, 52] on P-coalgebras within the type theory of Cubical Agda.

Definition 4.1 (Coalgebra ✓)**.** Fix $A : \mathsf{Type}$ and $B : A \to \mathsf{Type}$. For the polynomial endofunctor $P\,A\,B$, a *coalgebra* is a pair (X, γ) of a carrier X and a structure map $\gamma : X \to P\,A\,B\,X$.

$$\mathsf{P\text{-}Coalg}\, A\, B : \mathsf{Type}$$
$$\mathsf{P\text{-}Coalg}\, A\, B =_{\mathrm{def}} (X : \mathsf{Type}) \times (X \to P\,A\,B\,X)$$

Example 4.2 (Stream Coalgebra ✓)**.** Fix $A : \mathsf{Type}$ and consider the functor

$$P : \mathsf{Type} \to \mathsf{Type} \qquad P\,X =_{\mathrm{def}} A \times X.$$

A coalgebra for this functor is a pair (Z, γ) where $Z : \text{Type}$ and $\gamma : Z \to A \times Z$. We can understand this as a collection of streams over A with dynamics given by γ, which, by the universal property of products, amounts to a pair

$$\langle \gamma_1, \gamma_2 \rangle : (Z \to A) \times (Z \to Z).$$

Thus, we have a "stream system" given by a collection of streams, an output function γ_1 (which gives the head of a stream), and a transition function γ_2 (which gives the tail of a stream).

Remark 4.3. Fix $A : \text{Type}$ and $B : A \to \text{Type}$. Because the definition of polynomial endofunctor (Definition 3.1) is the dependent sum given by A and B, we can always present a coalgebra structure map $\gamma : Z \to P\,A\,B\,Z$ as a *shape observation* $\gamma_1 : Z \to A$ and a *position observation* $\gamma_2 : (z : Z) \to B\,(\gamma_1\,z) \to Z$.

Definition 4.4 (Coalgebra Homomorphism ✓)**.** Given a polynomial (A, B) and two coalgebras (X, γ) and (Y, δ) for $P\,A\,B$, a *coalgebra homomorphism* from (X, γ) to (Y, δ) is a function $f : X \to Y$ such that the following diagram commutes:

$$\begin{array}{ccc} X & \xrightarrow{\gamma} & PX \\ \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle Pf} \\ Y & \xrightarrow{\delta} & PY \end{array}$$

Formally, we define the type of coalgebra homomorphisms as the function space between the carriers paired with the proof of the above commutative diagram:

$$\begin{aligned} &\mathsf{isCoalgHom} : (X\ Y : \mathsf{P\text{-}Coalg}\,A\,B) \to (\mathsf{Carrier}\,X \to \mathsf{Carrier}\,Y) \to \mathsf{Type} \\ &\mathsf{isCoalgHom}\,(X, \gamma)\,(Y, \delta)\,f =_{\text{def}} \mathsf{P\text{-}map}\,f \circ \gamma = \delta \circ f \\ &\mathsf{CoalgHom} : \mathsf{P\text{-}Coalg}\,A\,B \to \mathsf{P\text{-}Coalg}\,A\,B \to \mathsf{Type} \\ &\mathsf{CoalgHom}\,X\,Y =_{\text{def}} (f : \mathsf{Carrier}\,X \to \mathsf{Carrier}\,Y) \times \mathsf{isCoalgHom}\,X\,Y\,f \end{aligned}$$

Example 4.5 (Stream Homomorphism ✓)**.** Homomorphisms between coalgebras are maps that preserve and reflect the dynamics of the source coalgebra. This means a homomorphism of stream coalgebras preserves outputs (heads) and transitions (tails):

$$\begin{array}{ccc} Z & \xrightarrow{\quad h \quad} & Y \\ \downarrow{\scriptstyle \langle \gamma_1, \gamma_2 \rangle} & & \downarrow{\scriptstyle \langle \sigma_1, \sigma_2 \rangle} \\ A \times Z & \xrightarrow{1_A \times h} & A \times Y \end{array}$$

Coalgebras support a notion of finality dual to the initiality of algebras. We introduce two notions that will become useful for our development: a final coalgebra, whose universal property pins down only the underlying function of the unique homomorphism, and a homotopy-final coalgebra, which additionally pins down the homomorphism witness.

Definition 4.6 (Final Coalgebra ✓)**.** Given a polynomial (A, B), a coalgebra $X : \mathsf{P\text{-}Coalg}\,A\,B$ is *final* if for every $Y : \mathsf{P\text{-}Coalg}\,A\,B$ there exists a coalgebra homomorphism $Y \to X$, and the underlying function of that homomorphism is unique.

$$\begin{aligned} &\mathsf{isFinal} : \mathsf{P\text{-}Coalg}\,A\,B \to \mathsf{Type} \\ &\mathsf{isFinal}\,X =_{\text{def}} (Y : \mathsf{P\text{-}Coalg}\,A\,B) \to (h : \mathsf{CoalgHom}\,Y\,X) \\ &\qquad\qquad \times \left((h' : \mathsf{CoalgHom}\,Y\,X) \to \mathsf{fst}\,h = \mathsf{fst}\,h'\right) \end{aligned}$$

Definition 4.7 (Homotopy-Final Coalgebra ✓)**.** Given a polynomial (A, B), a coalgebra $X : \mathsf{P\text{-}Coalg}\,A\,B$ is *homotopy-final* if for every $Y : \mathsf{P\text{-}Coalg}\,A\,B$ the type of coalgebra homomorphisms

$Y \to X$ is contractible.

$$\mathsf{isHFinal} : \text{P-Coalg}\, A\, B \to \mathsf{Type}$$
$$\mathsf{isHFinal}\, X =_{\text{def}} (Y : \text{P-Coalg}\, A\, B) \to \mathsf{isContr}(\mathsf{CoalgHom}\, Y\, X)$$

The reason for distinguishing between these two notions of finality is that proof assistant support for coinductive types that have homotopy-finality is not fully developed. Cubical Agda formalizations have confirmed finality [20] for containers, but whether coinductive records in Agda have homotopy-finality is open. As far as we know, whether coinductive types in Rocq have finality is open, and Lean would trivially have homotopy-finality (given finality) since it supports the uniqueness of identity proofs axiom. State-of-the-art results [8, 26] formulate M-types internally to type theory, but not in a fashion that is suitable for repair of native definitions. We return to this point when developing repair for coinductive types.

Example 4.8 (Final Stream Coalgebra ✓)**.** The `Stream` type is the final coalgebra for the stream functor $P\, S =_{\text{def}} A \times S$. Its coalgebra structure is just the two destructors bundled up, unfolding a stream into its head and tail:

```
outStream : Stream X -> P X 𝟙 (Stream X)
outStream s = (hd s , λ _ -> tl s)
```

The universal property of final coalgebras is what enables the construction of coprograms. The homomorphism from a coalgebra to a final coalgebra gives coinductive types a corecursion principle, a method for defining functions from any coalgebra to a final coalgebra by specifying how the function interacts with the destructors of the coalgebra. For `Stream`, the head is read off from γ, and the tail corecurses on the unfolded position:

```
Stream-corec : (C : P-Coalg X 𝟙) -> CoalgCarrier C -> Stream X
hd (Stream-corec C c) = fst (CoalgOut C c)
tl (Stream-corec C c) = Stream-corec C (snd (CoalgOut C c) tt)
```

That `Stream-corec` C is a homomorphism holds definitionally and it is the unique such function.

The verification of finality for `Stream` requires reasoning about the equality of two streams to prove the uniqueness of a coalgebra homomorphism. Equality for the infinite systems given by coalgebra (or coinductive types) requires a notion of indistinguishable behavior that is not captured by the usual intensional identity type: that type is inductively generated by refl, so its only closed proofs identify terms that are already definitionally equal, and it offers no way to build an identification of two infinite objects from an observation-by-observation comparison. To understand equality on coalgebras, we must define an adequate behavior-preserving relation on coalgebras.

Definition 4.9 (Bisimulation ✓)**.** Given a polynomial (A, B), fix a coalgebra $(X, \langle \gamma_1, \gamma_2 \rangle)$ for $P\, A\, B$ and a span $R : X \to X \to \mathsf{Type}$. The terms $x, y : X$ satisfy the one-step condition $\mathsf{Bisim}\, R\, x\, y$ when they have equal shape observations and R-related position observations.

$$\mathsf{Bisim}\, R : (x\, y : X) \to \mathsf{Type}$$
$$\mathsf{Bisim}\, R\, x\, y =_{\text{def}} (\text{b-shape} : \gamma_1\, x = \gamma_1\, y) \times \big((b : B\, (\gamma_1\, x)) \to R\, (\gamma_2\, x\, b)\, (\gamma_2\, y\, b^*)\big)$$

where b^* indicates a transport across b-shape.

The span R is a *bisimulation* when every R-related pair satisfies this one-step condition.

$$\mathsf{isBisim} : (R : X \to X \to \mathsf{Type}) \to \mathsf{Type}$$
$$\mathsf{isBisim}\, R =_{\text{def}} \forall \{x\, y\} \to R\, x\, y \to \mathsf{Bisim}\, R\, x\, y$$

Terms $x, y : X$ are then bisimilar when they are related by some bisimulation, some R with $\mathsf{isBisim}\, R$ and $R\, x\, y$. That this notion of bisimulation is a bisimulation in the categorical sense of Rutten [52] is proven in the subsequent proposition.

The definition of bisimulation given here is the single-carrier specialization of the definition typically given [52]. We use this specialization for simplicity; our use of the definition always transports relations over one carrier to another, so we have no need for a heterogeneous relation.

Proposition 4.10 (Coinduction Principle for Final Coalgebras ✓)**.** *Fix a final coalgebra (X, γ). The coinduction principle says that bisimilar elements are equal. Given a relation R that is a bisimulation, every R-related pair is identified by a path.*

$$\begin{aligned}\mathsf{coind} : (R : X \to X \to \mathsf{Type}) \to \mathsf{isBisim}\, R \\ \to \forall\{x\, y\} \to R\, x\, y \to x = y\end{aligned}$$

Proof sketch. This is the argument of Rutten [52], repeated in type theory by Ahrens et al. [8]. The argument goes by forming the coalgebra on $\widehat{R} =_{\text{def}} \Sigma\, x\, y.\, R\, x\, y$, whose projections are coalgebra homomorphisms into (X, γ), which must be the same, giving the desired identity by evaluating at x, y. □

The proof of Proposition 4.10 relies on $\mathsf{isBisim}\, R$ giving a coalgebra structure on the total space $(\widehat{R}, \gamma_R)$ whose projections $\mathsf{fst}, \mathsf{fst} \circ \mathsf{snd} : \widehat{R} \to X$ are coalgebra homomorphisms into (X, γ). Our lifting of the span R is the canonical one for a polynomial endofunctor $P\, A\, B$, in the sense that Bisim is the weakest one-step condition from which this follows: Bisim yields such a γ_R, and any other condition that does implies Bisim ✓.

Example 4.11 (Bisimulation on Streams ✓)**.** For streams the bisimulation data specializes to a relation R relating streams whose heads agree and whose tails are R-related:

```
record StreamBis (R : Stream X -> Stream X -> Type) (s t : Stream X) where
  field
    bhd : hd s = hd t
    btl : R (tl s) (tl t)
```

Coinduction then turns any such bisimulation into a path, defined by copattern matching on the path's head and tail:

```
Stream-coind : (R : Stream X -> Stream X -> Type)
            -> (∀ {s t} -> R s t -> StreamBis R s t)
            -> ∀ {s t} -> R s t -> s = t
hd (Stream-coind R isBis r i) = bhd (isBis r) i
tl (Stream-coind R isBis r i) = Stream-coind R isBis (btl (isBis r)) i
```

The ability to construct proofs of identity via copattern matching as in Examples 4.8 and 4.11 is the principal reason for our choice of Cubical Agda. The path type of cubical type theory is not an inductively generated type; it is defined as a function space out of the interval, which facilitates introduction by means other than refl. On the other hand, the usual identity type is amenable to the sort of breaking-down of structure that pattern matching on inductive types provides. Indeed, the bisimulation and the path type are equivalent in Cubical Agda (coind is an equivalence) [57], giving a type theory where equality is suited for coalgebraic reasoning.

4.2 Coalgebraic Repair

Finally, we have the tools in hand to cast coinductive repair in terms of final coalgebras. The goal of repair is to establish a relationship between types across which programs and proofs can be repaired. Coalgebra gives us the definitive notion of a relationship between coalgebras: our goal is to define a homomorphism between the coalgebras corresponding to `Stream` and `Stream'` so that we can transfer the programs and proofs between the two. However, as in the inductive case, the proof engineer would like an efficient repair mechanism that removes the source type as a dependency. We develop this in two parts: first with a coalgebra configuration, and then with an algorithm-friendly configuration.

Definition 4.12 ((Homotopy) Coalgebra Configuration ✓)**.** Given a polynomial (A, B), a final coalgebra structure on a carrier C packages a destructor map with a proof that the resulting coalgebra is final:

$$\mathsf{FinalCoalgOn}\, A\, B\, C =_{\mathrm{def}} (\gamma : C \to P\, A\, B\, C) \times \mathsf{isFinal}\, (C, \gamma).$$

A coalgebra configuration between carriers C and D is a pair of final coalgebra structures for the same polynomial, one on each carrier:

$$\mathsf{CoalgConfig}\, A\, B\, C\, D =_{\mathrm{def}} \mathsf{FinalCoalgOn}\, A\, B\, C \;\times\; \mathsf{FinalCoalgOn}\, A\, B\, D.$$

A coalgebra configuration determines a coalgebra isomorphism: each side's finality yields a unique homomorphism into the other, and the two round-trips are forced to be the identities by uniqueness. A homotopy coalgebra configuration is given by replacing the finality condition with isHFinal.

Now, we can use such a configuration to do the swap induced by the underlying equivalence on programs and proofs, rather than compose with the equivalence.

Example 4.13 (Coalgebraic Repair with Stream and Stream' ✓)**.** The configuration equips each carrier with the same universal property, so a program written with the corecursor and a proof written with the coinduction principle are repaired by swapping the universal property at which they are instantiated. The defining data of cons mentions its carrier only through the structure map, so it can be stated once, against an abstract coalgebra (E, out):

```
consSeed : P-Coalg X 𝟙 -> P-Coalg X 𝟙
consSeed (E , out) = ((X × E) ⊎ E , γ)
  where γ (inl (y , t)) = (y , λ _ -> inr t)
        γ (inr t)     = (fst (out t) , λ _ -> inr (snd (out t) tt))

cons-M : X -> Stream X -> Stream X
cons-M x s = Stream-corec (consSeed (Stream X , outStream )) (inl (x , s))

cons'-M : X -> Stream' X -> Stream' X
cons'-M x s = Stream'-corec (consSeed (Stream' X , outStream')) (inl (x , s))

streamComp-M : (s : Stream X) -> cons (hd s) (tl s) = s
streamComp-M s = Stream-coind etaR etaStep (refl , refl)

streamComp'-M : (s : Stream' X) -> cons' (fst (obs s)) (snd (obs s)) = s
streamComp'-M s = Stream'-coind etaR' etaStep' (refl , refl)
```

The coalgebra consSeed starts in state `inl (x , s)`, emits x, and then moves to `inr s`, from which it replays s through out. This mirrors the observation after Proposition 3.15: occurrences of the source type in defining data are further uses of its (co)algebra structure, repaired by the same swap. The bisimulations etaR/etaR′ likewise read their carriers only through the structure maps. The repaired program and proof carry no reference to the source type.

4.3 Coinductive Repair via Configurations

With the coalgebraic perspective on coinductive types, we can now develop the framework for repair algorithms for coinductive types. Akin to the standard presentation of W-types, we give coinductive types a set of type-theoretic rules, and refer to types satisfying these rules as M-types. The state-of-the-art work on M-types (in the formulation of rules) has not established their homotopy-finality. In fact, the contemporary work has focused on proving the finality and homotopy-finality for various formulations of M-types [8, 20, 26], none of which are rule-based presentations. We are primarily interested in developing repair for coinductive types, and thus are concerned with rule-based presentations that facilitate programming and proving. To develop repair for contemporary and future type theories, our development handles both a notion of M-types that corresponds to final coalgebras and a stronger notion that corresponds to homotopy-final coalgebras.

$$\frac{s : C \to A \qquad p : (c : C) \to B\,(s\,c) \to C}{\mathrm{corec}\,s\,p : C \to \mathsf{M}\,A\,B}\ \mathrm{corec} \qquad \frac{}{\beta_s : \mathrm{shape}\,(\mathrm{corec}\,s\,p\,c) = s\,c}\ \beta_s$$

$$\frac{}{\beta_p : \mathrm{pos}\,(\mathrm{corec}\,s\,p\,c)\ =_{\beta_s}\ \mathrm{corec}\,s\,p \circ (p\,c)}\ \beta_p$$

$$\frac{\begin{array}{c} h : C \to \mathsf{M}\,A\,B \\ \beta_s^h : (c : C) \to \mathrm{shape}\,(h\,c) = s\,c \qquad \beta_p^h : (c : C) \to \mathrm{pos}\,(h\,c) =_{\beta_s^h\,c} h \circ (p\,c) \end{array}}{\eta : h = \mathrm{corec}\,s\,p}\ \eta$$

$$\frac{h : C \to \mathsf{M}\,A\,B \qquad \beta_h : (c : C) \to \big(\mathrm{shape}\,(h\,c) = s\,c\big) \times \big(\mathrm{pos}\,(h\,c) = h \circ (p\,c)\big)}{(h, \beta_h)\ =\ (\mathrm{corec}\,s\,p,\ \beta)}\ \eta_\Sigma$$

$$\frac{R : \mathsf{M}\,A\,B \to \mathsf{M}\,A\,B \to \mathrm{Type} \qquad b : \mathrm{isBisim}\,R\,\langle \mathrm{shape}, \mathrm{pos}\rangle \qquad r : R\,x\,y}{\mathrm{coind}\,R\,b\,r : x = y}\ \mathrm{coind}$$

Fig. 1. Rules of the (homotopy) M-type structure over a polynomial (A, B): the corecursor corec, its computation rules β_s and β_p, the function-level uniqueness rule η and its Σ-level strengthening η_Σ, and the derived coinduction principle coind.

An M-type for a polynomial (A, B) is a carrier, written $\mathsf{M}\,A\,B$, equipped with destructors

$$\mathrm{shape} : \mathsf{M}\,A\,B \to A \qquad\qquad \mathrm{pos} : (x : \mathsf{M}\,A\,B) \to B\,(\mathrm{shape}\,x) \to \mathsf{M}\,A\,B$$

satisfying the rules corec, β_s, β_p, and η of Figure 1. The corecursor corec maps any coalgebra $(C, \langle s, p\rangle)$ into $\mathsf{M}\,A\,B$, and its two computation (β) rules say that destructing a corecursive value computes: β_s for the shape and β_p for the positions, which lives over β_s and is hence a dependent path. The uniqueness (η) rule states that any map agreeing with the source coalgebra under the destructors, that is, any coalgebra homomorphism, equals the corecursor. Strengthening η from the function level to the Σ level gives η_Σ, which we use for homotopy M-types (Definition 4.15). Every M-type also enjoys a coinduction principle coind, which says that any bisimulation entails a path; unlike the other rules, it is not assumed but derived (Definition 4.14).

Definition 4.14 (M-Types ✓)**.** The type HasMRules bundles the rules of M-types (Figure 1):

$$\begin{aligned}
&\mathrm{HasMRules} : (X : \mathrm{Type}) \to (\mathrm{shape} : X \to A) \to ((x : X) \to B\,(\mathrm{shape}\,x) \to X) \to \mathrm{Type} \\
&\mathrm{HasMRules}\,X\,\mathrm{shape}\,\mathrm{pos} =_{\mathrm{def}} \\
&\quad (\mathrm{corec} : \{C : \mathrm{Type}\}\,(s : C \to A)\,(p : (c : C) \to B\,(s\,c) \to C) \to C \to X) \\
&\quad \times (\beta_s : \forall\{C\,s\,p\,c\} \to \mathrm{shape}\,(\mathrm{corec}\,s\,p\,c) = s\,c) \\
&\quad \times (\beta_p : \forall\{C\,s\,p\,c\} \to \mathrm{pos}\,(\mathrm{corec}\,s\,p\,c) =_{\beta_s} \mathrm{corec}\,s\,p \circ (p\,c)) \\
&\quad \times \big(\eta : \forall\{C\,s\,p\}\,(h : C \to X) \to (\beta_s^h : (c : C) \to \mathrm{shape}\,(h\,c) = s\,c) \\
&\qquad \to ((c : C) \to \mathrm{pos}\,(h\,c) =_{\beta_s^h\,c} h \circ (p\,c)) \to h = \mathrm{corec}\,s\,p\big)
\end{aligned}$$

An *M-type* is a carrier with destructors for which HasMRules is inhabited. As we will see (Theorem 4.21), M-types are carriers of final coalgebras, and, as such, enjoy a coinduction principle, so we do not include a coinduction principle as a field of HasMRules.

We emphasize that M-types are not a primitive of our ambient theory: HasMRules is a property that a given carrier with given destructors may or may not enjoy. There exists an M-type for

any polynomial (A, B), and any M-type has a coalgebra given by $\langle \text{shape}, \text{pos} \rangle$, which enables the formation of bisimulations on M-types, as in the derived coind principle.

Concretely, we use the coinductive records of Cubical Agda to model these M-types and verify that they satisfy the given rules ✓.

```
record M (A : Type) (B : A -> Type) : Type where
  coinductive
  field
    shape : A
    pos   : B shape -> M A B

corec : (C : Type) (s : C -> A) (p : (c : C) -> B (s c) -> C) -> C -> M A B
shape (corec C s p c) = s c
pos   (corec C s p c) b = corec C s p (p c b)

M-HasMRules : HasMRules (M A B) shape pos
M-HasMRules = isFinal→HasMRules M-isFinal

M-coind = HasMRules-coind M-HasMRules
```

Now, we develop M-types that directly correspond to homotopy-final coalgebras.

Definition 4.15 (Homotopy M-Types ✓)**.** A homotopy M-type strengthens the η rule of an M-type from the function level to the Σ level: a homomorphism (h, β_h) agreeing with the source coalgebra under the destructors equals the corecursor bundle $(\text{corec}\, s\, p, \beta)$ by a Σ-path (η_Σ in Figure 1). Exactly as HasMRules bundles the plain rules, we bundle the M-rules with η_Σ in place of η into a single type HasHMRules X shape pos.

A homotopy M-type differs from an M-type only in this one strengthened rule. Note that while the existence of homotopy M-types in HoTT is known via the limit construction of Ahrens et al. [8], our development and prior work do not prove that they exist within Cubical Agda. Proving that the native `M` record is a homotopy M-type requires characterizing the path types of `M` at higher dimensions, which the state of the art [20, 26] does not do; indeed, our attempted formalization ran into a documented issue [1]. The obstruction disappears under additional hypotheses on the polynomial. For example, `Stream` A is a set whenever A is a set ✓. The same fact holds for M-types whose shape type is a set, but we do not pursue the general statement here.

Configurations expose the syntactic detail of a particular representation that a repair algorithm walks over. We give a configuration framework for coinductive types by focusing on destructors and outputs. We introduce the destructor-level signature and algebra, then define a single carrier's coinductive structure, and finally give a configuration as a pair of these.

Definition 4.16 (Destructor Algebra ✓)**.** The output telescope of a single destructor is described by a *destructor arity* DestrArity, generated by the same constructors Done, Nonrec, and Rec as ConstrArity (Definition 2.1). Reading an arity off at a carrier X gives the *outputs* Outputs $cs\, X$ of such a destructor, defined clause for clause as Args, with functorial action mapOutputs as for mapArgs. A CoSignature *sig* is a type Op of destructor labels together with an arity assignment arity : Op → DestrArity. Where a constructor algebra consumes arguments, a *destructor algebra* for *sig* on a carrier X interprets every label as a destructor exposing its outputs:

$$\text{DestrAlgebra}\, sig\, X =_{\text{def}} (op : \text{Op}) \to \big(X \to \text{Outputs}\, (\text{arity}\, op)\, X\big).$$

Definition 4.17 (Coinductive Configuration ✓)**.** A *coinductive coalgebra* carries a destructor algebra equipped with corecursion and computation rules. Formally, a CoindCoalg for *sig* on X packages:

- destr : DestrAlgebra *sig* X, the syntactic coalgebra on X;
- corec : DestrAlgebra *sig* $C \to C \to X$ for any source carrier C, the corecursor;

- corec-β, the per-destructor computation rule: for every source coalgebra δ and label op,
$$\mathsf{destr}\ op\ (\mathsf{corec}\ \delta\ c)\ =\ \mathsf{mapOutputs}\ (\mathsf{arity}\ op)\ (\mathsf{corec}\ \delta)\ (\delta\ op\ c);$$
- corec-η, function-level uniqueness: any $h : C \to X$ whose destructors agree with δ agrees with $\mathsf{corec}\ \delta$.

Again, in Cubical Agda the coinduction principle coind is derivable from these fields, so CoindCoalg already carries it implicitly. A *coinductive configuration* for *sig* between carriers C and D is a pair of coinductive coalgebras, one on each carrier:
$$\mathsf{CoConfig}\ sig\ C\ D =_{\mathrm{def}} \mathsf{CoindCoalg}\ sig\ C\ \times\ \mathsf{CoindCoalg}\ sig\ D.$$
Coinductive configurations admit a homotopy strengthening just like the M-type formulations.

Definition 4.18 (Homotopy Coinductive Configuration ✓)**.** An HCoindCoalg replaces the function-level corec-η with the Σ-level uniqueness rule η_Σ of a homotopy M-type (corec-uniq-Σ). A homotopy coinductive configuration is a pair of these:
$$\mathsf{HCoConfig}\ sig\ C\ D =_{\mathrm{def}} \mathsf{HCoindCoalg}\ sig\ C\ \times\ \mathsf{HCoindCoalg}\ sig\ D.$$
As with inductive configurations, we can verify that the coinductive configurations satisfy the correctness conditions a proof engineer would expect of a repair algorithm. These conditions are new correctness criteria for coinductive repair, but they follow the idea that the configuration behaves coherently with the underlying equivalence.

Proposition 4.19 (Correctness Conditions ✓)**.** *Fix a coinductive configuration* CoConfig *sig C D with destructor coalgebras* $\mathsf{destr}_C, \mathsf{destr}_D$*, corecursors* $\mathsf{corec}_C, \mathsf{corec}_D$*, and β laws. Let $\varphi : C \to D$ be the forward map of the equivalence induced by the configuration. It satisfies the coinductive duals of the PUMPKIN Pi coherence conditions.*

(1) destr_ok. *The destructors commute with φ: for every label op and $c : C$,*
$$\mathsf{destr}_D\ op\ (\varphi\ c)\ =\ \mathsf{mapOutputs}\ (\mathsf{arity}\ op)\ \varphi\ (\mathsf{destr}_C\ op\ c).$$
Since $\varphi = \mathsf{corec}_D\ \mathsf{destr}_C$ by construction, this is exactly its per-destructor β rule.

(2) iota_ok. *Each corecursor computes on its own destructors: this is exactly* corec-β*, one law for* corec_C *and one for* corec_D.

(3) corec_ok. *The corecursors commute with φ: for every source coalgebra ec* : DestrAlgebra *sig E and $e : E$,*
$$\varphi\ (\mathsf{corec}_C\ ec\ e)\ =\ \mathsf{corec}_D\ ec\ e.$$

(4) coind_ok. *The coinduction principles commute with φ: writing* $\mathsf{coind}_C, \mathsf{coind}_D$ *for coinduction on each side, for every bisimulation R on C and $r : R\ x\ y$,*
$$\mathsf{ap}\ \varphi\ (\mathsf{coind}_C\ r)\ =\ \mathsf{coind}_D\ (\varphi_* r),$$
*an equality of the parallel paths $\varphi\ x = \varphi\ y$, where φ_* transports the C-bisimulation to a D-bisimulation. It only holds for a homotopy configuration (*HCoConfig*).*

Proof sketch. The first three follow directly from the structure the configuration carries; the interesting part is coind_ok. It lifts the argument of Proposition 4.10 one dimension: a two-dimensional bisimulation R_2 on parallel paths $p, q : a = b$ yields $p = q$ ✓. The coalgebra now lives on $\Sigma\ a\ b\ (p\ q : a = b).\ R_2\ p\ q$, and in place of two projections there is, for each $k : \mathbb{I}$, a pair of homomorphisms sending (a, b, p, q, r) to $p\ k$ and $q\ k$. The Σ-level uniqueness of corecursors (corec-uniq-Σ) identifies all of these with the corecursor, coherently at $k = 0, 1$, which yields the square $p = q$ and is why coind_ok needs the homotopy tier. Applying this on D to $\mathsf{ap}\ \varphi\ (\mathsf{coind}_C\ r)$ and $\mathsf{coind}_D\ (\varphi_* r)$, which are related by a lift of R, gives coind_ok. □

Beyond generalizing Rutten's argument, the two-dimensional principle holds for any carrier satisfying the interface HasHMRules, not just a particular construction of M-types such as the

limit construction of Ahrens et al. [8], and it makes no truncation assumption. Working at the level of the interface is crucial for a theory of repair, which must apply to the coinductive types a proof assistant provides natively rather than to one chosen construction. The fourth condition, coind_ok, is the one condition that demands the homotopy tier. Streams over a set assemble into an HCoConfig ✓, so coind_ok holds in particular for the running `Stream`/`Stream'` example whenever the element type is a set.

The corec_ok condition is the coinductive form of repair by instantiation (Proposition 3.15), the dual of $\mathrm{fold}_C\,\gamma = \mathrm{fold}_D\,\gamma \circ \varphi$. It can take the source coalgebra *ec* verbatim because *ec* lives on a third type E and never mentions the repaired carriers, just as a fold's target algebra γ does not. By contrast, the data of elim_ok (a motive over C) and of coind_ok (a bisimulation on C) mention the carrier, and so must be transported along φ.

Example 4.20 (Stream and Stream' Repair via Configurations ✓)**.** Now we can phrase our running example through coinductive configurations.

```
StreamSig : CoSignature
StreamSig = MTypeSignature X 𝟙

streamCoindCoalg : CoindCoalg StreamSig (Stream X)
streamCoindCoalg = coalgBridgeBwd (outStream , Stream-isFinal)
stream'CoindCoalg : CoindCoalg StreamSig (Stream' X)
stream'CoindCoalg = coalgBridgeBwd (outStream' , Stream'-isFinal)

streamCoConfig : CoConfig StreamSig (Stream X) (Stream' X)
streamCoConfig = record { coindCoalgC = streamCoindCoalg; coindCoalgD = stream'CoindCoalg }
```

Given any source coalgebra, `unfold ind` corecurses into that component's carrier. Written against the abstract CoindCoalg, it names neither `Stream` nor `Stream'`.

```
unfold : CoindCoalg StreamSig E -> DestrAlgebra StreamSig S -> S -> E
unfold ind = CoindCoalg.corec ind
```

Repair is now executed through a swap of the configuration's components. `consVia` writes `cons` against an abstract component. Instantiating at `streamCoindCoalg` yields `cons` on `Stream`; the repair is the *same program* instantiated at the swapped component, now landing in `Stream'`.

```
consVia : CoindCoalg StreamSig E -> X -> E -> E
consVia ind x e = unfold ind δ (inl (x , e))
  where δ op (inl (y , t)) = y , (λ _ -> inr t) , tt
        δ op (inr t)    = mapOutputs _ inr (destr ind op t)

consCfg : X -> Stream X -> Stream X
consCfg = consVia streamCoindCoalg
consCfg' : X -> Stream' X -> Stream' X
consCfg' = consVia stream'CoindCoalg
```

The proof `streamComp` is repaired the same way, through the coinduction principle each component carries (Definition 4.17).

```
streamCompCfg : (s : Stream X) -> cons (hd s) (tl s) = s
streamCompCfg s = deriveCoind streamCoindCoalg etaR etaRIsBisim (refl , refl)

streamCompCfg' : (s : Stream' X) -> cons' (fst (obs s)) (snd (obs s)) = s
streamCompCfg' s = deriveCoind stream'CoindCoalg etaR' etaRIsBisim' (refl , refl)
```

This models the execution of a proof repair algorithm for coinductive types on the running programs and proofs. These repaired terms compute as we would expect. Observing one reduces the configuration scaffolding away, leaving native `Stream'` data. Each equation below holds definitionally.

```
consCfg'-head : (x : X) (s : Stream' X) -> fst (obs (consCfg' x s)) = x
consCfg'-next : (x : X) (s : Stream' X) -> fst (obs (snd (obs (consCfg' x s)))) = fst (obs s)
```

4.4 Connecting Coinductive Configurations and Coalgebra

We now characterize M-types as carriers of final coalgebras, setting the foundation for cementing our slogan that proof repair is an algebra isomorphism in the case of coinductive repair. We give the coalgebraic dual of Awodey et al.'s [12] characterization of W-types as homotopy-initial algebras. The novelty is the left-hand side of the equivalence below, a rule-based presentation of (homotopy) M-types, whereas prior treatments establish finality for coalgebraic or limit constructions of M-types [8, 20, 26], not for such an interface.

Theorem 4.21 ((Homotopy) M-Types are (Homotopy) Final Coalgebras)**.** *Given a polynomial* (A, B) *and a carrier* X *with destructors* $(s : X \to A,\ p : (x : X) \to B\,(s\,x) \to X)$*, having the (homotopy) M-rules is exactly the induced coalgebra* $(X, \langle s, p\rangle)$ *being (homotopy-)final.*

(1) (Final ✓) $\mathsf{HasMRules}\,X\,s\,p \ \simeq\ \mathsf{isFinal}\,(X, \lambda x.\,(s\,x, p\,x))$

(2) (Homotopy-final ✓) $\mathsf{HasHMRules}\,X\,s\,p \ \simeq\ \mathsf{isHFinal}\,(X, \lambda x.\,(s\,x, p\,x))$

Proof sketch. The corecursor together with its two β rules is exactly a coalgebra homomorphism out of any source, and η is exactly function-level uniqueness of that homomorphism, and conversely corec, β_s, β_p, and η are read back off the finality data. Both round-trips hold definitionally. The homotopy-final tier strengthens η from function-level uniqueness to contractibility of the coalgebra-homomorphism space. Contractibility transports either way and both sides are propositions, so the bi-implication is again an equivalence. □

On the inductive side, the correspondence between W-types and initial algebras requires technical higher-dimensional path algebra because the W-type data has a dependent eliminator that quantifies over all motives. For the coinductive story, the correspondence given by Theorem 4.21 is much easier to establish because of the absence of a dependent co-eliminator; instead, the difficulty is relocated to deriving that extra data in the proof of coind_ok.

Coinductive configurations have a canonical form we call *M-type configurations*. These configurations emulate M-types by capturing the coinductive structure of types with a single destructor that exposes a non-recursive and a recursive output.

Definition 4.22 (M-type Configuration ✓)**.** We define a one-destructor M-type signature: a single observation whose output telescope exposes a shape $a : A$ and then a $B\,a$-indexed family of recursive positions.

$$\mathsf{MTypeArity} : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to \mathbb{1} \to \mathsf{DestrArity}$$
$$\mathsf{MTypeArity}\,A\,B\,_ =_{\mathrm{def}} \mathsf{Nonrec}\,A\,(\lambda a \to \mathsf{Rec}\,(B\,a)\,\mathsf{Done})$$

$$\mathsf{MTypeSignature} : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to \mathsf{CoSignature}$$
$$\mathsf{MTypeSignature}\,A\,B\,.\mathsf{CoSignature.Op} =_{\mathrm{def}} \mathbb{1}$$
$$\mathsf{MTypeSignature}\,A\,B\,.\mathsf{CoSignature.arity} =_{\mathrm{def}} \mathsf{MTypeArity}\,A\,B$$

$$\mathsf{MTypeCoindCoalg} : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to (C : \mathsf{Type}) \to \mathsf{Type}$$
$$\mathsf{MTypeCoindCoalg}\,A\,B\,C =_{\mathrm{def}} \mathsf{CoindCoalg}\,(\mathsf{MTypeSignature}\,A\,B)\,C$$

An M-type configuration is the special coinductive configuration whose signature is an M-type signature, collapsing all per-destructor data into one destructor with one non-recursive output and one recursive output:

$$\mathsf{MTypeCoConfig} : (A : \mathsf{Type}) \to (A \to \mathsf{Type}) \to (C\,D : \mathsf{Type}) \to \mathsf{Type}$$
$$\mathsf{MTypeCoConfig}\,A\,B\,C\,D =_{\mathrm{def}} \mathsf{CoConfig}\,(\mathsf{MTypeSignature}\,A\,B)\,C\,D$$

M-type configurations collapse the detail of a configuration into the canonical one-destructor form over which we identify configurations with coalgebra isomorphisms. Every coinductive configuration is faithfully represented by an M-type configuration.

Lemma 4.23 (Coinductive Configurations Faithfully Embed into M-type Configurations ✓)**.** *Every coinductive configuration over an arbitrary coinductive signature sig determines an M-type configuration on the same carriers. We read off the polynomial from the signature:*

$$A \;=_{def}\; (op : \mathsf{Op}) \to \mathsf{Outputs}(\mathsf{arity}\, op)\ \mathbb{1}, \qquad B\, a \;=_{def}\; (op : \mathsf{Op}) \times \mathsf{recPos}(\mathsf{arity}\, op)\ (a\, op),$$

and the configuration converts to an M-type configuration for this (A, B)*:*

$$\mathsf{CoConfig}\ sig\ C\ D \to (A : \mathsf{Type}) \times ((B : (A \to \mathsf{Type})) \times \mathsf{MTypeCoConfig}\ A\ B\ C\ D).$$

The embedding is faithful: the forward and backward maps of the repair equivalence induced by the M-type configuration are equal to those induced by the original configuration.

We now state our main theorem over M-type configurations, using the coinductive dual of Definition 3.6 in two tiers matching the two notions of finality. A *coalgebra equivalence* $\mathsf{CoalgEquiv}\ X\ Y$ ✓ is a type equivalence of the carriers whose forward and backward maps both carry coalgebra-homomorphism witnesses. Its homotopy counterpart, $\mathsf{CoalgEquiv}^{\mathsf{H}}$ ✓, is the exact dual of Definition 3.6.

Theorem 4.24 (Coinductive Configurations are Coalgebra Isomorphisms ✓)**.** *Over the M-type signature, a coinductive configuration is a coalgebra isomorphism whose source is final.*

(1) (Final) A configuration is logically equivalent to a coalgebra isomorphism between final coalgebras:

$$\mathsf{CoConfig}\ (\mathsf{MTypeSignature}\ A\ B)\ C\ D \;\longleftrightarrow\; (\gamma_C : C \to P\ A\ B\ C) \times (\gamma_D : D \to P\ A\ B\ D) \times \mathsf{CoalgEquiv}\ (C, \gamma_C)\ (D, \gamma_D) \times \mathsf{isFinal}\ (C, \gamma_C)$$

(2) (Homotopy-final) A homotopy configuration is a coalgebra isomorphism between homotopy-final coalgebras:

$$\mathsf{HCoConfig}\ (\mathsf{MTypeSignature}\ A\ B)\ C\ D \;\simeq\; (\gamma_C : C \to P\ A\ B\ C) \times (\gamma_D : D \to P\ A\ B\ D) \times \mathsf{CoalgEquiv}^{\mathsf{H}}\ (C, \gamma_C)\ (D, \gamma_D) \times \mathsf{isHFinal}\ (C, \gamma_C)$$

Proof sketch. The first link is a carrier-level bridge given by our characterization of (homotopy) M-types as (homotopy-)final coalgebras (Theorem 4.21). We then identify a pair of final coalgebras with one final coalgebra plus a coalgebra isomorphism to the other. The two directions are mutually inverse exactly when the matched data is propositional. This holds at the homotopy-final tier, giving an equivalence. It fails at the final tier, where isFinal does not pin the homomorphism witness. □

This theorem is the coinductive completion of our slogan; a coinductive configuration witnesses that its carriers are final for a common polynomial functor. Repair between coinductive types is therefore natural precisely when the types have natural presentations as final coalgebras for the same functor, and the correspondence used for repair is the coalgebra isomorphism that finality forces to exist uniquely between them. At the homotopy-final tier a coalgebra isomorphism out of a final source is a proposition (as the proof above observes), so a homotopy coinductive configuration, too, determines a unique repair map. The changes handled by coinductive configurations are exactly those representable by a coalgebra isomorphism, the changes of destructor presentation that preserve the observable behavior of the type.

This development largely rests upon the discovery that proof repair for inductive types corresponds to algebra isomorphisms. Before the connection between repair and P-algebras was established, it was unclear how to extend proof repair to coinductive types, since the question was muddled by the many different implementation approaches taken by various proof assistants, none of which isolates what a coinductive type is. By focusing on the universal properties that coinductive types have, we are able to construct a high-level foundation for proof repair algorithms in a general manner, independent of any specific implementation approach.

5 Related Work

Proof Repair. Our work gives the first foundational treatment of proof repair. Proof repair was first introduced in 2018 [49, 51]. It is the proof engineering [47] analogue of program repair [10, 37]. Some of the early examples of proof repair [50] were based on the theory of algebraic ornaments [18, 19, 31, 36], which express a relation between indexed datatypes where one index type has richer information than the other. The class of proof repair algorithms that our foundations support comes from a line of work on automatically repairing proofs in Rocq in response to breaking changes in datatypes [46, 48, 50, 58]. That line of work focuses on algorithmic approaches, but leaves open questions about what a configuration represents and what makes a configuration "natural." The thesis work in particular [46] presents a diagram that gives intuition that inspired our foundations, but goes no further in developing that intuition. We formalize this line of work and answer these open questions. Supplementary automation beyond the core repair functionality, like decompilation, is out of the scope of our foundations.

Other proof repair algorithms exist that are out of the scope of our foundations. For example, Sisyphus [25] repairs separation logic proofs about imperative programs written in OCaml. Proof Mate [35] repairs proofs written in PVS. Preliminary work explores proof repair for JML [39]. All of these tools concern specification and proof languages that are quite different from the dependent type theory that we concern ourselves with.

Recently, neural methods for proof repair have begun to emerge, both as standalone tools [41, 42, 44, 45, 55, 59, 61] and as a way of augmenting other models for proofs [23, 30, 43]. The scope of these methods is defined by their training data, rather than by any foundational system. In the long run, foundations like ours may help in defining kinds of data that these tools might wish to support, or in building well-scoped and general neurosymbolic tools that integrate symbolic repair into other neural models (in the fashion of PALM [34]).

Proof Reuse and Transfer. Proof repair is one way of reusing proofs. The key conditions that yield repair rather than the more general reuse are that (1) proof repair is primarily concerned with adapting proofs to changes rather than with reusing proofs across developments, and (2) proof repair is most useful in the presence of imperfect foresight or breaking changes in dependencies that are outside of one's control. Giving up the first condition yields tools for proof reuse, like refinement frameworks [17] and transfer methods [16, 28, 53, 54]; giving up the second condition yields design principles that aim to make proofs reusable or robust to changes to begin with [15, 21, 47, 60].

The relationships between these tools and technologies have been discussed extensively in prose [16, 46, 47, 58]. Our foundations add a formal dimension to this discussion. By formalizing the connection between equivalences, configurations, and algebras, we draw an explicit bridge between proof repair and transfer methods whose inner workings use automatic transport. Moreover, with Proposition 3.15, we learn that proof repair satisfies its distinguishing conditions by exploiting a kind of representation independence [9]. This property may also hold of other "white-box" approaches to automatic transfer [54], which can be used for repair.

(Co)algebras and (Co)inductive Types. We briefly cover previous work on (co)algebras and (co)inductive types from which our development greatly benefited, with no claim of completeness.

Inductive types are a well-researched area within programming languages and type theory. The seminal work of Dybjer [22] gives a schema for inductive and recursive definitions within Martin-Löf type theory (MLTT). The work of Chapman et al. [14] presents a dependent type theory which gives inductive definitions via an encoding in a universe. The connection between inductive datatypes, algebras, and polynomial endofunctors has received significant attention. The primary inspiration for our work is the discovery by Awodey et al. [12] that the rules for W-types are equivalent to the existence of homotopy-initial algebras for polynomial endofunctors. Abbott et

al. [4] show that the notion of containers (similar to polynomial functors) can be used to represent all strictly positive types and the polymorphic functions between them. Damato et al. [20] generalize the container development to a setting without the uniqueness of identity proofs axiom, giving a formulation in intensional type theory and formalization in Cubical Agda.

Coinductive types and coalgebras are a prominent and growing research area. The work of Hagino [27] uses categorical logic to develop a simply typed lambda calculus with algebraic and coalgebraic datatypes. Basold and Geuvers [13] extend Hagino's categorical datatypes to a dependent type theory constructed solely from inductive and coinductive types, while leaving the (co)induction principles out. Ahrens et al. [8] establish that the existence of M-types (in the formulation of homotopy-final coalgebras) is derivable in HoTT in the presence of inductive types. In the course of developing models of non-wellfounded sets in HoTT, Gylterud et al. [26] characterize the identity type of the M-type (formulated as a homotopy-final coalgebra) as an indexed M-type. Most modern proof assistants feature some form of coinductive types, made possible by various research programs [5, 24, 33].

6 Conclusion

We have given the first foundational treatment of type-theoretic proof repair. Its fundamental construct, the configuration, is exactly an algebra isomorphism between initial algebras for a common polynomial functor (Theorem 3.11). This discovery greatly clarifies the foundational understanding of repair: the coherence conditions imposed by prior tools are consequences of initiality rather than side obligations, the repair equivalence a configuration determines is unique, and programs written against the universal property are repaired by re-instantiation alone (Proposition 3.15). Dualizing this account yielded the foundations of repair for a class of types no existing tool handles—coinductive types—together with new correctness criteria for coinductive repair and a novel characterization of (homotopy) M-types as (homotopy-)final coalgebras (Theorem 4.21).

Two aspects of the development deserve emphasis. First, none of it uses univalence: the inductive side lives in intensional MLTT with function extensionality, and the coinductive side uses Cubical Agda only for its path type, not for univalent principles. Since prior repair work was framed univalently, this broadens the class of type theories in which repair is known to be justified. Second, the algebraic framing is extensible. Higher inductive types, quotient inductive types, inductive-inductive types, and inductive-recursive types all admit algebraic presentations over richer signatures, and we expect configurations, their coherence conditions, and the uniqueness of repair to transfer along the same lines. We hope this foundation serves both as an explanation of why existing repair tools work and as a blueprint for the repair tools that do not yet exist.

Acknowledgments

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR0011262E019. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA. This work is also supported by the UK Advanced Research + Invention Agency (ARIA).

Disclosure of Use of AI in Research

We used AI tools (specifically, Claude) in our research for two tasks: (1) the standard refactoring of Agda proofs, and (2) making sense of the implications of Proposition 3.15. For the second task, after we had most of the main results, we asked Claude whether any of them could explain why repair can eliminate old types. This is a property of repair that had captured our curiosity, but for

which we had thus far failed to find a crisp, formal explanation. Claude responded by drawing a connection between our result of Proposition 3.15 and this property.

Data Availability Statement

The artifact for our paper is available on Zenodo.[2] It includes the Agda formalization of all of our results, as well as the Cubical Agda examples from Section 4. An explanation of the artifact and how to install it can be found in `artifact/README.md`. The explanation also includes a centralized table relating our informal and formal results (as an alternative presentation of the blue checkmarks throughout the paper that link to the formal results directly). Our data is additionally available in a public GitHub repository.[3]

References

[1] 2020. *cubical primitives should preserve guardedness · Issue #4740 · agda/agda.* https://github.com/agda/agda/issues/4740

[2] 2026. *Propositional Equality.* https://lean-lang.org/doc/reference/latest/Basic-Propositions/Propositional-Equality/#UIP

[3] 2026. *The Rocq Prover 9.2 documentation.* https://rocq-prover.org/doc/V9.2.0/refman/index.html

[4] Michael Abbott, Thorsten Altenkirch, and Neil Ghani. 2005. Containers: Constructing strictly positive types. *Theoretical Computer Science* 342, 1 (2005), 3–27. doi:10.1016/j.tcs.2005.06.002 Applied Semantics: Selected Topics.

[5] Andreas Abel, Brigitte Pientka, David Thibodeau, and Anton Setzer. 2013. Copatterns: programming infinite structures by observations. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Rome, Italy) *(POPL '13)*. Association for Computing Machinery, New York, NY, USA, 27–38. doi:10.1145/2429069.2429075

[6] Peter Aczel and Nax Mendler. 1989. A final coalgebra theorem. In *Category Theory and Computer Science* (Berlin, Heidelberg, 1989), David H. Pitt, David E. Rydeheard, Peter Dybjer, Andrew M. Pitts, and Axel Poigné (Eds.). Springer, 357–365. doi:10.1007/BFb0018361

[7] Jiří Adámek, Stefan Milius, and Lawrence S. Moss. 2025. *Initial Algebras and Terminal Coalgebras: The Theory of Fixed Points of Functors.* Cambridge University Press.

[8] Benedikt Ahrens, Paolo Capriotti, and Régis Spadotti. 2015. Non-Wellfounded Trees in Homotopy Type Theory. In *13th International Conference on Typed Lambda Calculi and Applications (TLCA 2015) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 38)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 17–30. doi:10.4230/LIPIcs.TLCA.2015.17

[9] Carlo Angiuli, Evan Cavallo, Anders Mörtberg, and Max Zeuner. 2021. Internalizing representation independence with univalence. *Proceedings of the ACM on Programming Languages* 5, POPL, Article 12 (2021). doi:10.1145/3434293

[10] Ashif Anwar. 2026. Neuro Symbolic Automated Program Repair: A Systematic Review of LLM-Based and Symbolic Techniques. *International Journal of Computer Engineering and Data Science (IJCEDS)* 5, 1 (Mar. 2026), 1–16. https://www.ijceds.com/ijceds/article/view/118

[11] Steve Awodey. 2010. *Category Theory* (2nd ed.). Oxford University Press, Inc., USA.

[12] Steve Awodey, Nicola Gambino, and Kristina Sojakova. 2017. Homotopy-Initial Algebras in Type Theory. 63, 6 (2017), 1–45. doi:10.1145/3006383

[13] Henning Basold and Herman Geuvers. 2016. Type Theory based on Dependent Inductive and Coinductive Types. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science* (New York, NY, USA) *(LICS '16)*. Association for Computing Machinery, New York, NY, USA, 327–336. doi:10.1145/2933575.2934514

[14] James Chapman, Pierre-Évariste Dagand, Conor McBride, and Peter Morris. 2010. The gentle art of levitation. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming* (Baltimore, Maryland, USA) *(ICFP '10)*. Association for Computing Machinery, New York, NY, USA, 3–14. doi:10.1145/1863543.1863547

[15] Adam Chlipala. 2013. *Certified Programming with Dependent Types - A Pragmatic Introduction to the Coq Proof Assistant.* MIT Press. http://mitpress.mit.edu/books/certified-programming-dependent-types

[16] Cyril Cohen, Enzo Crance, and Assia Mahboubi. 2024. Trocq: proof transfer for free, with or without univalence. In *European Symposium on Programming*. Springer, 239–268. doi:10.1007/978-3-031-57262-3_10

[17] Cyril Cohen, Maxime Dénès, and Anders Mörtberg. 2013. Refinements for Free!. In *Certified Programs and Proofs*, Georges Gonthier and Michael Norrish (Eds.). Springer International Publishing, Cham, 147–162. doi:10.1007/978-3-319-03545-1_10

[2] https://doi.org/10.5281/zenodo.23326858

[3] https://github.com/trey3p/proof-repair-coalgebra

[18] Pierre-Evariste Dagand and Conor McBride. 2012. Transporting functions across ornaments. In *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming* (Copenhagen, Denmark) *(ICFP '12)*. Association for Computing Machinery, New York, NY, USA, 103–114. doi:10.1145/2364527.2364544

[19] Pierre-Evariste Dagand and Conor McBride. 2013. A Categorical Treatment of Ornaments. In *Proceedings of the 2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '13)*. IEEE Computer Society, USA, 530–539.

[20] Stefania Damato, Thorsten Altenkirch, and Axel Ljungström. 2025. Formalising Inductive and Coinductive Containers. doi:10.48550/arXiv.2409.02603 arXiv:2409.02603 [cs.LO]

[21] Benjamin Delaware, William Cook, and Don Batory. 2011. Product Lines of Theorems. In *Proceedings of the 2011 ACM International Conference on Object Oriented Programming Systems Languages and Applications* (Portland, Oregon, USA) *(OOPSLA 2011)*. ACM, New York, NY, USA, 595–608. doi:10.1145/2048066.2048113

[22] Peter Dybjer. 1994. Inductive families. *Form. Asp. Comput.* 6, 4 (July 1994), 440–465. doi:10.1007/BF01211308

[23] Emily First, Markus N. Rabe, Talia Ringer, and Yuriy Brun. 2023. Baldur: Whole-Proof Generation and Repair with Large Language Models. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2023)*. Association for Computing Machinery, New York, NY, USA, 1229–1241. doi:10.1145/3611643.3616243

[24] Eduardo Giménez. 1994. Codifying Guarded Definitions with Recursive Schemes. In *Selected Papers from the International Workshop on Types for Proofs and Programs (TYPES '94)*. Springer-Verlag, Berlin, Heidelberg, 39–59.

[25] Kiran Gopinathan, Mayank Keoliya, and Ilya Sergey. 2023. Mostly Automated Proof Repair for Verified Libraries. *Proc. ACM Program. Lang.* 7, PLDI, Article 107 (June 2023), 25 pages. doi:10.1145/3591221

[26] Hakon Robbestad Gylterud, Elisabeth Stenholm, and Niccolò Veltri. 2026. Terminal Coalgebras and Non-wellfounded Sets in Homotopy Type Theory. Volume 22, Issue 2 (2026), 14325. doi:10.46298/lmcs-22(2:35)2026 arXiv:2001.06696 [math.LO]

[27] Tatsuya Hagino. 1987. A Typed Lambda Calculus with Categorical Type Constructors. In *Category Theory and Computer Science*. Springer-Verlag, Berlin, Heidelberg, 140–157.

[28] Brian Huffman and Ondřej Kunčar. 2013. Lifting and Transfer: A Modular Design for Quotients in Isabelle/HOL. In *Certified Programs and Proofs*, Georges Gonthier and Michael Norrish (Eds.). Springer International Publishing, Cham, 131–146. doi:10.1007/978-3-319-03545-1_9

[29] Bart Jacobs and Jan Rutten. 2011. *An introduction to (co)algebra and (co)induction.* Cambridge University Press, 38–99.

[30] Prithwish Jana, Kaan Kale, Ahmet Ege Tanriverdi, Cruise Song, Sriram Vishwanath, and Vijay Ganesh. 2026. ProofBridge: Auto-Formalization of Natural Language Proofs in Lean via Joint Embeddings. arXiv:2510.15681 [cs.LO] https://arxiv.org/abs/2510.15681

[31] HSIANG-SHANG KO and JEREMY GIBBONS. 2016. Programming with ornaments. *Journal of Functional Programming* 27 (2016), e2. doi:10.1017/S0956796816000307

[32] DEXTER KOZEN and ALEXANDRA SILVA. 2017. Practical coinduction. *Mathematical Structures in Computer Science* 27, 7 (2017), 1132–1152. doi:10.1017/S0960129515000493

[33] Bastiaan Laarakker, Daniël Otten, and Benno van den Berg. 2026. Constructing (Co)inductive Types via Large Sizes. doi:10.48550/arXiv.2602.18921 arXiv:2602.18921 [cs.LO]

[34] Minghai Lu, Benjamin Delaware, and Tianyi Zhang. 2024. Proof Automation with Large Language Models. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering* (Sacramento, CA, USA) *(ASE '24)*. Association for Computing Machinery, New York, NY, USA, 1509–1520. doi:10.1145/3691620.3695521

[35] Paolo Masci and Aaron Dutle. 2022. Proof Mate: An Interactive Proof Helper for PVS (Tool Paper). In *NASA Formal Methods: 14th International Symposium, NFM 2022, Pasadena, CA, USA, May 24–27, 2022, Proceedings* (Pasadena, CA, USA). Springer-Verlag, Berlin, Heidelberg, 809–815. doi:10.1007/978-3-031-06773-0_44

[36] Conor McBride. 2011. Ornamental algebras, algebraic ornaments. https://plv.mpi-sws.org/plerg/papers/mcbride-ornaments-2up.pdf

[37] Martin Monperrus. 2018. Automatic Software Repair: A Bibliography. *ACM Comput. Surv.* 51, 1, Article 17 (Jan. 2018), 24 pages. doi:10.1145/3105906

[38] Leonardo de Moura and Sebastian Ullrich. 2021. The Lean 4 Theorem Prover and Programming Language. In *Automated Deduction – CADE 28: 28th International Conference on Automated Deduction, Virtual Event, July 12–15, 2021, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 625–635. doi:10.1007/978-3-030-79876-5_37

[39] Amirfarhad Nilizadeh, Gary T. Leavens, and David R. Cok. 2023. Automated Reasoning Repair. In *Proceedings of the 24th ACM International Workshop on Formal Techniques for Java-like Programs* (Berlin, Germany) *(FTfJP '22)*. Association for Computing Machinery, New York, NY, USA, 11–14. doi:10.1145/3611096.3611099

[40] Ulf Norell. 2009. Dependently typed programming in Agda. In *Proceedings of the 4th International Workshop on Types in Language Design and Implementation* (Savannah, GA, USA) *(TLDI '09)*. Association for Computing Machinery, New York, NY, USA, 1–2. doi:10.1145/1481861.1481862

[41] Azim Ospanov, Farzan Farnia, and Roozbeh Yousefzadeh. 2026. APOLLO: Automated LLM and Lean Collaboration for Advanced Formal Reasoning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*. https://openreview.net/forum?id=fxDCgOruk0

[42] Manooshree Patel, Bartosz Piotrowski, Leopold Haller, and Hugh James Leather. 2026. ProofRepairBench: Exploring Proof Repair in Lean. In *ICLR 2026 Workshop: VerifAI-2: The Second Workshop on AI Verification in the Wild*. https://openreview.net/forum?id=6SwWVNwEJK

[43] Santhana Srinivasan R and Maithilee Patawar. 2026. LAMP: Lean-based Agentic framework with MCP and Proof Repair. arXiv:2606.28841 [cs.LO] https://arxiv.org/abs/2606.28841

[44] Thomas Reichel. 2024. *Neural approaches to theorem search & proof repair*. Master's thesis. University of Illinois at Urbana-Champaign.

[45] Tom Reichel, R. Wesley Henderson, Andrew Touchet, Andrew Gardner, and Talia Ringer. 2023. Proof Repair Infrastructure for Supervised Models: Building a Large Proof Repair Dataset. In *14th International Conference on Interactive Theorem Proving (ITP 2023) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 268)*, Adam Naumowicz and René Thiemann (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 26:1–26:20. doi:10.4230/LIPIcs.ITP.2023.26

[46] Talia Ringer. 2021. *Proof Repair*. Ph. D. Dissertation. University of Washington. https://digital.lib.washington.edu/researchworks/handle/1773/47429

[47] Talia Ringer, Karl Palmskog, Ilya Sergey, Milos Gligoric, and Zachary Tatlock. 2019. QED at Large: A Survey of Engineering of Formally Verified Software. *Found. Trends Program. Lang.* 5, 2–3 (Sept. 2019), 102–281. doi:10.1561/2500000045

[48] Talia Ringer, RanDair Porter, Nathaniel Yazdani, John Leo, and Dan Grossman. 2021. Proof repair across type equivalences. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (Virtual, Canada) *(PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 112–127. doi:10.1145/3453483.3454033

[49] Talia Ringer, Nathaniel Yazdani, John Leo, and Dan Grossman. 2018. Adapting Proof Automation to Adapt Proofs. 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, 115–129. doi:10.1145/3167094

[50] 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 für Informatik, Dagstuhl, Germany, 26:1–26:19. doi:10.4230/LIPIcs.ITP.2019.26

[51] Valentin Robert. 2018. *Front-end tooling for building and maintaining dependently-typed functional programs*. Ph. D. Dissertation. UC San Diego.

[52] J. J. M. M. Rutten. 2000. Universal coalgebra: a theory of systems. 249, 1 (2000), 3–80. doi:10.1016/S0304-3975(00)00056-6

[53] Nicolas Tabareau, Éric Tanter, and Matthieu Sozeau. 2018. Equivalences for Free: Univalent Parametricity for Effective Transport. *Proc. ACM Program. Lang.* 2, ICFP, Article 92 (jul 2018), 29 pages. doi:10.1145/3236787

[54] Nicolas Tabareau, Éric Tanter, and Matthieu Sozeau. 2021. The Marriage of Univalence and Parametricity. *J. ACM* 68, 1, Article 5 (jan 2021), 44 pages. doi:10.1145/3429979

[55] Amer Tahat, David Hardin, Adam Petz, and Perry Alexander. 2024. Proof repair utilizing large language models: A case study on the copland remote attestation proofbase. In *International Conference on Bridging the Gap between AI and Reality*. Springer, 145–166.

[56] The Univalent Foundations Program. 2013. *Homotopy Type Theory: Univalent Foundations of Mathematics*. https://homotopytypetheory.org/book, Institute for Advanced Study.

[57] Andrea Vezzosi, Anders Mörtberg, and Andreas Abel. 2019. Cubical agda: a dependently typed programming language with univalence and higher inductive types. *Proc. ACM Program. Lang.* 3, ICFP, Article 87 (July 2019), 29 pages. doi:10.1145/3341691

[58] Cosmo Viola, Max Fan, and Talia Ringer. 2025. Proof Repair across Quotient Type Equivalences. *Proc. ACM Program. Lang.* 9, OOPSLA2, Article 386 (Oct. 2025), 26 pages. doi:10.1145/3763164

[59] Evan Wang, Simon Chess, Daniel Lee, Siyuan Ge, Ajit Mallavarapu, Jarod Alper, and Vasily Ilin. 2026. Learning to Repair Lean Proofs from Compiler Feedback. arXiv:2602.02990 [cs.LG] https://arxiv.org/abs/2602.02990

[60] Doug Woos, James R. Wilcox, Steve Anton, Zachary Tatlock, Michael D. Ernst, and Thomas Anderson. 2016. Planning for Change in a Formal Verification of the Raft Consensus Protocol. In *Proceedings of the 5th ACM SIGPLAN Conference on Certified Programs and Proofs* (St. Petersburg, FL, USA) *(CPP 2016)*. ACM, New York, NY, USA, 154–165. doi:10.1145/2854065.2854081

[61] Jun Yang, Yuechun Sun, Yi Wu, Rodrigo Caridad, Yongwei Yuan, Jianan Yao, Shan Lu, and Kexin Pei. 2026. ExVerus: Verus Proof Repair via Counterexample Reasoning. In *Forty-third International Conference on Machine Learning*. https://openreview.net/forum?id=FNDkXA0OUJ