

Formal Metatheory of Second-Order Abstract Syntax

MARCELO FIORE* and DMITRIJ SZAMOZVANCEV†, University of Cambridge, UK

Despite extensive research both on the theoretical and practical fronts, formalising, reasoning about, and implementing languages with variable binding is still a daunting endeavour – repetitive boilerplate and the overly complicated metatheory of capture-avoiding substitution often get in the way of progressing on to the actually interesting properties of a language. Existing developments offer some relief, however at the expense of inconvenient and error-prone term encodings and lack of formal foundations.

We present a mathematically-inspired language-formalisation framework implemented in Agda. The system translates the description of a syntax signature with variable-binding operators into an intrinsically-encoded, inductive data type equipped with syntactic operations such as weakening and substitution, along with their correctness properties. The generated metatheory further incorporates metavariables and their associated operation of metasubstitution, which enables second-order equational/rewriting reasoning. The underlying mathematical foundation of the framework – initial algebra semantics – derives compositional interpretations of languages into their models satisfying the semantic substitution lemma by construction.

CCS Concepts: • **Software and its engineering** → **Formal software verification**; **Functional languages**; • **Theory of computation** → **Type theory**; *Equational logic and rewriting*.

Additional Key Words and Phrases: Agda, abstract syntax, language formalisation, category theory

1 INTRODUCTION

Programming and logic research papers that introduce and study new languages and calculi with variable binding typically gloss over the associated notion of capture-avoiding substitution – it is often taken as standard along with its correctness properties. Nevertheless, the representation of variables and substitution becomes a major roadblock when one attempts to formalise such languages in proof assistants, as substituting into terms with variable binders involves dealing with variable shadowing and capture. Some ways of encoding variables (e.g. using strings or numeric de Bruijn indices) lead to fragile and error-prone implementations of syntactic operations: for instance, binding links are easily broken if one increments the wrong index or forgets to modify a string. Other approaches – such as type- and scope-safe de Bruijn indices – place static guarantees on the correctness of the implementation but require significant boilerplate and seemingly superfluous generality that puts the interesting metatheoretic proofs on hold. Either way, the end result is a formalisation littered with ad hoc definitions and lemmas about weakening, substitution, etc., all of which need to be tweaked if the syntax is changed even slightly.

We tackle this unsatisfactory state of affairs by stepping back and considering the mathematical foundations of *second-order abstract syntax*; that is, abstract syntax with variable binding and parametrised metavariables. Using an approach derived from the presheaf model developed by [Fiore, Plotkin, and Turi \[1999\]](#) and [Fiore \[2008\]](#), we build a powerful and entirely generic framework for language formalisation that extracts the maximum amount of syntactic metatheory from a second-order syntax description with minimal boilerplate and user effort. From a concise textual specification of a typed syntax with binding operators, such as the following one for the simply-typed lambda calculus (STLC)

*Research partially supported by EPSRC grant EP/V002309/1.

†Research supported by EPSRC grant EP/R513180/1.

type	term
$N : 0\text{-ary}$	$\text{app} : \alpha \succ \beta \quad \alpha \rightarrow \beta$
$_ \succ _ : 2\text{-ary}$	$\text{lam} : \alpha.\beta \quad \rightarrow \alpha \succ \beta$

with base type N and function type \succ , our system generates Agda code for: (i) a grammar of types and an intrinsically-typed data type of terms; (ii) operations of weakening and substitution together with their correctness properties; (iii) a record that, when instantiated with a mathematical model, induces a semantic interpretation of the syntax in the model that preserves substitution; (iv) a term constructor for parametrised metavariables and their associated operation of metasubstitution; and (v) an equational/rewriting theory that can be instantiated with the axioms of the syntax to obtain a library for second-order equational/rewriting reasoning.

1.1 Background

The framework presented in this paper builds on a long series of practical and theoretical developments in the study of abstract syntax, incorporating elements from both lines of research.

Intrinsic typing. Among the numerous language-formalisation strategies available for a variety of proof assistants, the type- and scope-safe, intrinsically-typed encoding of typed syntax with variable binding popularised by Altenkirch and Reus [1999], Benton et al. [2012], and, more recently, by Allais, Atkey, Chapman, McBride, and McKinna [2021] (henceforth cited as AACMM [2021]) stands out as an excellent fit for dependently-typed metalanguages like Agda [Norell 2009] and Coq [The Coq Development Team 2004]. It represents terms of a syntax as a type- and context-indexed family \mathcal{X} of sets: for a type α and typing context (list of types) Γ , the set $\mathcal{X} \alpha \Gamma$ consists of the terms t that satisfy the typing judgment $\Gamma \vdash t : \alpha$. In an intrinsically-typed formalisation this indexed family \mathcal{X} is generated inductively: for example, application in the syntax of the STLC corresponds to a constructor app that combines a function term $f : \mathcal{X} (\alpha \succ \beta) \Gamma$ and an argument term $a : \mathcal{X} \alpha \Gamma$ into $\text{app}(f, a) : \mathcal{X} \beta \Gamma$. As such, constructors of the syntax are directly encoded as their own typing rules, making ill-typed and ill-scoped terms unrepresentable.

In contrast, an extrinsic encoding resembles the way type systems are presented in the research literature: one has a grammar of raw terms t and an inductively-defined well-typedness relation $\Gamma \vdash t : \alpha$. While both approaches work and are widely used, dependent types and intrinsic encoding complement each other very well. For instance, in the introductory textbook *Programming Language Foundations in Agda*, Kokke, Siek, and Wadler [2020] present the formalisation of the STLC in both styles, clearly highlighting and advocating the superiority of intrinsic typing.

Another advantage of intrinsic typing is that the need for separate type-preservation proofs of syntactic operations is eliminated: for example, defining the single-variable substitution operation $\text{sub}_1 : \mathcal{X} \alpha \Gamma \rightarrow \mathcal{X} \beta (\alpha \cdot \Gamma) \rightarrow \mathcal{X} \beta \Gamma$ immediately establishes the fact that substitution preserves typing. Just as the typing rules of terms are baked into their syntax, the well-typedness proofs of operations are baked into their definition. The downside is that implementing such operations often requires significant effort, with sub_1 being a prime example: it is not definable from first principles, because the induction hypothesis (viz. the recursive call) is not general enough to handle binding terms, whose fresh bound variables get added to the end of the typing context and thereby “cover” the free variable one wants to substitute for. Instead, someone unfamiliar with the approach has to walk through a long and frustrating path of trial and error. For instance: trying to define weakening and exchange (which are usually employed in the pen-and-paper well-typedness proof) is similarly futile; generalising to substitution for a variable in the middle of the context works for binding terms but not for variables; further generalising to simultaneous substitution still requires weakening; single-variable weakening cannot be implemented directly either,

and must be derived from variable renaming. Even worse is that reasoning about single-variable substitution (proving the syntactic or semantic substitution lemmas, for example) forces one to prove similar properties about renaming, weakening, and simultaneous substitution, leading to a tedious and bloated formalisation that is specific to a particular syntax.

Generic traversals. If we accept the hoops needed to jump through to implement single-variable substitution, we can look for common patterns to abstract out – after all, renaming and substitution both proceed by traversing a term and replacing variables with *data*, either a variable or a term. McBride [2005] builds on this observation to axiomatise the properties that such *data* need possess in order to define a generic term-traversal function that can be instantiated to either renaming or substitution. Allais et al. [2017] generalise traversals to semantic interpretations in models of the syntax, recovering renaming and substitution as interpretations into the syntactic model. They also give a pattern for proving simulation and fusion properties of traversals, which capture – among other things – the correctness properties of renaming and substitution. In AACMM [2021], the authors further generalise their work on the STLC to an arbitrary universe of second-order syntaxes, producing a robust and flexible language-formalisation framework.

Presheaf model. The presheaf model of second-order abstract syntax is a category-theoretic approach to enriching typed languages with variable binding by means of substitution and metasubstitution structures, developed by Fiore and collaborators; see e.g. [Fiore et al. 1999; Fiore 2008; Fiore and Hamana 2013; Fiore and Hur 2010; Fiore and Mahmoud 2010]. It too represents syntax as a type- and context-indexed family of sets, but further requires it to have a functorial renaming structure upon which substitution, metasubstitution, etc. structures are built in a systematic categorical way. While it is an elegant and powerful mathematical theory, its adoption in language-formalisation frameworks has only been incidental: several concepts central to the abstract categorical approach (coends, colimits, etc.) are hard to represent in a dependently-typed setting and a faithful reproduction of the theory would need to heavily rely on a formalised category-theory library such as `agda-categories` [Hu and Carette 2021] or `UniMath` [Voevodsky et al. 2014].

Recent work by Borthelle et al. [2020] proposed that the presheaf model be adapted to a setting of sorted families by axiomatising the renaming structure and its interaction with initial algebra semantics – ideas which have been explored informally by Allais et al. [2017] and Kaiser et al. [2018] as well. To note, however, is that their skew-monoidal [Szlachányi 2012] approach still depends on categorical theory and results that lead to an impractical formalisation: for example, the syntax of terms is presented by means of the colimit of an ω -chain (and therefore represented by equivalence classes), rather than by an inductive data type.

1.2 Contributions

Our work aims to place the constructions and empirical observations of intrinsically-typed term encodings on a formal grounding, thereby bridging the gap between theory and practice. We are closely guided by the presheaf model of second-order abstract syntax, adapting and re-working its categorical theory (definitions and results) to a dependently-typed setting without sacrificing practicality and generality over arbitrary second-order signatures.

- In Section 2 we give the basic mathematical framework of sorted families, renaming coalgebras, and substitution monoids employed in the library. These notions are adapted from the presheaf model [Fiore et al. 1999] which, on its own, is too abstract and high-level to be used as a basis for practical language formalisation. However, with this new (and mathematically justifiable) change of perspective, we can translate all the theory of the presheaf approach to type- and

context-indexed families and, in doing so, shine a light upon the constructions and properties commonly encountered in type- and scope-safe treatments of syntax.

- **Section 3** outlines the main conceptual tool we rely on: initial algebra semantics. Crucially, not only do we invoke initiality to derive generic traversals such as renaming and substitution, we also use it to prove their correctness, including the associativity of substitution and its interaction with renaming. The notion of a Σ -monoid is introduced; it axiomatises the structure of a model of the syntax where terms can be interpreted in a compositional and substitution-preserving manner. We also discuss metavariables and metasubstitution [Fiore 2008], together with their application to second-order equational/rewriting reasoning [Fiore and Hur 2010].
- After developing the entirely signature-generic metatheory, in **Section 4** we present the translation of second-order syntax descriptions to signature endofunctors and give two approaches of encoding the initial algebra of terms for a signature.
- Finally, in **Section 5**, we showcase the features of our framework in two extended examples: the syntax of the STLC together with its sound denotational and operational semantics, and the equational theory of partial differentiation.

The paper is intended for an audience familiar with dependently-typed proof assistants and the struggles of language formalisation. Even though they play a vital role in the development, we chose to minimise references to advanced categorical concepts. In the listings we use “retouched Agda” which omits inessential implementation details to make the code cleaner, without impacting clarity (for example, we hide some implicit arguments and record declarations, and use shading instead of `cong` in equational proofs). The full formalisation with extended proofs and examples can be found on the [project website](#).

2 MATHEMATICAL FOUNDATIONS

We start by setting up the abstract mathematical foundations of our work, with the aim of decoupling the renaming and substitution structures from a particular syntax of terms. **Section 2.1** lists the standard definitions for contexts, families, and variables, highlighting their important categorical properties. In **Section 2.2** we introduce the coalgebraic and monoidal views of renaming and substitution, and in **Section 2.3** we consider properties of maps parametrised by a substitution mapping which will be important players in the development of initial algebra semantics.

2.1 Contexts and families

The definitions of contexts and variables are well-known from intrinsically-typed treatments of syntax. Rather than using named variable-type pairs, contexts are lists of types that come from a fixed set T , and variables are typed and scoped de Bruijn indices into the context: `new` points to the first element of the context, while `old`(v) points to the variable v in an extended context.

```

data Ctx : Set where
  () : Ctx
  _:_ : (α : T) → (Γ : Ctx) → Ctx

data I : T → Ctx → Set where
  new : I α (α · Γ)
  old : I β Γ → I β (α · Γ)

```

We call context-indexed sets of type $\text{Ctx} \rightarrow \text{Set}$ *families*, and sort-indexed families (such as I above) a *sorted family*: for an arbitrary sorted family $\mathcal{X} : T \rightarrow \text{Ctx} \rightarrow \text{Set}$, an element in the set $\mathcal{X} \alpha \Gamma$ can be seen as an \mathcal{X} -term of type α with free variables in context Γ . Maps between sorted families are sort- and context-indexed families of functions, and together they form the category \mathbf{Fam}_s .

```

Family : Set1      Familys : Set1      _→_ : Familys → Familys → Set
Family = Ctx → Set  Familys = T → Family  X → Y = {α : T}{Γ : Ctx} → X α Γ → Y α Γ

```

The concatenation of two contexts is denoted $\Gamma \dot{+} \Delta$. An associated endofunctor on families is that of *context extension*: a $(\delta \Theta \mathcal{X})$ -term in context Γ is an \mathcal{X} -term in the extended context $\Theta \dot{+} \Gamma$.

$$\begin{array}{l} \underline{\dot{+}}_{_} : \mathbf{Ctx} \rightarrow \mathbf{Ctx} \rightarrow \mathbf{Ctx} \\ \emptyset \quad \dot{+} \Delta = \Delta \\ (\alpha \cdot \Gamma) \dot{+} \Delta = \alpha \cdot (\Gamma \dot{+} \Delta) \end{array} \qquad \begin{array}{l} \delta : \mathbf{Ctx} \rightarrow \mathbf{Family}_s \rightarrow \mathbf{Family}_s \\ \delta \Theta \mathcal{X} \alpha \Gamma = \mathcal{X} \alpha (\Theta \dot{+} \Gamma) \end{array}$$

2.1.1 Cocartesian structure of contexts. Fundamental to the algebraic study of abstract syntax is the *category of contexts* \mathbb{F} , with contexts Γ, Δ as objects and type-preserving mappings between variables of two contexts as morphisms. We will call these morphisms *renamings*; intuitively, a renaming $\rho : \Gamma \rightsquigarrow \Delta$ assigns a variable $\rho(v) : \alpha$ in context Δ to each variable $v : \alpha$ in Γ . Renamings can be generalised to arbitrary *context maps* $\sigma : \Gamma \dashv\vdash [\mathcal{X}] \rightarrow \Delta$ that assign an \mathcal{X} -term $\sigma(v) : \mathcal{X} \alpha \Delta$ to each variable $v : \alpha$ in Γ – a renaming is then expressed as an \mathcal{I} -valued context map.

$$\begin{array}{l} \underline{\dashv\vdash}_{_} : \mathbf{Ctx} \rightarrow \mathbf{Family}_s \rightarrow \mathbf{Ctx} \rightarrow \mathbf{Set} \\ \Gamma \dashv\vdash [\mathcal{X}] \rightarrow \Delta = \{\alpha : T\} \rightarrow \mathcal{I} \alpha \Gamma \rightarrow \mathcal{X} \alpha \Delta \end{array} \qquad \begin{array}{l} \underline{\rightsquigarrow}_{_} : \mathbf{Ctx} \rightarrow \mathbf{Ctx} \rightarrow \mathbf{Set} \\ \Gamma \rightsquigarrow \Delta = \Gamma \dashv\vdash [\mathcal{I}] \rightarrow \Delta \end{array}$$

The category of contexts \mathbb{F} is *cocartesian*, with context concatenation $\dot{+}$ serving as coproduct. The injection maps $\Gamma \rightsquigarrow (\Gamma \dot{+} \Delta)$ and $\Delta \rightsquigarrow (\Gamma \dot{+} \Delta)$ and the universal copairing (definable for arbitrary \mathcal{X} -valued context maps) $\mathbf{copair} : (\Gamma \dashv\vdash [\mathcal{X}] \rightarrow \Theta) \rightarrow (\Delta \dashv\vdash [\mathcal{X}] \rightarrow \Theta) \rightarrow (\Gamma \dot{+} \Delta) \dashv\vdash [\mathcal{X}] \rightarrow \Theta$ can be constructed by pattern-matching on the contexts and variables, and repeated applications of **old**. A useful special case of copairing is adding a single term to a substitution, which corresponds to the standard *cons* operation in the theory of explicit substitutions [Abadi et al. 1989].

$$\begin{array}{l} \mathbf{add} : (\mathcal{X} : \mathbf{Family}_s) \rightarrow \mathcal{X} \alpha \Delta \rightarrow \Gamma \dashv\vdash [\mathcal{X}] \rightarrow \Delta \rightarrow (\alpha \cdot \Gamma) \dashv\vdash [\mathcal{X}] \rightarrow \Delta \\ \mathbf{add} \mathcal{X} t \sigma = \mathbf{copair} \mathcal{X} (\lambda \{ \mathbf{new} \rightarrow t \}) \sigma \end{array}$$

2.1.2 Bicartesian closed and linear structures of families. Since families are indexed sets, they inherit the bicartesian closed structure of **Set**: we take sums and products of families via pointwise disjoint unions and Cartesian products, and also define family exponentials pointwise as $(\mathcal{X} \Rightarrow \mathcal{Y}) \alpha \Gamma \triangleq \mathcal{X} \alpha \Gamma \rightarrow \mathcal{Y} \alpha \Gamma$ (since there are no functoriality requirements, relevant in the case of presheaf exponentials, to be satisfied). We also have the following *linear exponential* (derived from the Day [1970] internal hom) which will play an important role in metasubstitution:

$$\begin{array}{l} \underline{\dashv\vdash}_{_} : \mathbf{Family}_s \rightarrow \mathbf{Family}_s \rightarrow \mathbf{Family}_s \\ (\mathcal{X} \dashv\vdash \mathcal{Y}) \alpha \Gamma = \{\Delta : \mathbf{Ctx}\} \rightarrow \mathcal{X} \alpha \Delta \rightarrow \mathcal{Y} \alpha (\Delta \dot{+} \Gamma) \end{array}$$

2.2 Renaming and substitution

Our next step is to precisely identify the structure required on a sorted family to support substitution; this way, defining substitution for a particular syntax will amount to equipping the family of terms of the syntax with such structure. Since, in practice, capture-avoiding substitution into syntactic terms involves weakening (a form of variable renaming), we are also looking for a way to axiomatise renaming as a basic fundamental notion. Our guiding principle throughout the development will be to characterise renaming, substitution, and other metatheoretic operations – along with their correctness laws – as natural, recognisable categorical constructions in sorted families.

2.2.1 Renaming structure. A sorted family \mathcal{X} admits renaming when \mathcal{X} -terms in one variable context can be compatibly transformed to \mathcal{X} -terms in another. The corresponding renaming operation lifts a renaming map $\Gamma \rightsquigarrow \Delta$ to a function $\mathcal{X} \alpha \Gamma \rightarrow \mathcal{X} \alpha \Delta$. This amounts to the requirement that the family $\mathcal{X} \alpha$ be a (covariant) *presheaf* on \mathbb{F} (i.e. a functor from \mathbb{F} to **Set**), and imposing such structure on every family would give rise to the presheaf model of Fiore et al. [1999]. Here, instead of going down this route, we refactor the type of the renaming operation on families as follows:

$$\mathbf{r} : \{\alpha : \mathbb{T}\}\{\Gamma : \mathbf{Ctx}\} \rightarrow \mathcal{X} \alpha \Gamma \rightarrow (\{\Delta : \mathbf{Ctx}\} \rightarrow (\Gamma \rightsquigarrow \Delta) \rightarrow \mathcal{X} \alpha \Delta)$$

Note that the codomain of \mathbf{r} can be expressed as a function of α and Γ . By abstracting the renaming-dependence as a *modal operator* \square [Allais et al. 2021], we may then express the renaming operation as a map of families $\mathcal{X} \rightarrow \square \mathcal{X}$, internalising the functorial action as a \mathbf{Fam}_s -morphism.

$$\begin{aligned} \square & : \mathbf{Family}_s \rightarrow \mathbf{Family}_s \\ \square \mathcal{X} \alpha \Gamma & = \{\Delta : \mathbf{Ctx}\} \rightarrow (\Gamma \rightsquigarrow \Delta) \rightarrow \mathcal{X} \alpha \Delta \end{aligned}$$

The \square modality is a comonad and a coalgebra for it has to respect the identity and composition of renamings. This leads to our first observation: a sorted family has renaming structure when equipped with a \square -coalgebra structure. The **Coalg** record collects these requirements.

$$\begin{aligned} \text{record } \mathbf{Coalg} (\mathcal{X} : \mathbf{Family}_s) : \mathbf{Set} \text{ where} \\ \text{field } \mathbf{r} & : \mathcal{X} \rightarrow \square \mathcal{X} \\ \text{counit} & : \{t : \mathcal{X} \alpha \Gamma\} \rightarrow \mathbf{r} t \text{ id} \equiv t \\ \text{comult} & : \{\rho : \Gamma \rightsquigarrow \Delta\}\{\varrho : \Delta \rightsquigarrow \Theta\}\{t : \mathcal{X} \alpha \Gamma\} \rightarrow \mathbf{r} t (\varrho \circ \rho) \equiv \mathbf{r} (\mathbf{r} t \rho) \varrho \end{aligned}$$

Since context transformations such as weakening, contraction, etc. correspond to renaming maps, the associated structural rules for a \square -coalgebra \mathcal{X} can be derived via renaming:

$$\begin{aligned} \mathbf{wkl} : \mathcal{X} \alpha \Gamma \rightarrow \mathcal{X} \alpha (\Gamma \dot{+} \Delta) & \quad \mathbf{wkr} : \mathcal{X} \alpha \Delta \rightarrow \mathcal{X} \alpha (\Gamma \dot{+} \Delta) & \quad \mathbf{contr} : \mathcal{X} \alpha (\Gamma \dot{+} \Gamma) \rightarrow \mathcal{X} \alpha \Gamma \\ \mathbf{wkl} t = \mathbf{r} t (\text{inl } \Delta) & \quad \mathbf{wkr} t = \mathbf{r} t (\text{inr } \Gamma) & \quad \mathbf{contr} t = \mathbf{r} t (\text{copair } \mathcal{I} \text{ id id}) \end{aligned}$$

The natural notion of a transformation between \square -coalgebras is a *homomorphism*: a map $\mathcal{X} \rightarrow \mathcal{Y}$ that preserves the coalgebra structures of \mathcal{X} and \mathcal{Y} .

$$\begin{aligned} \text{record } \mathbf{Coalg} \Rightarrow (\mathcal{X}^\square : \mathbf{Coalg} \mathcal{X})(\mathcal{Y}^\square : \mathbf{Coalg} \mathcal{Y})(f : \mathcal{X} \rightarrow \mathcal{Y}) : \mathbf{Set} \text{ where} \\ \text{field } \langle \mathbf{r} \rangle : \{\rho : \Gamma \rightsquigarrow \Delta\}\{t : \mathcal{X} \alpha \Gamma\} \rightarrow f (\mathcal{X}.\mathbf{r} t \rho) \equiv \mathcal{Y}.\mathbf{r} (f t) \rho \end{aligned}$$

2.2.2 Pointed structure. If \mathcal{X} is a sorted family of terms for a second-order abstract syntax, there must be a way to coerce variables into \mathcal{X} -terms. Sorted families with such a coercion $\eta : \mathcal{I} \rightarrow \mathcal{X}$ will be called *pointed*. If the underlying family of a \square -coalgebra is pointed, we may also impose the requirement that the point is compatible with renaming. We characterise pointed coalgebras and their point-preserving homomorphisms as records:

$$\begin{aligned} \text{record } \mathbf{Coalg}_* (\mathcal{X} : \mathbf{Family}_s) : \mathbf{Set} \text{ where} & \quad \text{record } \mathbf{Coalg}_* \Rightarrow (\mathcal{X}_*^\square : \mathbf{Coalg}_* \mathcal{X})(\mathcal{Y}_*^\square : \mathbf{Coalg}_* \mathcal{Y}) \\ \text{field } \square : \mathbf{Coalg} \mathcal{X} \ ; \ \eta : \mathcal{I} \rightarrow \mathcal{X} & \quad (f : \mathcal{X} \rightarrow \mathcal{Y}) : \mathbf{Set} \text{ where} \\ \mathbf{r} \circ \eta : \{v : \mathcal{I} \alpha \Gamma\}\{\rho : \Gamma \rightsquigarrow \Delta\} \rightarrow & \quad \text{field } \square \Rightarrow : \mathbf{Coalg} \Rightarrow \mathcal{X}_*^\square \mathcal{Y}_*^\square f \\ \mathbf{r} (\eta v) \rho \equiv \eta (\rho v) & \quad \langle \eta \rangle : \{v : \mathcal{I} \alpha \Gamma\} \rightarrow f (\mathcal{X}.\eta v) \equiv \mathcal{Y}.\eta v \end{aligned}$$

An example of a pointed \square -coalgebra is the family of variables, with renaming implemented as application. The point η of a pointed \square -coalgebra is itself a pointed \square -coalgebra homomorphism.

2.2.3 Substitution structure. The \square modality parametrises a family by a renaming. We now generalise this by parametrising a family \mathcal{Y} by an arbitrary \mathcal{X} -valued context map. This is captured as a binary operation on families which we call the *internal substitution hom* of \mathcal{X} and \mathcal{Y} :

$$\begin{aligned} \llbracket _ , _ \rrbracket & : \mathbf{Family}_s \rightarrow \mathbf{Family}_s \rightarrow \mathbf{Family}_s \\ \llbracket \mathcal{X} , \mathcal{Y} \rrbracket \alpha \Gamma & = \{\Delta : \mathbf{Ctx}\} \rightarrow (\Gamma \dashv \llbracket \mathcal{X} \rrbracket \rightarrow \Delta) \rightarrow \mathcal{Y} \alpha \Delta \end{aligned}$$

A renaming operation has, equivalently, the type $\mathcal{X} \rightarrow \llbracket \mathcal{I} , \mathcal{X} \rrbracket$, and a substitution operation on \mathcal{X} has the type $\mathcal{X} \rightarrow \llbracket \mathcal{X} , \mathcal{X} \rrbracket$. The internal hom has a left adjoint \odot , called the *substitution tensor product*; that is, maps of the form $\mathcal{X} \rightarrow \llbracket \mathcal{Y} , \mathcal{Z} \rrbracket$ are naturally isomorphic to maps of the form $\mathcal{X} \odot \mathcal{Y} \rightarrow \mathcal{Z}$ via “uncurrying”. These operators equip \mathbf{Fam}_s with a skew-monoidal closed structure

[Street 2013] (full monoidality would require that e.g. the unitor $\mathcal{I} \odot \mathcal{X} \rightarrow \mathcal{X}$ is invertible, which is not the case since \odot is defined as a dependent sum, rather than a coend [Fiore et al. 1999]).

$$\begin{aligned} _ \odot _ &: \text{Family}_s \rightarrow \text{Family}_s \rightarrow \text{Family}_s \\ (\mathcal{X} \odot \mathcal{Y}) \alpha \Delta &= \Sigma[\Gamma \in \text{Ctx}] (\mathcal{X} \alpha \Gamma \times (\Gamma -[\mathcal{Y}] \rightarrow \Delta)) \end{aligned}$$

Expressed using the tensor product, the substitution operation has the type $\mathcal{X} \odot \mathcal{X} \rightarrow \mathcal{X}$ – it combines a term $\mathcal{X} \alpha \Gamma$ and substitution map $\Gamma -[\mathcal{X}] \rightarrow \Delta$ into a term $\mathcal{X} \alpha \Delta$. To be proper, it must also be associative, and be compatible with the point of \mathcal{X} when it has one. If we package this structure on a family \mathcal{X} into a record – a point and substitution operation $\mathcal{I} \xrightarrow{\eta} \mathcal{X} \xleftarrow{\mu} \mathcal{X} \odot \mathcal{X}$ satisfying unit and associativity laws – we end up with precisely a *monoid* in $(\mathbf{Fam}_s, \mathcal{I}, \odot)$. Monoids can be equivalently expressed using the internal hom as $\mathcal{I} \xrightarrow{\eta} \mathcal{X} \xrightarrow{\mu} \llbracket \mathcal{X}, \mathcal{X} \rrbracket$, which is the presentation we prefer for technical reasons: the metatheory proofs would not go through if we had used \odot .

```
record Mon (M : Familys) : Set where
  field η : I → M
        μ : M → ⟨ M , M ⟩

  lunit : {σ : Γ -[ M ] → Δ} {v : I α Γ} → μ (η v) σ ≡ σ v
  runit  : {t : M α Γ} → μ t η ≡ t
  assoc  : {σ : Γ -[ M ] → Δ} {ς : Δ -[ M ] → Θ} {t : M α Γ} →
           μ (μ t σ) ς ≡ μ t (λ v → μ (σ v) ς)
```

The multiplication μ represents simultaneous substitution, replacing every variable in context Γ with an M -term in Δ . In practice (e.g. in β -reduction) one often uses one- or two-variable substitution for the last variable or variables in the context, which is derived using `add`:

$$\begin{aligned} \llbracket _ / _ \rrbracket &: M \alpha \Gamma \rightarrow M \beta (\alpha \cdot \Gamma) \rightarrow M \beta \Gamma & \llbracket _ , _ / _ \rrbracket_2 &: M \alpha \Gamma \rightarrow M \beta \Gamma \rightarrow M \tau (\alpha \cdot \beta \cdot \Gamma) \rightarrow M \tau \Gamma \\ \llbracket s / _ \rrbracket t &= \mu t (\text{add } M s \eta) & \llbracket s_1 , s_2 / _ \rrbracket_2 t &= \mu t (\text{add } M s_1 (\text{add } M s_2 \eta)) \end{aligned}$$

Monoid homomorphisms preserve the unit and multiplication. When M is a family associated with an inductively defined syntax, N is a model of the syntax, and f is a map $M \rightarrow N$, the preservation of multiplication $\langle \mu \rangle : \{\sigma : \Gamma -[\mathcal{M}] \rightarrow \Delta\} \{t : M \alpha \Gamma\} \rightarrow f(M.\mu t \sigma) \equiv N.\mu (f t) (f \circ \sigma)$ expresses the *semantic substitution lemma*: the interpretation of substitution in the syntax is the substitution of interpretations in the model. We give its familiar form for single-variable substitution as an example below. The fact that f commutes with `add` is established by function extensionality, case analysis on the variable, and preservation of the unit $\langle \eta \rangle : \{v : \mathcal{I} \alpha \Gamma\} \rightarrow f(M.\eta v) \equiv N.\eta v$.

$$\begin{aligned} \text{sub-lemma} &: (s : M \alpha \Gamma) (t : M \beta (\alpha \cdot \Gamma)) \rightarrow f(M.\llbracket s / _ \rrbracket t) \equiv N.\llbracket f s / _ \rrbracket (f t) \\ \text{sub-lemma } s t &= \text{trans } \langle \mu \rangle (\text{cong } (N.\mu (f t)) (\text{ext } \lambda \{ \text{new} \rightarrow \text{refl} ; (\text{old } y) \rightarrow \langle \eta \rangle \})) \end{aligned}$$

Every monoid has an induced pointed \square -coalgebra instance $M_*^\square : \mathbf{Coalg}_*$, where renaming is implemented by substituting variables for variables, and pointed \square -coalgebra laws follow from monoid axioms. But, what if M already comes with a coalgebra structure? We will revisit this question after examining some general properties of maps into internal homs.

2.3 Parametrised maps

Working with families with coalgebra structure lets us be precise about when renaming is required in constructions: for example, we need the coalgebra structure for weakening on \mathcal{X} , but not for copairing of \mathcal{X} -valued context maps. However, as Fiore et al. [1999] show, the interaction between the substitution tensor and the presheaf structure forms a core part of the model theory that will

need to be recast in our coalgebraic setting. To achieve this, we introduce an important collection of properties for maps into internal homs (equivalently, out of substitution tensors).

2.3.1 Coalgebraic maps. Maps of the form $f : \mathcal{X} \rightarrow \llbracket \mathcal{P}, \mathcal{Y} \rrbracket$ play a vital role in the mathematical theory, since they transform \mathcal{X} to \mathcal{Y} while altering the variable context according to a parameter family \mathcal{P} . Many syntactic operations are of this form: for example, as we have already mentioned, renaming $\mathcal{X} \rightarrow \llbracket \mathcal{I}, \mathcal{X} \rrbracket$ and substitution $\mathcal{X} \rightarrow \llbracket \mathcal{X}, \mathcal{X} \rrbracket$ are maps parametrised by \mathcal{I} and \mathcal{X} , respectively. If all three families $\mathcal{X}, \mathcal{P}, \mathcal{Y}$ have pointed coalgebra structure, there are two ways in which a map f can be compatible with the various renaming operations: it can preserve renaming, or identify terms that get renamed via the coalgebra structure and via the context map. Furthermore, one can state a compatibility law between the points of all three families. Maps which satisfy these three properties are called *coalgebraic*.

record Coalgebraic ($f : \mathcal{X} \rightarrow \llbracket \mathcal{P}, \mathcal{Y} \rrbracket$): **Set** where

$$\begin{aligned} \text{field } \text{rof} & : \{\sigma : \Gamma \dashv \llbracket \mathcal{P} \rrbracket \rightarrow \Delta\} \{\varrho : \Delta \rightsquigarrow \Theta\} \{t : \mathcal{X} \alpha \Gamma\} \rightarrow \mathcal{Y}.r (f t \sigma) \varrho \equiv f t (\lambda v \rightarrow \mathcal{P}.r (\sigma v) \varrho) \\ \text{for} & : \{\rho : \Gamma \rightsquigarrow \Delta\} \{\varsigma : \Delta \dashv \llbracket \mathcal{P} \rrbracket \rightarrow \Theta\} \{t : \mathcal{X} \alpha \Gamma\} \rightarrow f (\mathcal{X}.r t \rho) \varsigma \equiv f t (\varsigma \circ \rho) \\ \text{fo}\eta & : \{v : \mathcal{I} \alpha \Gamma\} \rightarrow f (\mathcal{X}.\eta v) \mathcal{P}.\eta \equiv \mathcal{Y}.\eta v \end{aligned}$$

Though the hom of pointed families is not in general pointed, the codomain of coalgebraic maps has a pointed \square -coalgebra structure, and the map itself is a pointed \square -coalgebra homomorphism.

$$\begin{aligned} \text{Cod}_*^\square & : \text{Coalg}_* \llbracket \mathcal{P}, \mathcal{Y} \rrbracket \\ \text{Cod}_*^\square & = \text{record} \{ \eta = \lambda v \sigma \rightarrow f (\mathcal{X}.\eta v) \sigma \\ & \quad ; \square = \text{record} \{ r = \lambda h \rho \sigma \rightarrow h (\sigma \circ \rho) ; \text{counit} = \text{refl} ; \text{comult} = \text{refl} \} \\ & \quad ; \text{ro}\eta = \lambda \{v\}\{\rho\} \rightarrow \text{ext } \lambda \sigma \rightarrow \text{begin } f (\mathcal{X}.\eta v) (\sigma \circ \rho) \equiv \langle \text{for} \rangle \\ & \quad \quad \quad f (\mathcal{X}.r (\mathcal{X}.\eta v) \rho) \sigma \equiv \langle \mathcal{X}.\text{ro}\eta \rangle \\ \text{f}_*^\square \Rightarrow & : \text{Coalg}_* \Rightarrow \mathcal{X}_*^\square \text{Cod}_*^\square f \\ \text{f}_*^\square \Rightarrow & = \text{record} \{ \square \Rightarrow = \text{record} \{ \langle r \rangle = \text{ext } (\lambda \sigma \rightarrow \text{for}) ; \langle \eta \rangle = \text{refl} \} \end{aligned}$$

Examples of coalgebraic maps are the renaming map for a pointed \square -coalgebra:

$$\begin{aligned} r^c & : (\mathcal{X}_*^\square : \text{Coalg}_* \mathcal{X}) \rightarrow \text{Coalgebraic } \mathcal{X}_*^\square \mathcal{I}_*^\square \mathcal{X}_*^\square \mathcal{X}.r \\ r^c \mathcal{X}_*^\square & = \text{record} \{ \text{rof} = \text{sym comult} ; \text{for} = \text{sym comult} ; \text{fo}\eta = \text{ro}\eta \} \end{aligned}$$

as well as the multiplication μ for a substitution monoid and the structural map $j : \mathcal{I} \rightarrow \llbracket \mathcal{X}, \mathcal{X} \rrbracket$ of skew-closed categories that corresponds to application (mapping v to $\sigma \mapsto \sigma(v)$). Later on, we will be able to show that universal parametrised maps from an initial syntactic algebra are coalgebraic, which will be a prerequisite for proving the substitution axioms.

2.3.2 Lifting and strength. The definition of simultaneous substitution by hand involves traversing the host term, applying the substitution map to variables, and recursing into subterms. The difficulties arise when the subterm has a newly bound variable, such as in the body of a λ -abstraction. To avoid variable shadowing, substitution must map the newly bound variable to itself; to avoid variable capture, the terms to be substituted must not contain free occurrences of the new variable. Fortunately, intrinsically-typed encoding guides the complex de Bruijn ‘‘arithmetic’’ required and guarantees to maintain type- and scope-safety.

The critical step is applying a substitution $\sigma : \Gamma \dashv \llbracket \mathcal{P} \rrbracket \rightarrow \Delta$ to a term in an extended context $\mathcal{P} \beta (\alpha \cdot \Gamma)$, without disturbing the newly bound variable α . The usual name for the transformation needed on σ is *lifting* [Altenkirch and Reus 1999; Benton et al. 2012], which can be generalised to arbitrary extensions as $\text{lift } \Theta \sigma : (\Theta \dagger \Gamma) \dashv \llbracket \mathcal{P} \rrbracket \rightarrow (\Theta \dagger \Delta)$. The inductive definition of lifting

requires both a point (to map newly bound variables to themselves) and renaming (to weaken the context of the recursive call), so we demand the structure of a pointed \square -coalgebra on \mathcal{P} :

$$\begin{aligned} \text{lift} &: (\Theta : \text{Ctx}) \rightarrow (\Gamma \text{ -}[\mathcal{P}] \rightarrow \Delta) \rightarrow (\Theta \dot{+} \Gamma) \text{ -}[\mathcal{P}] \rightarrow (\Theta \dot{+} \Delta) \\ \text{lift } \emptyset & \quad \sigma v \quad = \sigma v \\ \text{lift } (\tau \cdot \Theta) \sigma \text{ new} & = \mathcal{P}.\eta \text{ new} \\ \text{lift } (\tau \cdot \Theta) \sigma (\text{old } v) & = \mathcal{P}.r (\text{lift } \Theta \sigma v) \text{ old} \end{aligned}$$

While this definition works, it expresses `lift` as a `Set`-level transformation of context maps, rather than as a morphism of sorted families – as such, it counterposes our goal of representing metasynthetic operations purely algebraically. Consequently, it is not a priori obvious what laws `lift` should satisfy. Benton et al. [2012, Section 4] demonstrate the intricate inter-dependence of `lift` with the other structures by listing eight laws that concern the interaction amongst lifting, renaming, and substitution, all of which must be proved in order, with each property building on top of the previous ones. It is these sorts of auxiliary definitions and ad-hoc lemmas that we wish to avoid in our systematic, categorical development of abstract syntax.

The key to demystifying `lift` lies in the categorical concept of *cotensorial strength* [Kock 1971] for a sorted-family endofunctor F : namely, an operation of type $F(\llbracket \mathcal{X}, \mathcal{Y} \rrbracket) \rightarrow \llbracket \mathcal{X}, F\mathcal{Y} \rrbracket$ satisfying certain coherence laws. The analogous tensorial strength $F(\mathcal{X}) \odot \mathcal{Y} \rightarrow F(\mathcal{X} \odot \mathcal{Y})$ plays a central role in the presheaf model of abstract syntax [Fiore 2008], since it captures the intuition of pushing a substitution into a syntactic structure represented by the endofunctor F . Most notably, `lift` attains a natural place within the categorical model as it can be used to implement strength for the instance $F = \delta \Theta$: the resulting operation $\delta \Theta \llbracket \mathcal{P}, \mathcal{X} \rrbracket \rightarrow \llbracket \mathcal{P}, \delta \Theta \mathcal{X} \rrbracket$ is then responsible for pushing a substitution $\sigma: \Gamma \text{ -}[\mathcal{P}] \rightarrow \Delta$ under a binder of variables in context Θ .

In our setting of families and coalgebras, a more refined notion of strength is required (similar to the structural strength of Borthelle et al. [2020]). A *coalgebraic strength* for F is a transformation

$$\text{str} : F(\llbracket \mathcal{P}, \mathcal{X} \rrbracket) \rightarrow \llbracket \mathcal{P}, F\mathcal{X} \rrbracket$$

where \mathcal{P} is a pointed \square -coalgebra and \mathcal{X} is a sorted family. The operation must be *natural* in both components: that is, it must suitably commute with the functorial mapping of any pointed \square -coalgebra homomorphism $f: \mathcal{Q} \rightarrow \mathcal{P}$ and family of maps $g: \mathcal{X} \rightarrow \mathcal{Y}$.

$$\begin{aligned} \text{str-nat}_1 &: (f_*^\square \Rightarrow \text{Coalg}_*^\square \mathcal{Q}_*^\square \mathcal{P}_*^\square f) (h : F(\llbracket \mathcal{P}, \mathcal{X} \rrbracket) \alpha \Gamma) (\sigma : \Gamma \text{ -}[\mathcal{Q}] \rightarrow \Delta) \rightarrow \\ & \quad \text{str } \mathcal{P}_*^\square \mathcal{X} h (f \circ \sigma) \equiv \text{str } \mathcal{Q}_*^\square \mathcal{X} (F_1 (\lambda h' \sigma' \rightarrow h' (f \circ \sigma')) h) \sigma \end{aligned}$$

$$\begin{aligned} \text{str-nat}_2 &: (g : \mathcal{X} \rightarrow \mathcal{Y}) (h : F(\llbracket \mathcal{P}, \mathcal{X} \rrbracket) \alpha \Gamma) (\sigma : \Gamma \text{ -}[\mathcal{P}] \rightarrow \Delta) \rightarrow \\ & \quad \text{str } \mathcal{P}_*^\square \mathcal{Y} (F_1 (\lambda h' \sigma' \rightarrow g (h' \sigma')) h) \sigma \equiv F_1 g (\text{str } \mathcal{P}_*^\square \mathcal{X} h \sigma) \end{aligned}$$

The strength also satisfies a unit law `str-unit`: $(h : F(\llbracket \mathcal{I}, \mathcal{X} \rrbracket) \alpha \Gamma) \rightarrow \text{str } \mathcal{I}_*^\square \mathcal{X} h \text{ id} \equiv F_1 (i \mathcal{X}) h$ and an associativity law with respect to the two other structural transformations of skew-closed categories: the unit $i = \lambda h \rightarrow h \text{ id} : \llbracket \mathcal{I}, \mathcal{X} \rrbracket \rightarrow \mathcal{X}$ and the (curried) composition of internal homs $\mathbb{L} = \lambda h \sigma \varsigma \rightarrow h (\lambda v \rightarrow \sigma v \varsigma) : \llbracket \mathcal{Y}, \mathcal{Z} \rrbracket \rightarrow \llbracket \llbracket \mathcal{X}, \mathcal{Y} \rrbracket, \llbracket \mathcal{X}, \mathcal{Z} \rrbracket \rrbracket$. The associativity law has to be stated in terms of coalgebraic maps $f: \mathcal{P} \rightarrow \llbracket \mathcal{Q}, \mathcal{R} \rrbracket$, since the usual pentagon identity [Kock 1971, Lemma 1.3] is too strict in the skew-closed setting (see `str-assoc` in the diagram below). This generalised associativity law neatly combines with naturality to give the following powerful corollary, given both as Agda code and – for a clearer presentation – as a commutative diagram:

$$\begin{aligned}
\text{str-dist} : \{f : \mathcal{P} \rightarrow \llbracket Q, \mathcal{R} \rrbracket\} (f^c : \text{Coalgebraic } \mathcal{P}_*^\square \mathcal{Q}_*^\square \mathcal{R}_*^\square f) \\
(h : F \llbracket \mathcal{R}, \mathcal{X} \rrbracket \alpha \Gamma) (\sigma : \Gamma \dashv \llbracket \mathcal{P} \rrbracket \rightarrow \Delta) (\zeta : \Delta \dashv \llbracket Q \rrbracket \rightarrow \Theta) \rightarrow \\
\text{str } \mathcal{R}_*^\square \mathcal{X} h (\lambda v \rightarrow f (\sigma v) \zeta) \\
\equiv \text{str } \mathcal{Q}_*^\square \mathcal{X} (\text{str } \mathcal{P}_*^\square \llbracket Q, \mathcal{X} \rrbracket (F_1 (\lambda h \sigma \zeta \rightarrow h (\lambda v \rightarrow f (\sigma v) \zeta)) h) \sigma) \zeta
\end{aligned}$$

$$\begin{array}{ccccc}
F \llbracket \mathcal{R}, \mathcal{X} \rrbracket & \xrightarrow{FL} & F \llbracket \llbracket Q, \mathcal{R} \rrbracket, \llbracket Q, \mathcal{X} \rrbracket \rrbracket & \xrightarrow{F \llbracket f, \text{id} \rrbracket} & F \llbracket \mathcal{P}, \llbracket Q, \mathcal{X} \rrbracket \rrbracket \\
\downarrow \text{str}_{\mathcal{R}, \mathcal{X}} & & \text{str}_{\llbracket Q, \mathcal{R} \rrbracket, \llbracket Q, \mathcal{X} \rrbracket} \downarrow & \text{str-nat}_1 & \downarrow \text{str}_{\mathcal{P}, \llbracket Q, \mathcal{X} \rrbracket} \\
& & \llbracket \llbracket Q, \mathcal{R} \rrbracket, F \llbracket Q, \mathcal{X} \rrbracket \rrbracket & \xrightarrow{\llbracket f, \text{id} \rrbracket} & \llbracket \mathcal{P}, F \llbracket Q, \mathcal{X} \rrbracket \rrbracket \\
& \text{str-assoc} & & & \downarrow \llbracket \text{id}, \text{str}_{Q, \mathcal{X}} \rrbracket \\
\llbracket \mathcal{R}, F \mathcal{X} \rrbracket & \xrightarrow{\llbracket \cdot \rrbracket} & \llbracket \llbracket Q, \mathcal{R} \rrbracket, \llbracket Q, F \mathcal{X} \rrbracket \rrbracket & \xrightarrow{\llbracket f, \text{id} \rrbracket} & \llbracket \mathcal{P}, \llbracket Q, F \mathcal{X} \rrbracket \rrbracket
\end{array}$$

As f is coalgebraic, its codomain $\llbracket Q, \mathcal{R} \rrbracket$ is a pointed \square -coalgebra (Section 2.3.1) and thus is a valid first component to $\text{str}_{\llbracket Q, \mathcal{R} \rrbracket, \llbracket Q, \mathcal{X} \rrbracket}$. Remarkably, for different choices of f , the str-dist corollary above generalises all four lifting lemmas given by Benton et al., fulfilling our goal of placing lift and its laws on a formal foundation.

The axiomatisation of strength as a categorical concept is required for the initiality and freeness theorems of Section 3. The strength for a signature endofunctor (Section 4.2) will invariably derive from the **Strength** instance of the context extension endofunctor δ , whose implementation makes use of lift : to feed a context map $\Gamma \dashv \llbracket \mathcal{P} \rrbracket \rightarrow \Delta$ into a hom $\llbracket \mathcal{P}, \mathcal{X} \rrbracket \alpha (\Theta \dot{+} \Gamma)$, one has to extend both its domain and codomain with Θ . The strength proofs feature many of the intricate properties we axiomatised earlier, such as homomorphism and coalgebraic laws. We refer the reader to the [formalisation](#) for the details.

2.3.3 Coalgebraic monoids. We revisit the question posed at the end of Section 2.2.3. As shown there, every substitution monoid is a \square -coalgebra, since renaming with $\rho : \Gamma \rightsquigarrow \Delta$ can be implemented as substitution with $\eta \circ \rho : \Gamma \dashv \llbracket \mathcal{X} \rrbracket \rightarrow \Delta$. This may suggest that renaming for a syntax can be derived from substitution, and so that one only needs to define the latter operation. Attempting this will be futile, however: substitution into terms with variable binding will require weakening (as part of lift), a special case of renaming. Consequently, to equip a family of terms with substitution structure, one needs to have already shown that it is a pointed \square -coalgebra.

This a priori renaming structure will not necessarily be equivalent to the one induced by substitution (unlike in the presheaf model, where the two are identified by the quotiented dependent sum used in the definition of the substitution tensor product), but their equivalence is a prerequisite of the free monoid proof in Theorem 3.1 below. We overcome this conflict by axiomatising families with compatible pointed coalgebra and monoid structures, called *coalgebraic monoids*:

$$\begin{aligned}
&\text{record } \text{CoalgMon} (\mathcal{X} : \text{Family}_s) : \text{Set } \text{where} \\
&\quad \text{field } \mathcal{X}_*^\square : \text{Coalg}_* \mathcal{X} \\
&\quad \mathcal{X}^M : \text{Mon } \mathcal{X} \\
&\quad \eta\text{-compat} : \{v : \mathcal{I} \alpha \Gamma\} \rightarrow \eta_*^\square v \equiv \eta^M v \\
&\quad \mu\text{-compat} : \{\rho : \Gamma \rightsquigarrow \Delta\} \{t : \mathcal{X} \alpha \Gamma\} \rightarrow \mathbf{r} t \rho \equiv \mu t (\eta^M \circ \rho)
\end{aligned}$$

The compatibility laws ensure that the existing \mathcal{X}_*^\square and derived $\mathcal{X}^M \cdot \square$ pointed \square -coalgebra structures on \mathcal{X} are equivalent, and, in particular, can be exchanged in the first component of str .

We have now set up all the mathematical foundations needed for the development. Next, we move on to the central conceptual tool of our framework: initial algebra semantics.

3 INITIAL ALGEBRA SEMANTICS

Initial algebra semantics [Goguen et al. 1976] is one of the most useful concepts in the study of data types and functional programming. It stems from the observation that every inductive data type \mathbb{T} corresponds to an endofunctor $F_{\mathbb{T}}$ of which the data type is an initial algebra. For example, the data type \mathbb{N} of natural numbers has associated endofunctor $F_{\mathbb{N}}(A) = 1 + A$ and comes equipped with an isomorphism $1 + \mathbb{N} \xrightarrow{\cong} \mathbb{N}$. Furthermore, for any $F_{\mathbb{N}}$ -algebra A – equivalently, a type A with an element $z : A$ and an operation $s : A \rightarrow A$ – one has a unique function $\text{rec}_{\mathbb{N}}(z, s) : \mathbb{N} \rightarrow A$ that satisfies $\text{rec}_{\mathbb{N}}(z, s)(0) = z$ and $\text{rec}_{\mathbb{N}}(z, s)(n+1) = s(\text{rec}_{\mathbb{N}}(z, s)(n))$. This is usually called the *recursor*, *fold*, or *catamorphism* for \mathbb{N} and captures the process of consuming an inhabitant of the inductive type by recursion on its structure.

This idea extends to inductively defined families of types and endofunctors thereon [Altenkirch et al. 2015], such as a data type of intrinsically-typed terms. In particular, the recursor out of an initial family of terms can then be seen as a compositional semantic interpretation map: for example, an interpretation of the STLC in any of its models. As we explore next, initial algebra semantics is not only useful for semantic purposes: it can also be used to implement syntactic operations and prove their laws for free.

3.1 Signature algebras and monoids

Just as the categorical theory, our framework is entirely signature-generic. The metatheory developed here can be freely instantiated for any second-order signature, encompassing algebraic theories, computational calculi, logics, etc. Thus, for the remainder of the section, we fix a sorted-family endofunctor $\Sigma : \mathbf{Fam}_s \rightarrow \mathbf{Fam}_s$ with an instance of coalgebraic strength $\Sigma : \mathbf{Str}$. Algebras $\Sigma \mathcal{A} \rightarrow \mathcal{A}$ for this endofunctor represent families that support the operations of the signature. For example, for the signature Σ_{Δ} of the STLC, a Σ_{Δ} -algebra is a family \mathcal{A} equipped with operations $\text{app} : \mathcal{A} (\alpha \succ \beta) \Gamma \times \mathcal{A} \alpha \Gamma \rightarrow \mathcal{A} \beta \Gamma$ and $\text{abs} : \mathcal{A} \beta (\tau \cdot \Gamma) \rightarrow \mathcal{A} (\alpha \succ \beta) \Gamma$.

We also consider Σ -algebras with substitution structure, known as Σ -monoids [Fiore et al. 1999; Fiore and Saville 2017]. These are families with algebra structure $\Sigma \mathcal{X} \rightarrow \mathcal{X}$ and monoid structure $\mathcal{I} \rightarrow \mathcal{X} \rightarrow \langle \mathcal{X}, \mathcal{X} \rangle$ which are compatible with each other; the compatibility is expressed using the coalgebraic strength str on Σ , used to swap the constructor application with the parametrisation by a context map. Note that str uses the pointed \square -coalgebra structure derived from substitution.

record $\Sigma \text{Mon} (\mathcal{X} : \mathbf{Family}_s) : \mathbf{Set}$ **where**

field $\mathcal{X}^{\text{M}} : \mathbf{Mon} \mathcal{X}$

alg : $\Sigma \mathcal{X} \rightarrow \mathcal{X}$

$\mu \langle \text{alg} \rangle : \{ \sigma : \Gamma \rightarrow \Delta \} \{ t : \Sigma \mathcal{X} \alpha \Gamma \} \rightarrow$

$\mu (\text{alg } t) \sigma \equiv \text{alg} (\text{str } \mathcal{X}_*^{\square} \mathcal{X} (\Sigma_1 \mu t) \sigma)$

$$\begin{array}{ccc}
 \Sigma \mathcal{X} & \xrightarrow{\Sigma \mu} & \Sigma \langle \mathcal{X}, \mathcal{X} \rangle \xrightarrow{\text{str}_{\mathcal{X}, \mathcal{X}}} \langle \mathcal{X}, \Sigma \mathcal{X} \rangle \\
 \text{alg} \downarrow & & \downarrow \langle \text{id}, \text{alg} \rangle \\
 \mathcal{X} & \xrightarrow{\mu} & \langle \mathcal{X}, \mathcal{X} \rangle
 \end{array}$$

3.2 Algebras with metavariables

Our formalisation is novel in that it also incorporates *parametrised metavariables* [Aczel 1978; Hamana 2004; Fiore 2008]: terms that take the syntactic form $\mathfrak{m} \langle -, \dots, - \rangle$ and are to be imagined as “named holes with slots”, with the “hole” named \mathfrak{m} standing for an unspecified term with ℓ open variables and the “slots” containing terms that occupy the open variables. They too come from a sorted family: for instance, the family \mathfrak{B} containing the metavariables $\mathfrak{m} : \mathfrak{B} \beta [\alpha]$ and $\mathfrak{n} : \mathfrak{B} \alpha []$ can be used to construct terms such as $r = (\lambda x : \alpha. \mathfrak{m} \langle x \rangle) (\mathfrak{n} \langle \rangle)$ and $t = \mathfrak{m} \langle \mathfrak{n} \langle \rangle \rangle$. Metavariables let us abstractly reason about generic terms of some specific syntactic structure: for example, the notion of β -equivalence can be axiomatised as the identification of all common instances of r and t above. For the rest of this section until Section 3.5, we also fix a sorted family \mathfrak{X} of metavariables.

A Σ -algebra structure map on \mathcal{A} captures all the constructors of a second-order syntax, but for full generality, one also needs to account for both variables and metavariables. These are respectively represented by a point $\mathbf{var} : \mathcal{I} \rightarrow \mathcal{A}$ and a metavariable operator $\mathbf{mvar} : \mathfrak{X} \rightarrow \langle \mathcal{A}, \mathcal{A} \rangle$. In elementary terms, the latter is a function that associates an \mathcal{A} -term $\mathbf{mvar} \ m \ \varepsilon : \mathcal{A} \ \tau \ \Gamma$ to every parametrised metavariable $m : \mathfrak{X} \ \tau \ \Pi$ and metavariable environment $\varepsilon : \Pi \ -[\mathcal{A}] \rightarrow \Gamma$. For example, the term $m \langle t, s \rangle : \mathcal{A} \ \tau \ \Gamma$ for a metavariable $m : \mathfrak{X} \ \tau \ [\alpha, \beta]$, and terms $t : \mathcal{A} \ \alpha \ \Gamma$ and $s : \mathcal{A} \ \beta \ \Gamma$ is represented by $\mathbf{mvar} \ m \ \lambda \{ \mathbf{new} \rightarrow t ; \mathbf{old} \ \mathbf{new} \rightarrow s \}$. Families that support this structure will be called (Σ, \mathfrak{X}) -meta-algebras, with the expected notion of homomorphism between them.

```
record MetaAlg (A : Familys) : Set where
  field alg :  $\Sigma \mathcal{A} \rightarrow \mathcal{A}$ 
        var :  $\mathcal{I} \rightarrow \mathcal{A}$ 
        mvar :  $\mathfrak{X} \rightarrow \langle \mathcal{A}, \mathcal{A} \rangle$ 
```

```
record MetaAlg $\Rightarrow$  (Az : MetaAlg A) (Bz : MetaAlg B) (f : A → B) : Set where
  field <alg> : {t :  $\Sigma \mathcal{A} \ \alpha \ \Gamma$ } → f (A.alg t) ≡ B.alg ( $\Sigma f$  t)
        <var> : {v :  $\mathcal{I} \ \alpha \ \Gamma$ } → f (A.var v) ≡ B.var v
        <mvar> : {m :  $\mathfrak{X} \ \alpha \ \Pi$ }{ $\varepsilon : \Pi \ -[ \mathcal{A} ] \rightarrow \Gamma$ } → f (A.mvar m  $\varepsilon$ ) ≡ B.mvar m (f ∘  $\varepsilon$ )
```

Such meta-algebras and their homomorphisms form a category **MetaAlg**, whose initial object – whenever it exists – will be denoted $\mathbb{T} \ \mathfrak{X}$ (or just \mathbb{T} , if the metavariable family is clear from the context) with structural maps **alg**, **var** and **mvar**. The universal property of initial objects states that there is a unique meta-algebra homomorphism **sem** : $\mathbb{T} \rightarrow \mathcal{A}$ into any meta-algebra \mathcal{A} . Note that, varying \mathfrak{X} , \mathbb{T} acts as a mapping from families \mathfrak{X} to (Σ, \mathfrak{X}) -meta-algebras $\mathbb{T} \ \mathfrak{X}$. Adapted from [Fiore \[2008, Theorem 2\]](#), we have the following main result in the study of abstract syntax:

THEOREM 3.1. *The initial meta-algebra $\mathbb{T} \ \mathfrak{X}$ is the free Σ -monoid on \mathfrak{X} .*

The most interesting implication of this concise statement is that the substitution structure of the free Σ -monoid is induced purely by initiality. Since the initial meta-algebra \mathbb{T} corresponds to the inductively defined grammar of terms (see [Section 4.3](#)), it follows that the syntactic structure fully determines the action of substitution. The theorem formally captures the observation that much of syntactic metatheory is uninteresting boilerplate, and the methodology of initial algebra semantics will allow us to extract substitution, compositional interpretations, and related correctness laws from the syntax for free. We outline the proof of the theorem in the rest of the section.

3.3 Parametrised interpretations

The initial algebra approach explains and validates our adherence to formality and efforts to internalise syntactic operations as categorical constructions at the level of sorted families. Since they are maps out of an initial meta-algebra, the renaming $\mathbb{T} \rightarrow \square \mathbb{T}$ and the substitution $\mathbb{T} \rightarrow \langle \mathbb{T}, \mathbb{T} \rangle$ operations may be induced by initiality as soon as we show that their codomains $\square \mathbb{T}$ and $\langle \mathbb{T}, \mathbb{T} \rangle$ are meta-algebras. This will be derived from the following lemma.

LEMMA 3.2. *Given a pointed \square -coalgebra \mathcal{P} , a Σ -algebra \mathcal{A} , and family maps $\varphi : \mathcal{P} \rightarrow \mathcal{A}$ and $\chi : \mathfrak{X} \rightarrow \langle \mathcal{A}, \mathcal{A} \rangle$, the internal hom $\langle \mathcal{P}, \mathcal{A} \rangle$ acquires a (Σ, \mathfrak{X}) -meta-algebra structure.*

The proof is encapsulated in the **Traversal** module, instantiations of which will give rise to homomorphic initial algebra interpretations $\mathbb{T} \rightarrow \langle \mathcal{P}, \mathcal{A} \rangle$. It crucially relies on the coalgebraic strength when defining the structure map $\Sigma \langle \mathcal{P}, \mathcal{A} \rangle \rightarrow \langle \mathcal{P}, \Sigma \mathcal{A} \rangle \rightarrow \langle \mathcal{P}, \mathcal{A} \rangle$. A simple corollary (derived by instantiating \mathcal{P} with \mathcal{I}) is that if \mathcal{A} is a meta-algebra, then so is $\square \mathcal{A}$.

```

module Traversal ( $\mathcal{P}_*^\square : \mathbf{Coalg}_* \mathcal{P}$ ) ( $\mathit{alg}_{\mathcal{A}} : \Sigma \mathcal{A} \rightarrow \mathcal{A}$ ) ( $\varphi : \mathcal{P} \rightarrow \mathcal{A}$ ) ( $\chi : \mathfrak{X} \rightarrow \llbracket \mathcal{A}, \mathcal{A} \rrbracket$ ) where
   $\mathit{Trav}^\Sigma : \mathbf{MetaAlg} \llbracket \mathcal{P}, \mathcal{A} \rrbracket$ 
   $\mathit{Trav}^\Sigma = \mathbf{record} \{$ 
     $\mathit{alg} = \lambda h \ \sigma \rightarrow \mathit{alg}_{\mathcal{A}} (\mathbf{str} \mathcal{P}_*^\square \mathcal{A} h \sigma)$ 
     $;$   $\mathit{var} = \lambda v \ \sigma \rightarrow \varphi (\sigma v)$ 
     $;$   $\mathit{mvar} = \lambda m \ \varepsilon \ \sigma \rightarrow \chi m (\lambda v \rightarrow \varepsilon v \ \sigma) \}$ 

```

One may be tempted to immediately induce the substitution map $\mathbb{T} \rightarrow \llbracket \mathbb{T}, \mathbb{T} \rrbracket$ as a \mathbb{T} -parametrised traversal into \mathbb{T} . However, that would require a pointed \square -coalgebra instance on \mathbb{T} , which we do not yet have. Our formalism very concretely exhibits the dependence of substitution on renaming.

3.4 Σ -monoid structure by initiality

The construction of the renaming and substitution maps on the initial meta-algebra \mathbb{T} make extensive use of *initiality*: maps of the form $\mathbb{T} \rightarrow \mathcal{A}$ are uniquely induced by equipping \mathcal{A} with a meta-algebra structure. The proofs of the renaming and substitution laws are then established by proving that the (composite) maps that correspond to the two sides of an equation are meta-algebra homomorphisms and must therefore be equal.

As an example, the renaming map $\mathit{ren} : \mathbb{T} \rightarrow \square \mathbb{T}$ is induced as the unique homomorphism from $\mathbb{T} \in \mathbf{MetaAlg}$ to the sorted family $\square \mathbb{T}$ regarded as a meta-algebra by means of [Lemma 3.2](#). The counit law $\mathit{ren} t \ \mathit{id} = t$ amounts to showing that the mapping $t \mapsto \mathit{ren} t \ \mathit{id} : \mathbb{T} \rightarrow \mathbb{T}$ is a meta-algebra homomorphism so that – by initiality of \mathbb{T} – must be equal to the identity $t \mapsto t$. Similarly, proving that the \mathbf{Fam}_s -morphisms defined as $t, \rho, \varrho \mapsto \mathit{r} t (\varrho \circ \rho)$ and $t, \rho, \varrho \mapsto \mathit{r} (\mathit{r} t \ \rho) \ \varrho$ of type $\mathbb{T} \rightarrow \llbracket \mathcal{I}, \llbracket \mathcal{I}, \mathbb{T} \rrbracket \rrbracket$ are meta-algebra homomorphisms implies that they must be equal, giving us the comultiplication law. The proofs (available in the [formalisation](#)) depend on the structure-preservation properties of ren , the strength laws $\mathbf{str}\text{-unit}$ and $\mathbf{str}\text{-nat}_2$, and the corollary $\mathbf{str}\text{-dist}$ applied to the coalgebraic map $\mathit{j} : \mathcal{I} \rightarrow \llbracket \mathcal{I}, \mathcal{I} \rrbracket$.

We therefore have a pointed \square -coalgebra instance $\mathbb{T}_*^\square : \mathbf{Coalg}_* \mathbb{T}$ that is then used as a traversal parameter to induce the substitution map $\mathit{sub} : \mathbb{T} \rightarrow \llbracket \mathbb{T}, \mathbb{T} \rrbracket$. Initiality once again helps us prove the substitution monoid laws abstractly; in fact, the reasoning steps very closely resemble those needed in the proof of the renaming structure, and they can all be presented in the form of clear, categorical proofs. Once the instances $\mathbb{T}^M : \mathbf{Mon} \mathbb{T}$ and $\mathbb{T}^{\text{OM}} : \mathbf{CoalgMon} \mathbb{T}$ are derived, we further establish that \mathbb{T} is a Σ -monoid, the proof of which relies on the fact that coalgebraic monoids identify the given coalgebra instance \mathbb{T}_*^\square with the one induced from substitution.

The effort put into identifying the notions of coalgebraic map and strength ([Section 2.3.2](#)) pays off repeatedly: the comultiplication law, coalgebraic axioms, and substitution associativity (the *fusion lemmas* of [Allais et al. \[2017, Section 9.2\]](#)) are all established using $\mathbf{str}\text{-dist}$, instantiated at different coalgebraic maps (the application map j , renaming ren and substitution sub) to derive the generalised forms of the four lifting laws listed by [Benton et al. \[2012\]](#). The renaming and substitution operations and their correctness laws are derived in an elegant, mathematically-motivated manner, with no auxiliary definitions and ad-hoc lemmas, or additional reasoning machinery needed (such as the bisimulation/fusion framework developed by [AACMM \[2021\]](#)).

3.5 Free Σ -monoid structure

In the previous section we generically derived a Σ -monoid structure for $\mathbb{T} \mathfrak{X}$, for any family of metavariables \mathfrak{X} . This gives us a lawful substitution operation on $\mathbb{T} \mathfrak{X}$ that can be used in further developments, such as an operational semantics or equational theory. We can go further, however, and characterise $\mathbb{T} \mathfrak{X}$ as the *free Σ -monoid on \mathfrak{X}* : given any Σ -monoid \mathcal{M} and metavariable interpretation $\omega : \mathfrak{X} \rightarrow \mathcal{M}$, there is a unique Σ -monoid homomorphism $\omega^\# : \mathbb{T} \mathfrak{X} \rightarrow \mathcal{M}$ that satisfies $\omega^\#(\mathbf{mvar} \ m \ \mathbf{var}) = \omega(m)$ for every metavariable $m : \mathfrak{X} \tau \Pi$. As expected, $\omega^\#$ is constructed

by initiality, using the fact that the Σ -monoid \mathcal{M} is a meta-algebra with metavariable operator $\mathfrak{X} \xrightarrow{\omega} \mathcal{M} \xrightarrow{\mu} \llbracket \mathcal{M}, \mathcal{M} \rrbracket$. Initiality is also used to establish that the unique map $\omega^\#$ preserves substitution, i.e. that it is a Σ -monoid homomorphism.

Freeness induces a free-forgetful adjunction between the categories of sorted families and of Σ -monoids, and makes \mathbb{T} into the *free Σ -monoid monad* on sorted families, whose Kleisli extension $(\mathfrak{X} \rightarrow \mathbb{T} \mathfrak{Y}) \rightarrow (\mathbb{T} \mathfrak{X} \rightarrow \mathbb{T} \mathfrak{Y})$ acts as a form, albeit limited, of *metasubstitution* [Hamana 2004]: occurrences of metavariables from a family \mathfrak{X} in a term of $\mathbb{T} \mathfrak{X}$ get replaced with terms of $\mathbb{T} \mathfrak{Y}$ according to a mapping $\mathfrak{X} \rightarrow \mathbb{T} \mathfrak{Y}$. This important notion and its generalisation are discussed next.

3.6 Metasubstitution

A main difference between metasubstitution and object-level substitution is that the former is *capture-permitting*: a metavariable \mathfrak{m} in a term $\lambda x : \mathbb{N}. \mathfrak{m}\langle x \rangle$ can be replaced with terms that contain free occurrences of x ; so that valid instances could be $\lambda x. x$, $\lambda x. x \cdot (x + 1)$, or $\lambda x. \mathfrak{p}\langle x, x \rangle$ for a metavariable $\mathfrak{p} : \mathfrak{B} \mathbb{N} [\mathbb{N}, \mathbb{N}]$. The aforementioned metasubstitution map $(\mathfrak{X} \rightarrow \mathbb{T} \mathfrak{Y}) \rightarrow (\mathbb{T} \mathfrak{X} \rightarrow \mathbb{T} \mathfrak{Y})$ is limited in that the only free variables that the mapping $\zeta : \mathfrak{X} \rightarrow \mathbb{T} \mathfrak{Y}$ may refer to are the parameters of the metavariable: for example, in the open term $x : \mathbb{N}, y : \mathbb{N} \vdash \mathfrak{m}\langle y \rangle$ one cannot instantiate \mathfrak{m} with $x + y$, since x is not a parameter of \mathfrak{m} and is therefore not in scope in the output of ζ . As we wish metasubstitution to model “textual replacement”, this is an unnatural restriction.

Our goal then is to capture the following intuition [Fiore 2008]: the term that replaces a metavariable \mathfrak{m} can feature both the parameters of \mathfrak{m} , and all the object-level variables in scope at the occurrence of \mathfrak{m} in the term. In elementary terms, the type of the operation may be given as

$$\mathbf{msub} : \forall \{ \alpha \Gamma \} \rightarrow (t : \mathbb{T} \mathfrak{X} \alpha \Gamma) \rightarrow (\zeta : \forall \{ \tau \Pi \} \rightarrow \mathfrak{X} \tau \Pi \rightarrow \mathbb{T} \mathfrak{Y} \tau (\Pi \dot{+} \Gamma)) \rightarrow \mathbb{T} \mathfrak{Y} \alpha \Gamma$$

As usual, we aim to represent this operation as a morphism of families, and ideally derive it by initiality. Key to this is recognising the type $\forall \{ \Pi \} \rightarrow \mathfrak{X} \tau \Pi \rightarrow \mathbb{T} \mathfrak{Y} \tau (\Pi \dot{+} \Gamma)$ as the linear exponential of families $(\mathfrak{X} \multimap \mathbb{T} \mathfrak{Y}) \tau \Gamma$ (Section 2.1.2). The derived definition $[\mathfrak{X} \multimap \mathfrak{Y}] \Gamma \triangleq \forall \{ \tau \} \rightarrow (\mathfrak{X} \multimap \mathfrak{Y}) \tau \Gamma$ combines two sorted families into an unsorted one. We also modify the family exponential to take an unsorted family as first argument; that is, overloading notation: $(X \Rightarrow \mathfrak{Y}) \alpha \triangleq X \Rightarrow (\mathfrak{Y} \alpha)$. The type of metasubstitution may be now succinctly expressed as

$$\mathbf{msub} : \mathbb{T} \mathfrak{X} \rightarrow ([\mathfrak{X} \multimap \mathbb{T} \mathfrak{Y}] \Rightarrow \mathbb{T} \mathfrak{Y})$$

We can then use the following general result to induce \mathbf{msub} by initiality: given a Σ -monoid \mathcal{M} , the family $[\mathfrak{X} \multimap \mathcal{M}] \Rightarrow \mathcal{M}$ acquires a (Σ, \mathfrak{X}) -meta-algebra structure. However, this only applies if we assume an additional property of the signature endofunctor Σ : it comes with an *exponential strength* $\mathbf{estr} : \Sigma(X \Rightarrow \mathfrak{Y}) \rightarrow (X \Rightarrow \Sigma \mathfrak{Y})$ (equivalent to the usual Cartesian strength) for every unsorted \square -coalgebra X (i.e. unsorted family X with an operation $X \Gamma \rightarrow (\Gamma \rightsquigarrow \Delta) \rightarrow X \Delta$). With this in place, the meta-algebra structure on $[\mathfrak{X} \multimap \mathcal{M}] \Rightarrow \mathcal{M}$ is given as follows:

$$\begin{aligned} \mathbf{MSub}^\Sigma : (\mathfrak{X} : \mathbf{Family}_s) &\rightarrow \Sigma \mathbf{Mon} \mathcal{M} \rightarrow \mathbf{MetaAlg} \mathfrak{X} ([\mathfrak{X} \multimap \mathcal{M}] \Rightarrow \mathcal{M}) \\ \mathbf{MSub}^\Sigma \mathfrak{X} \mathcal{M}^{\Sigma \mathbf{M}} = \mathbf{record} \{ &\mathbf{alg} = \lambda t \quad \zeta \rightarrow \mathcal{M}. \mathbf{alg} (\mathbf{estr} [\mathfrak{X} \multimap \mathcal{M}. \square]^\square \mathcal{M} t \zeta) \\ &; \mathbf{var} = \lambda v \quad \zeta \rightarrow \mathcal{M}. \eta v \\ &; \mathbf{mvar} = \lambda \mathfrak{m} \varepsilon \zeta \rightarrow \mathcal{M}. \mu (\zeta \mathfrak{m}) (\mathbf{copair} \mathcal{M} (\lambda v \rightarrow \varepsilon v \zeta) \mathcal{M}. \eta) \} \end{aligned}$$

- For \mathbf{alg} , we use \mathbf{estr} to swap the Σ -algebra structure and dependence on the metasubstitution map. We make use of the fact that $[\mathfrak{X} \multimap \mathcal{P}]$ is an unsorted coalgebra if \mathcal{P} is a sorted coalgebra.
- For \mathbf{var} , we ignore ζ entirely and use the point of \mathcal{M} .
- For \mathbf{mvar} , we start by looking up the metavariable $\mathfrak{m} : \mathfrak{X} \tau \Pi$ in the linear metasubstitution map $\zeta : [\mathfrak{X} \multimap \mathcal{M}] \Gamma$, obtaining the term $\zeta \mathfrak{m} : \mathcal{M} \tau (\Pi \dot{+} \Gamma)$ in an extended context. We need to tweak this further to get the required output in $\mathcal{M} \tau \Gamma$, which we do by applying a substitution

$(\Pi \vdash \Gamma) \dashv \text{M} \rightarrow \Gamma$ to $\zeta \mathfrak{m}$. It is constructed as the copairing of the unit $\mathcal{M}.\eta : \Gamma \dashv \text{M} \rightarrow \Gamma$ and the substitution map $\Pi \dashv \text{M} \rightarrow \Gamma$ which looks up a variable in context Π in \mathfrak{m} 's metavariable environment $\varepsilon : \Pi \dashv ([\mathfrak{X} \dashv \mathcal{P}] \Rightarrow \mathcal{M}) \rightarrow \Gamma$ and applies the resulting parametrised term to ζ .

This latter specification is quite a mouthful: metasubstitution (which is derived by initiality into the meta-algebra $\mathbf{MSub}^{\mathfrak{X}} \mathfrak{T}^{\Sigma \mathcal{M}}$) is less a form of “textual replacement” and more an intricate surgical procedure with several interconnected parts. The metavariable-preservation property of `msub` illuminates the role of recursively applying metasubstitution to the elements of the metavariable environment ε , then substituting the terms into the parameters of $\zeta \mathfrak{m}$:

$$\mathbf{msub} (\mathbf{mvar} \mathfrak{m} \varepsilon) \zeta \equiv \mathbf{sub} (\zeta \mathfrak{m}) (\mathbf{copair} (\mathbb{T}\mathfrak{Y})) (\lambda v \rightarrow \mathbf{msub} (\varepsilon v) \zeta) \mathbf{var}$$

As an illustration of metasubstitution, consider the open term $x : \mathbb{N} \vdash \lambda y : \mathbb{N}. \mathfrak{a}\langle x + 1, \mathfrak{b}\langle y \rangle \rangle : \mathbb{N}$ and the metasubstitution mapping $\zeta = (\mathfrak{a}\langle m, n \rangle \mapsto \mathfrak{c}\langle m \rangle \times n; \mathfrak{b}\langle m \rangle \mapsto \mathfrak{c}\langle m + x \rangle)$ in the global context $x : \mathbb{N}$ (note that the latter term includes both the parameter m and the free variable x). The evaluation of the metasubstitution proceeds as follows (where `sub` $t [\cdot \cdot \cdot]$ and `msub` $t \langle \langle \cdot \cdot \cdot \rangle \rangle$ denote substitution and metasubstitution into t , respectively):

$$\begin{aligned} & \mathbf{msub} (\lambda y : \mathbb{N}. \mathfrak{a}\langle x + 1, \mathfrak{b}\langle y \rangle \rangle) \langle \langle \zeta \rangle \rangle \\ \equiv & \lambda y : \mathbb{N}. \mathbf{msub} (\mathfrak{a}\langle x + 1, \mathfrak{b}\langle y \rangle \rangle) \langle \langle \mathbf{wk} \zeta \rangle \rangle & \textcircled{1} \\ \equiv & \lambda y : \mathbb{N}. \mathbf{sub} (\mathfrak{c}\langle m \rangle \times n) [m \mapsto x + 1, n \mapsto \mathbf{msub} (\mathfrak{b}\langle y \rangle) \langle \langle \mathbf{wk} \zeta \rangle \rangle] & \textcircled{2} \\ \equiv & \lambda y : \mathbb{N}. \mathbf{sub} (\mathfrak{c}\langle m \rangle \times n) [m \mapsto x + 1, n \mapsto \mathbf{sub} (\mathfrak{c}\langle m + x \rangle) [m \mapsto y]] & \textcircled{3} \\ \equiv & \lambda y : \mathbb{N}. \mathbf{sub} (\mathfrak{c}\langle m \rangle \times n) [m \mapsto x + 1, n \mapsto \mathfrak{c}\langle y + x \rangle] & \textcircled{4} \\ \equiv & \lambda y : \mathbb{N}. \mathfrak{c}\langle x + 1 \rangle \times \mathfrak{c}\langle y + x \rangle & \textcircled{5} \end{aligned}$$

In step $\textcircled{1}$ we push the metasubstitution under the binder. A crucial component of this step is the weakening of the terms in the metasubstitution mapping ζ represented here by `wk` ζ . Indeed, since the local context changes from $x : \mathbb{N}$ to $y : \mathbb{N}, x : \mathbb{N}$, the de Bruijn index of the variable x has to be shifted without altering the parameters m, n (for example, $m : \mathbb{N}, x : \mathbb{N} \vdash \mathfrak{c}\langle m + x \rangle$ is renamed to $m : \mathbb{N}, y : \mathbb{N}, x : \mathbb{N} \vdash \mathfrak{c}\langle m + x \rangle$ by shifting the index of x). Although this modification makes no difference with the named-variable representation displayed here, in practice it becomes an explicit application of `ren` – implemented as part of the exponential strength for δ . It is worth noting the delicate interplay between `ren`, `sub`, and `msub`. In step $\textcircled{2}$ we apply the metasubstitution to \mathfrak{a} by looking up the term $\mathfrak{c}\langle m \rangle \times n$, substituting the contents of \mathfrak{a} 's metavariable environment for m and n , and recursively metasubstituting `wk` ζ into $\mathfrak{b}\langle y \rangle$ (and into $x + 1$, where it is a no-op). The mappings of variables x, y to themselves are omitted. Steps $\textcircled{3}$ and $\textcircled{4}$ evaluate the recursive calls, applying an object-level substitution to $\mathfrak{c}\langle m + x \rangle$ that replaces the parameter m with the variable y (and the global variable x with itself). The final substitution is evaluated at step $\textcircled{5}$.

3.7 Equational systems

An immediate application of metasubstitution is building generic *equational systems* for second-order languages [Fiore and Hur 2010; Fiore and Mahmoud 2010]. By specifying the axioms of an equational theory with the use of metavariables, one can use metasubstitution to extract axiom instances between terms and apply rewrite rules within compound expressions. For example, β -equivalence arises from instances of the axiom

$$\mathfrak{b} : [\alpha]\beta, \mathfrak{a} : [\alpha] \triangleright \emptyset \vdash (\lambda x : \alpha. \mathfrak{b}\langle x \rangle) \mathfrak{a}\langle \rangle \approx \mathfrak{b}\langle \mathfrak{a}\langle \rangle \rangle : \beta$$

with terms (of the appropriate types and contexts) metasubstituted for the metavariables \mathfrak{a} and \mathfrak{b} . Generic equality can be then directly encoded in Agda as the smallest equivalence relation closed under a given relation `Axiom` : $\forall (\mathfrak{X} \Gamma \{\alpha\}) \rightarrow \mathbb{T} \mathfrak{X} \alpha \Gamma \rightarrow \mathbb{T} \mathfrak{X} \alpha \Gamma \rightarrow \mathbf{Set}$ and metasubstitution:

$\text{data } _ \triangleright _ \vdash _ \approx _ : (\mathfrak{X} : \text{Family}_s)(\Gamma : \text{Ctx})\{\alpha : T\} \rightarrow \mathbb{T} \mathfrak{X} \alpha \Gamma \rightarrow \mathbb{T} \mathfrak{X} \alpha \Gamma \rightarrow \text{Set}_1$ **where**
 $\text{eq} : t \equiv s \rightarrow \mathfrak{X} \triangleright \Gamma \vdash t \approx s$
 $\text{sy} : \mathfrak{X} \triangleright \Gamma \vdash t \approx s \rightarrow \mathfrak{X} \triangleright \Gamma \vdash s \approx t$
 $\text{tr} : \mathfrak{X} \triangleright \Gamma \vdash t \approx s \rightarrow \mathfrak{X} \triangleright \Gamma \vdash s \approx u \rightarrow \mathfrak{X} \triangleright \Gamma \vdash t \approx u$
 $\text{ax} : \text{Axiom } \mathfrak{X} \Gamma t s \rightarrow \mathfrak{X} \triangleright \Gamma \vdash t \approx s$
 $\text{ms} : \mathfrak{X} \triangleright \Gamma \vdash t \approx s \rightarrow (\zeta \xi : [\mathfrak{X} \rightarrow \mathbb{T} \mathfrak{Y}] \Gamma) \rightarrow$
 $(\forall \{\tau \Pi\}(m : \mathfrak{X} \tau \Pi) \rightarrow \mathfrak{Y} \triangleright \Pi \vdash \Gamma \vdash \zeta m \approx \xi m) \rightarrow \mathfrak{Y} \triangleright \Gamma \vdash \text{msub } t \zeta \approx \text{msub } s \xi$

The `ms` constructor expresses that whenever two terms t and s are equivalent, and two instantiations for their metavariable context ζ and ξ are equivalent (for every metavariable), then performing the metasubstitution also gives equivalent terms. Using the equivalence constructors one can derive useful proof combinators and a library for equational reasoning; for example, `ax \approx` equates two terms via the instantiation of an axiom:

$\text{ax}\approx : \text{Axiom } \mathfrak{X} \Gamma t s \rightarrow (\zeta : [\mathfrak{X} \rightarrow \mathbb{T} \mathfrak{Y}] \Gamma) \rightarrow \mathfrak{Y} \triangleright \Gamma \vdash \text{msub } t \zeta \approx \text{msub } s \zeta$
 $\text{ax}\approx a \zeta = \text{ms} (\text{ax } a) \zeta \zeta (\lambda _ \rightarrow \text{eq refl})$

The biggest gains, however, come from not having to tediously encode the congruence rules for every subterm of every term of the syntax. To rewrite a deeply nested subexpression, we simply mark its location in the term with a “typed hole” implemented as a distinguished metavariable \odot , and use `ms` to instantiate it with the two sides of an equality rule: for example, $f \approx g$ and $(\odot a) \approx (\odot a)$ imply that $f a = \text{msub } (\odot a) (\lambda \{\odot\} \rightarrow f) \stackrel{\text{ms}}{\approx} \text{msub } (\odot a) (\lambda \{\odot\} \rightarrow g) = g a$. Further examples of this and other combinators are given in [Section 5](#). An important future development is generically proving the soundness and completeness of second-order equational logic.

We thus conclude our abstract development of initial algebra semantics and move on to the discussion of second-order signatures.

4 GENERIC SIGNATURES

The abstract development discussed so far was entirely generic over the second-order signature and term syntax. In this section we discuss how endofunctors Σ are constructed from descriptions of syntax signatures, the benefits and drawbacks of various term representations, and how we leverage code generation to turn our library into a practical framework for language formalisation. Thanks to our modular implementation, we have several choices in each of the following matters:

- how to encode the signature of a second-order syntax (e.g. binding algebras [\[Fiore et al. 1999\]](#), indexed containers [\[Altenkirch et al. 2015\]](#), AACMM [\[2021\]](#)-style [Descriptions](#));
- how to convert the signature into a Fam_s endofunctor Σ (e.g. polynomial functors [\[Fiore 2012; Arkor and Fiore 2020\]](#), higher- or first-order argument collections, [Desc](#) interpretations);
- how to define the `data` type for the initial (Σ, \mathfrak{X}) -meta-algebra (implicit or explicit encodings).

Each of these have their benefits and drawbacks, and we identify three choices that work particularly well together and combine convenience, flexibility, and appropriate computational behaviour.

4.1 Binding signatures

Binding signatures were introduced by [Aczel \[1978\]](#) as a generalisation of standard algebraic signatures to languages with variable binding. Our formalisation will use the typed variant of the notion as given in [\[Fiore and Hur 2010\]](#):

Definition 4.1. A second-order signature $\Sigma = (T, O, |-|)$ is specified by a set of sorts T , a set of operators O , and an arity function $|-| : O \rightarrow \text{List } ((\text{List } T) \times T) \times T$.

For an operator o , the tuple $|o| = ((\vec{\alpha}_1, \beta_1), (\vec{\alpha}_2, \beta_2), \dots, (\vec{\alpha}_n, \beta_n), \tau)$ consists of the output sort $\tau \in T$, and a list of n arguments of type β_i , with each argument binding variables as listed in $\vec{\alpha}_i$. We write this more concisely as $o: [\vec{\alpha}_1]\beta_1, \dots, [\vec{\alpha}_n]\beta_n \rightarrow \tau$, omitting empty bound variable lists.

Example 4.2. The second-order signature of the simply-typed λ -calculus over a base type has the set of types T generated inductively from a base type N and a binary function type \succ , and the two type-indexed families of operators

$$\text{app}_{\alpha, \beta} : (\alpha \succ \beta), \alpha \rightarrow \beta \quad \text{lam}_{\alpha, \beta} : [\alpha]\beta \rightarrow (\alpha \succ \beta)$$

This definition of signatures can be adapted to Agda almost verbatim, using `Ctx` in place of `List T`:

```
record Signature (O : Set) : Set₁ where
  constructor sig
  field |_| : O → List (Ctx × T) × T
  Arity : O → List (Ctx × T)
  Arity o = proj₁ | o |
  Sort : O → T
  Sort o = proj₂ | o |
```

The set T of types (or sorts) is normally given as an inductive data type, and O as an enumeration of operators. For example, the STLC has the following sorts and operator symbols:

```
data ΛT : Set where
  N : ΛT
  _>_ : ΛT → ΛT → ΛT
data Λ₀ : Set where
  appₒ lamₒ : {α β : ΛT} → Λ₀
```

The `Signature` instances are direct translations of [Example 4.2](#) above. One can use some simple shorthands for specifying arguments and bound variables to make the declaration concise.

```
Λ:Sig : Signature Λ₀
Λ:Sig = sig λ where (appₒ {α}{β}) → (t₀ α > β), (t₀ α)   ↦₂ β
                  (lamₒ {α}{β}) → (α t₁ β)                ↦₁ α > β
```

4.2 Signature endofunctor

The signature contains all the information needed to determine the syntactic structure of a language. To make use of the abstract development in [Section 3](#), we need to convert a `Signature` into a sorted-family endofunctor Σ , which captures the way in which constructors of the syntax are associated with arguments. For example, given the signature of the STLC, an element of $\Sigma X \beta \Gamma$ may be the operator `app` associated with two X -terms $f: X (\alpha \succ \beta) \Gamma$ and $a: X \alpha \Gamma$, while an element of $\Sigma X (\alpha \succ \beta) \Gamma$ may be the operator `lam` with a term $b: X \beta (\alpha \cdot \Gamma)$.

For technical reasons, that we will expand upon later, we choose to represent the “collection” of arguments of an operator as a tuple of terms. An alternative would be a higher-order encoding as a mapping from an argument index to a term (similar to the implementation of substitutions as context maps); however, even though constructing the `Strength` for such a representation would be easier, it complicates the initiality proof which we wish to keep as simple as possible.

```
Arg : List (Ctx × T) → Familyₛ → Family
Arg [] X Γ = ⊤
Arg ((Θ , τ) :: as) X Γ = δ Θ X τ Γ × Arg as X Γ
```

Note the use of the context extension endofunctor δ : it is used to add the new variables bound by an argument to the the global context, making all variables available in the body of the binder.

We are ready to give the definition of the signature endofunctor for a signature $(T, O, \text{Arity}, \text{Sort})$:

$$\begin{aligned} \Sigma &: \text{Family}_s \rightarrow \text{Family}_s \\ \Sigma \mathcal{X} \alpha \Gamma &= \Sigma[o \in O] (\alpha \equiv \text{Sort } o \times \text{Arg } (\text{Arity } o) \mathcal{X} \Gamma) \end{aligned}$$

An element of the set $\Sigma \mathcal{X} \alpha \Gamma$ is a dependent tuple consisting of an operator symbol $o \in O$, a proof that the output sort of the operator is α , and a tuple of \mathcal{X} -terms for each operand of the operator of the type and extension context given by the operator arity. For example, an element of $\Sigma \mathcal{X} \beta \Gamma$ is $(\text{app}, \text{refl}, (f, t, \text{tt}))$, for terms $f: \mathcal{X} (\alpha \succ \beta) \Gamma$ and $t: \mathcal{X} \alpha \Gamma$. One can suppress the tt for operators of positive arity by adding a case for a singleton argument list in the definition of Arg , and use a pattern synonym [Pickering et al. 2016] to hide the refl element, writing $\text{app}: (f, t)$ for the above.

The only other construction we need is the **Strength** instance for Σ . The family \mathcal{X} in a $(\Sigma \mathcal{X})$ -term is only used in the argument list, so the Σ -strength can be easily derived from the Arg -strength $\text{Arg as } (\mathcal{P}, \mathcal{X}) \rightarrow (\mathcal{P}, \text{Arg as } \mathcal{X})$. Applying strength to the tuple of arguments simply applies it to every component of type $\delta \Theta (\mathcal{P}, \mathcal{X}) \tau \Gamma$ – and this is nothing but the strength instance for δ which we constructed back in Section 2.3.2.

$$\begin{aligned} \text{str}^A &: (\mathcal{P}_*^\square : \text{Coalg}_* \mathcal{P})(\mathcal{X} : \text{Family}_s)(\text{as} : \text{List } (\text{Ctx} \times T)) \rightarrow \text{Arg as } (\mathcal{P}, \mathcal{X}) \rightarrow (\mathcal{P}, \text{Arg as } \mathcal{X}) \\ \text{str}^A \mathcal{P}_*^\square \mathcal{X} [] & \quad \text{tt} \quad \sigma = \text{tt} \\ \text{str}^A \mathcal{P}_*^\square \mathcal{X} ((\Theta, \tau) :: \text{as}) (h, \text{hs}) & \sigma = (\delta\text{-Str.str } \Theta \mathcal{P}_*^\square \mathcal{X} h \sigma), (\text{str}^A \mathcal{P}_*^\square \mathcal{X} \text{ as } \text{hs } \sigma) \end{aligned}$$

The strength laws are similarly established by pointwise application of the appropriate $\delta\text{-Str}$ fields to the elements of the argument tuple. Extending Arg -strength to Σ is easy, since the operation does not modify the operator or sort equality proof. In addition to $\Sigma\text{-Str}$ below, we also have an instance of exponential strength for Σ derived via weakening.

$\Sigma\text{-Str} : \text{Strength } \Sigma F$

$$\begin{aligned} \Sigma\text{-Str} &= \text{record} \{ \text{str} = \lambda \mathcal{P}_*^\square \mathcal{X} (o, e, a) \sigma \rightarrow o, e, (\text{str}^A \mathcal{P}_*^\square \mathcal{X} (\text{Arity } o) a \sigma) \\ & \quad ; \text{str-nat}_1 = \lambda f_*^\square \Rightarrow (o, e, a) \sigma \rightarrow \text{cong } (o, e, _) (\text{str}^A\text{-nat}_1 f_*^\square \Rightarrow (\text{Arity } o) a \sigma) \\ & \quad ; \text{str-nat}_2 = \lambda g \quad (o, e, a) \sigma \rightarrow \text{cong } (o, e, _) (\text{str}^A\text{-nat}_2 g \quad (\text{Arity } o) a \sigma); \dots \} \end{aligned}$$

4.3 Term syntax

The final piece of the puzzle is constructing the initial (Σ, \mathfrak{X}) -meta-algebra \mathbb{T} from a second-order signature endofunctor. Such initial algebras correspond to inductive **data** types whose constructors combine \mathbb{T} -terms into other \mathbb{T} -terms, allowing for arbitrarily nested syntactic structure.

We have several choices in the approach we take. One is to treat the tuples $(\text{op} : (a_1, \dots, a_n))$ as the terms of the syntax, directly encoding the Σ -algebra structure as a unified term constructor $\text{con} : \Sigma \mathbb{T} \rightarrow \mathbb{T}$; the other is the more common implementation with one constructor for each operator, applicable for signatures with a finite set of operators.

Implicit encoding. Alongside the Σ -algebraic structure, terms of a second-order syntax also have to include constructors for variables and metavariables. This suggests the following generic data type of terms for an arbitrary signature:

$$\begin{aligned} \text{data } \mathbb{T} \text{m} : \text{Family}_s \text{ where} & & \text{data } \text{Sub } (\mathcal{X} : \text{Family}_s) : \text{Ctx} \rightarrow \text{Ctx} \rightarrow \text{Set} \text{ where} \\ \text{con} &: \Sigma \mathbb{T} \text{m } \tau \Gamma & \rightarrow \mathbb{T} \text{m } \tau \Gamma & \bullet &: \text{Sub } \mathcal{X} \emptyset \Gamma \\ \text{var} &: \mathcal{I} \tau \Gamma & \rightarrow \mathbb{T} \text{m } \tau \Gamma & _ \leftarrow _ &: \mathcal{X} \alpha \Gamma \rightarrow \text{Sub } \mathcal{X} \Pi \Gamma \rightarrow \text{Sub } \mathcal{X} (\alpha \cdot \Pi) \Gamma \\ \text{mvar} &: \mathfrak{X} \tau \Pi \rightarrow \text{Sub } \mathbb{T} \text{m } \Pi \Gamma \rightarrow \mathbb{T} \text{m } \tau \Gamma & & & \end{aligned}$$

Note the use of **Sub** (defined on the right above) in place of a context map to represent the metavariable environment. It is a first-order, inductive encoding of a simultaneous substitution of terms in context Γ for every variable in context Π , which, while isomorphic to context maps (with conversion functions **lookup** and **tabulate**), is a more appropriate choice for syntax. Recalling that

metasubstitution recurses into the metavariable environment, in a higher-order representation the recursive call would get suspended in the body of a λ -abstraction and lead to terms with unevaluated expressions. With context maps, the application $\text{sem}(\text{mvar } m(\lambda\{\text{new} \rightarrow \text{con}(\text{op} : t)\}))$ would only normalise to $\text{mvar } m(\lambda\{\text{new} \rightarrow \text{sem}(\text{con}(\text{op} : t))\})$, and sem would not get pushed further under the constructor con unless the environment function is applied to new . In contrast, using the first-order encoding Sub , each element of the sequence gets fully normalised, so, as desired, $\text{sem}(\text{mvar } m(\text{con}(\text{op} : t) \blacktriangleleft \bullet))$ reduces to $\text{mvar } m(\text{con}(\text{op} : \text{sem } t) \blacktriangleleft \bullet)$. The same reasoning was behind our choice to represent operator arguments as tuples, rather than higher-order assignments – though the strength instance would have been simpler, specifying arguments via case-analysis is cumbersome, and syntactic operations only recurse one layer deep into subterms.

The proof of initiality of Tm asks us to define a recursive function $\text{sem} : \text{Tm} \rightarrow \mathcal{A}$ for any meta-algebra \mathcal{A} , translating constructors of Tm to the algebra and (meta)variable structure of \mathcal{A} . There is an issue with the obvious definition: recursively interpreting the subterms of a constructor involves mapping sem over a tuple of terms, which Agda does not recognise as a terminating function call. Allais et al. [2021, Section 4] also encountered this problem and proposed the use of Agda’s *sized types* [Abel 2010] (which, unfortunately, are logically unsound in the current Agda version 2.6.2) to mark a term as strictly “larger” than its subterms. The workaround suits their needs of defining sem – at the expense of having to carry around the size index everywhere – but it does not extend to proving the uniqueness of sem which our framework closely relies on (see Pitts [2019] for the technical details of the same issue in the context of general F -algebras).

The solution to this issue is surprisingly simple: instead of recursively interpreting the subterms using a functorial mapping $(\mathcal{X} \rightarrow \mathcal{Y}) \rightarrow (\text{Arg as } \mathcal{X} \Gamma \rightarrow \text{Arg as } \mathcal{Y} \Gamma)$ (and similarly for Sub), we inline the transformations as the mutually recursive functions \mathbb{A} and \mathbb{S} that apply $\text{sem} : \text{Tm} \rightarrow \mathcal{A}$ to subterms directly. The termination checker is satisfied without the need for sized types!

$$\begin{array}{ll}
 \mathbb{A} : \forall \text{as} \rightarrow \text{Arg as } \text{Tm } \Gamma \rightarrow \text{Arg as } \mathcal{A} \Gamma & \mathbb{S} : \text{Sub Tm } \Pi \Gamma \rightarrow \Pi \text{ } \neg [\mathcal{A}] \rightarrow \Gamma \\
 \mathbb{A} [] \quad \text{tt} = \text{tt} & \mathbb{S} (t \blacktriangleleft \sigma) \text{new} = \text{sem } t \\
 \mathbb{A} (a :: \text{as}) (t, \text{ts}) = (\text{sem } t, \mathbb{A} \text{ as } \text{ts}) & \mathbb{S} (t \blacktriangleleft \sigma) (\text{old } v) = \mathbb{S} \sigma v \\
 \\
 \text{sem} (\text{con } (o, e, a)) = \text{alg } (o, e, \mathbb{A} (\text{Arity } o) a) \\
 \text{sem} (\text{var } v) = \text{var } v \\
 \text{sem} (\text{mvar } m \ \varepsilon) = \text{mvar } m (\mathbb{S} \ \varepsilon)
 \end{array}$$

The proof that sem is a unique (Σ, \mathfrak{X}) -meta-algebra homomorphism is also quite straightforward, only requiring a few mutually inductive lemmas about \mathbb{A} and \mathbb{S} . We then have the following result:

THEOREM 4.3. *Tm is an initial (Σ, \mathfrak{X}) -meta-algebra.*

We can thus instantiate our previous metatheory with this concrete data type to get access to substitution and its correctness properties, sound compositional interpretations in models, etc.

Explicit encoding. The implicit encoding achieves our conceptual goals: it is a first-order initial meta-algebra for an arbitrary second-order signature. Its main practical disadvantage is that the term syntax is closely tied to the algebraic framework and forces users to adopt a cumbersome encoding of terms, rather than the natural and elegant “one constructor for each typing rule” principle of intrinsic typing. Pattern synonyms may be used to simplify the surface syntax, though we found them quite fragile when working with parametrised modules – not to mention that pattern synonyms are untyped and feel inherently “hacky” for something as important as the terms of a formalised language. Ideally, we would like users to be able to adopt our framework as seamlessly as possible, perhaps even plugging it into an existing formalisation without changing the fundamental data type of syntactic terms.

This is very much possible, thanks to our rigid separation between signatures, endofunctors and initial meta-algebras. Defining the initial meta-algebra instance for an existing data type is not an especially laborious task: one needs a recursive initial interpretation function, a homomorphism proof, and an inductive uniqueness proof. Metavariables still require the mutually recursive transformation \mathbb{S} (defined as before and so omitted here), but now that one is manually recursing into subterms, the transformation \mathbb{A} is not needed.

$\text{data } \Lambda : \text{Family}_s \text{ where}$ $\text{var} : \mathcal{I} \tau \Gamma \rightarrow \Lambda \tau \Gamma$ $\text{mvar} : \mathcal{X} \tau \Pi \rightarrow \text{Sub } \Lambda \Pi \Gamma \rightarrow \Lambda \tau \Gamma$ $\text{app} : \Lambda (\alpha \succ \beta) \Gamma \rightarrow \Lambda \alpha \Gamma \rightarrow \Lambda \beta \Gamma$ $\text{lam} : \Lambda \beta (\alpha \cdot \Gamma) \rightarrow \Lambda (\alpha \succ \beta) \Gamma$	$\text{sem} : \Lambda \rightarrow \mathcal{A}$ $\text{sem} (\text{var } v) = \text{var } v$ $\text{sem} (\text{mvar } m \ \varepsilon) = \text{mvar } m (\mathbb{S} \ \varepsilon)$ $\text{sem} (\text{app } g \ a) = \text{alg} (\text{app}_o : \text{sem } g, \text{sem } a)$ $\text{sem} (\text{lam } b) = \text{alg} (\text{lam}_o : \text{sem } b)$
--	--

The homomorphism instance uses a simple lemma $\mathbb{S}\text{-tab}$ about the interaction of \mathbb{S} and `tabulate` (for all context maps ε , $\mathbb{S} (\text{tabulate } \varepsilon) = \text{sem} \circ \varepsilon$), and the Σ -algebra homomorphism proof, which is satisfied merely by pattern-matching on the operand and sort equality proof.

$$\begin{aligned} \text{sem}^\Sigma &\Rightarrow : \text{MetaAlg} \Rightarrow \Lambda^\Sigma \mathcal{A}^\Sigma \text{sem} \\ \text{sem}^\Sigma &= \text{record} \{ \langle \text{alg} \rangle = \lambda \{t = t\} \rightarrow \langle \text{alg} \rangle t ; \langle \text{var} \rangle = \text{refl} \\ &\quad ; \langle \text{mvar} \rangle = \lambda \{m = m\} \{\varepsilon\} \rightarrow \text{cong} (\text{mvar } m) (\mathbb{S}\text{-tab } \varepsilon) \} \\ \text{where } \langle \text{alg} \rangle &: (t : \Sigma \Lambda \alpha \Gamma) \rightarrow \text{sem} (\Lambda^\Sigma . \text{alg } t) \equiv \mathcal{A}^\Sigma . \text{alg} (\Sigma_1 \text{sem } t) \\ &\quad \langle \text{alg} \rangle (\text{app}_o : _) = \text{refl} \\ &\quad \langle \text{alg} \rangle (\text{lam}_o : _) = \text{refl} \end{aligned}$$

The uniqueness proof – that `sem` is equal to any meta-algebra homomorphism $g : \Lambda \rightarrow \mathcal{A}$ – involves the mutually inductive lemma $\mathbb{S}\text{-lu}$ and the inverse property of `tabulate` and `lookup` in the metavariable case; everything else follows from the homomorphism properties of g .

$$\begin{aligned} \mathbb{S}\text{-lu} &: (\varepsilon : \text{Sub } \Lambda \Pi \Gamma) (v : \mathcal{I} \alpha \Pi) \rightarrow \mathbb{S} \ \varepsilon \ v \equiv g (\text{lookup } \varepsilon \ v) \\ \mathbb{S}\text{-lu} (t \triangleleft \varepsilon) \text{new} &= \text{sem}! \ t \\ \mathbb{S}\text{-lu} (t \triangleleft \varepsilon) (\text{old } v) &= \mathbb{S}\text{-lu } \varepsilon \ v \\ \text{sem}! &: (t : \Lambda \alpha \Gamma) \rightarrow \text{sem } t \equiv g \ t \\ \text{sem}! (\text{var } v) &= \text{sym } \langle \text{var} \rangle \\ \text{sem}! (\text{mvar } m \ \varepsilon) \text{rewrite } \mathbb{S}\text{-lu } \varepsilon &= \text{trans} (\text{sym } \langle \text{mvar} \rangle) (\text{cong} (g \circ \text{mvar } m) (\text{tabulate } \varepsilon)) \\ \text{sem}! (\text{app } f \ a) \text{rewrite } \text{sem}! \ f \mid \text{sem}! \ a &= \text{sym } \langle \text{alg} \rangle \\ \text{sem}! (\text{lam } b) \text{rewrite } \text{sem}! \ b &= \text{sym } \langle \text{alg} \rangle \end{aligned}$$

At the expense of minimal extra boilerplate, one is able to prove that the inductively defined family Λ is an initial meta-algebra, and instantiate our metatheory with this result. The `Theory` module associated with an initial meta-algebra exports every definition and lemma given in the framework for easy access; for example, the coveted single-variable substitution operation is directly accessible as `Theory.[_]` : $\Lambda \alpha \Gamma \rightarrow \Lambda \beta (\alpha \cdot \Gamma) \rightarrow \Lambda \beta \Gamma$, with no extra work required.

4.4 Code generation

The next logical step is eliminating the need to write boilerplate altogether and allow users to go from a direct specification of the second-order signature to the formalised metatheory of an explicitly encoded data type of terms automatically. Since the initiality proof for the explicit encoding is rather formulaic and much of it is signature-independent, one may as well generate the associated Agda code from a syntax description using a simple Python script. This signature-to-Agda compiler takes a simple textual specification of a type and term signature, and produces the Agda code for the `Signature` declaration and initiality proof outlined in Sections 4.1 and 4.3. For instance,

one can give the signature of our running STLC example as follows (recall the introduction), also optionally specifying infix symbols and fixity information:

type	term
$N : 0\text{-ary}$	$\text{app} : \alpha \succ \beta \quad \alpha \rightarrow \beta \quad _ \$ _ l20$
$_ \succ _ : 2\text{-ary} \mid r30$	$\text{lam} : \alpha.\beta \quad \rightarrow \alpha \succ \beta \quad \lambda _ r10$

Saving this in a file `stlc` and running `python soas.py stlc` produces the files `Signature.agda` and `Syntax.agda` with all the required imports, declarations of the signature ΛSig and explicit inductive data type of terms Λ , and the initiality proof. One is thereby free to proceed with the interesting parts of language formalisation like operational and denotational semantics right away.

The compiler is a fairly straightforward Python program that parses the specification language and produces formatted Agda files using the built-in string templating system. Its simplicity is one of its advantages: the boilerplate code it generates is systematic and minimal (1-2 lines of code for every type constructor and 6-7 lines for every term constructor), so manual testing on a wide range of examples gives us sufficient confidence in the robustness and correctness of the script. Unlike most other code generation solutions (some listed in [Section 6.1](#)) that produce the entire syntax-specific formalisation, including fragile and impenetrable de Bruijn arithmetic proofs, we leverage the fully syntax-generic metatheory implemented in the library and generate just enough code (namely the initiality proof, which is less boilerplate and more an elegant categorical argument) to instantiate it. Consequently, the compiler output is concise, readable, and easy to maintain.

Having examined the generic metatheory, construction of signatures, and the term syntax, we conclude with some demonstrations of the framework used in practice.

5 EXAMPLES

Our framework is flexible and unopinionated: it can be plugged into any intrinsically-typed formalisation of a second-order calculus and equip the syntax with the often-needed operations of weakening and substitution, and their corresponding laws. It also helps users in defining evaluation functions, interpreters, and syntactic translations in a concise and provably-correct manner. In this section, we give two extended examples of how the library may be used.

5.1 Computational calculi

The STLC has been our running example throughout the paper, and our framework is a great playground for experimenting with various extensions of it, be it with new types, terms, or equations. Below is a list of constructs that can be easily represented and compiled into Agda. The semantics of such a language would be rather complicated, but its syntax is still just a second-order signature.

type	term		
$N : 0\text{-ary}$	$\text{app} : \alpha \succ \beta \quad \alpha \rightarrow \beta$	$\rightarrow \alpha \otimes \beta$	$\text{pair} : \alpha \quad \beta \rightarrow \alpha \otimes \beta$
$_ \succ _ : 2\text{-ary}$	$\text{lam} : \alpha.\beta \quad \rightarrow \alpha \succ \beta$	$\rightarrow \alpha$	$\text{fst} : \alpha \otimes \beta \rightarrow \alpha$
$_ \otimes _ : 2\text{-ary}$	$\text{let} : \alpha \quad \alpha.\beta \rightarrow \beta$	$\rightarrow \beta$	$\text{snd} : \alpha \otimes \beta \rightarrow \beta$
$_ \oplus _ : 2\text{-ary}$	$\text{fix} : \alpha \succ \alpha \rightarrow \alpha$	$\rightarrow \alpha \oplus \beta$	$\text{inl} : \alpha \rightarrow \alpha \oplus \beta$
$\neg _ : 1\text{-ary}$	$\text{throw} : \alpha \quad \neg \alpha \rightarrow \beta$	$\rightarrow \alpha \oplus \beta$	$\text{inr} : \beta \rightarrow \alpha \oplus \beta$
$T : 1\text{-ary}$	$\text{callcc} : (\neg \alpha).\alpha \rightarrow \alpha$	$\rightarrow \gamma$	$\text{case} : \alpha \oplus \beta \quad \alpha.\gamma \quad \beta.\gamma \rightarrow \gamma$
	$\text{return} : \alpha \rightarrow T \alpha$	$\rightarrow N$	$\text{ze} : \rightarrow N$
	$\text{bind} : T \alpha \quad \alpha.(T \beta) \rightarrow T \beta$	$\rightarrow N$	$\text{su} : N \rightarrow N$
		$\rightarrow \alpha$	$\text{nrec} : N \quad \alpha \quad (\alpha, N).\alpha \rightarrow \alpha$

We will use the minimal fragment of STLC (with `app` and `lam`) to showcase the construction of models and interpretations. The CCC model of the STLC [Lambek 1980] in the category **Set** interprets types as sets, and terms $\Gamma \vdash t : \alpha$ as functions from the interpretation of Γ to the interpretation of α . Interpretation of contexts can be higher-order or first-order (as a Cartesian product of type interpretations) – with the higher-order encoding the model definition is remarkably concise.

```

[ ] :  $\Lambda \Gamma \rightarrow \text{Set}$ 
[  $\mathbb{N}$  ] =  $\mathbb{N}$ 
[  $\alpha \succ \beta$  ] = [  $\alpha$  ]  $\rightarrow$  [  $\beta$  ]

[ ]c :  $\text{Ctx} \rightarrow \text{Set}$ 
[  $\Gamma$  ]c =  $\forall \{\alpha\} \rightarrow \mathcal{I} \alpha \Gamma \rightarrow [ \alpha ]$ 

_+_: [  $\alpha$  ]  $\rightarrow$  [  $\Gamma$  ]c  $\rightarrow$  [  $\alpha \cdot \Gamma$  ]c
( $a^+ \gamma$ ) new =  $a$ 
( $a^+ \gamma$ ) old  $v$  =  $\gamma v$ 

Env $\Sigma$  :  $\Sigma \text{Mon Env}$ 
Env $\Sigma$  = record
{  $\text{M} = \text{record} \{ \eta = \lambda v \gamma \rightarrow \gamma v ; \mu = \lambda t \sigma \delta \rightarrow t (\lambda v \rightarrow \sigma v \delta)$ 
; lunit = refl ; runit = refl ; assoc = refl }
; alg =  $\lambda \{ (\text{app}_o : f, a) \gamma \rightarrow f \gamma (a \gamma)$ 
;  $(\text{lam}_o : b) \gamma \rightarrow \lambda a \rightarrow b (a^+ \gamma) \}$ 
;  $\mu(\text{alg}) = \lambda \{ (\text{app}_o : \_) \rightarrow \text{refl}$ 
;  $(\text{lam}_o : b) \rightarrow \text{ext}^2 \lambda \delta a \rightarrow \text{cong } b (\text{dext}$ 
 $\lambda \{ \text{new} \rightarrow \text{refl} ; \text{old } v \rightarrow \text{refl} \} \}$  }

module Env = FreeMonoid  $\emptyset$  Env $\Sigma$  ( $\lambda ()$ )

Env :  $\text{Family}_s$ 
Env  $\alpha \Gamma = [ \Gamma ]^c \rightarrow [ \alpha ]$ 
eval :  $\Lambda_0 \rightarrow \text{Env}$ 
eval = Env.sem

```

Here we restrict to the sorted family $\Lambda_0 = \Lambda \emptyset$ of λ -terms without metavariables for simplicity. The `eval` function interprets λ -terms as Agda programs; for example, `eval ($\lambda \lambda x_1$) ($\lambda ()$)` (where the 1st de Bruijn index `var` (`old new`) is denoted x_1) normalises to the Agda function $\lambda x y \rightarrow x$. Since it is derived by initiality, the interpretation is compositional and satisfies the semantic substitution lemma *by construction*. This is an enormous time-saver, since proving the soundness of substitution is often one of the most tedious steps required for the development of denotational semantics – the binding terms are the usual suspects, forcing one to reason about semantics of lifting, weakening, renaming, etc. Our framework does all the heavy lifting, allowing users to move on to less bureaucratic proofs. For example, after defining the predicate `Val` satisfied by value terms of the form λb for $b : \Lambda_0 \beta (\alpha \cdot \Gamma)$, it takes minimal effort to equip the language with an intrinsically-typed call-by-value reduction relation and a proof that it preserves the interpretation of terms:

```

data  $\rightsquigarrow$  :  $\Lambda_0 \alpha \Gamma \rightarrow \Lambda_0 \alpha \Gamma \rightarrow \text{Set}$  where
   $\zeta$ - $\$1$  : { $f g : \Lambda_0 (\alpha \succ \beta) \Gamma$ } { $a : \Lambda_0 \alpha \Gamma$ }  $\rightarrow f \rightsquigarrow g \rightarrow f \$ a \rightsquigarrow g \$ a$ 
   $\zeta$ - $\$2$  : { $f : \Lambda_0 (\alpha \succ \beta) \Gamma$ } { $a b : \Lambda_0 \alpha \Gamma$ }  $\rightarrow \text{Val } f \rightarrow a \rightsquigarrow b \rightarrow f \$ a \rightsquigarrow f \$ b$ 
   $\beta$ - $\lambda$  : { $t : \Lambda_0 \alpha \Gamma$ } { $b : \Lambda_0 \beta (\alpha \cdot \Gamma)$ }  $\rightarrow \text{Val } t \rightarrow (\lambda b) \$ t \rightsquigarrow [ t / ] b$ 

sound : { $t s : \Lambda_0 \alpha \Gamma$ }  $\rightarrow t \rightsquigarrow s \rightarrow (\gamma : [ \Gamma ]^c) \rightarrow \text{eval } t \gamma \equiv \text{eval } s \gamma$ 
sound ( $\zeta$ - $\$1$   $r$ )  $\gamma$  rewrite sound  $r \gamma = \text{refl}$ 
sound ( $\zeta$ - $\$2$   $r$ )  $\gamma$  rewrite sound  $r \gamma = \text{refl}$ 
sound ( $\beta$ - $\lambda$  { $t$ } { $b$ }  $\_$ )  $\gamma$  rewrite Env.sub-lemma  $t b$ 
= cong (eval  $b$ ) (dext  $\lambda \{ \text{new} \rightarrow \text{refl} ; \text{old } v \rightarrow \text{refl} \}$ )

```

Note the use of the freely-obtained single-variable substitution $[t /] b$ in the β - λ axiom, and the invocation of `sub-lemma` which translates `eval ([$t /] b) \gamma$ to Env.[eval $t /] (eval s) \gamma$. The STLC is of course a basic calculus, but implementing its syntax, operational and denotational semantics from scratch still involves a considerable amount of effort, mostly spent defining and reasoning about substitution. Leveraging our framework, this standard but relatively robust formalisation is possible in fewer than 100 lines of Agda code. We can also continue in the standard way with proofs like progress, determinacy, normalisation, etc. – if a syntactic property like associativity of substitution is required, it is likely to be present in the Theory module of the library already.`

5.2 Partial differentiation

Another example of a second-order calculus is the axiomatisation of partial differentiation laid out by Plotkin [2020]. The syntax consists of the first-order theory of commutative rings with a second-order partial-differentiation operator $\text{PDiff}(x.e\langle x \rangle, d)$, interpreted as the partial derivative of the expression $e\langle x \rangle$ with respect to x , evaluated at d (which has no free occurrences of x). This is usually denoted $\left. \frac{\partial e\langle x \rangle}{\partial x} \right|_{x=d}$. To differentiate $e\langle x \rangle$ without evaluation, one renames x to a dummy variable w , differentiates $e\langle w \rangle$ with respect to w , then evaluates the result at $w = x$ – thus, the notation $\frac{\partial}{\partial x} e\langle x \rangle$ is taken as abbreviating $\left. \frac{\partial e\langle w \rangle}{\partial w} \right|_{w=x}$.

The signature of rings augmented with the partial-differentiation operator can be readily expressed as an unsorted syntax description. We manually implement some derived operators in Agda, such as ∂_0 and ∂_1 (with symbols `d0` and `d1`), respectively the partial derivatives w.r.t. the first and second variable. Note the use of weakening below, available from the `Theory` module.

```

term
  zero : * | 0
  one  : * | 1
  inv  : * → * | ⊖_ r50
  add  : * * → * | ⊕_ l20
  mult : * * → * | ⊗_ l40
  pdiff : ** * → * | ∂_

  ∂0_ : PD  $\mathfrak{X}$  * (* ·  $\Gamma$ ) → PD  $\mathfrak{X}$  * (* ·  $\Gamma$ )
  ∂0 e = ∂ (Theory.wk  $\mathfrak{X}$  e) | x0
  ∂1_ : PD  $\mathfrak{X}$  * (* · * ·  $\Gamma$ ) → PD  $\mathfrak{X}$  * (* · * ·  $\Gamma$ )
  ∂1 e = ∂ (Theory.wk  $\mathfrak{X}$  e) | x1

```

The equational theory may also be included in the syntax description, listed both as explicit equations of the form ‘(name) metavar \triangleright vars \vdash $\text{expr}_1 = \text{expr}_2$ ’, and algebraic properties of operators. Plotkin’s use of function variables matches up with parametrised metavariables, so the axioms of the paper can be directly translated to our formalism. A few examples are given below.

```

theory
  ‘zero’ unit of ‘add’
  ‘mult’ distributes over ‘add’
  (∂⊕) a : *      ▷ x : *      ⊢ d0 (add (x,a)) = one
  (∂C) f : (*,*) * ▷ x : *   y : * ⊢ d1 (d0 (f[x,y])) = d0 (d1 (f[x,y]))

```

The generated Agda modules include the intrinsically-typed syntax of semirings with partial differentiation (as an inductive sorted family `PD`), and the generic equational reasoning framework of Section 3.7 instantiated with the data type `_▷_⊢_≈_A_` generated from the axiom descriptions. It employs some syntactic sugar for the purposes of readability and ease of use. Instead of defining a named inductive family of named metavariables for every axiom, we build up a context of metavariables in-place with the notation $\llbracket \Pi_1 \vdash \tau_1 \rrbracket \llbracket \Pi_2 \vdash \tau_2 \rrbracket \cdots \llbracket \Pi_n \vdash \tau_n \rrbracket$, and refer to the metavariables using (alphabetic) de Bruijn indices $\mathbf{a}, \mathbf{b}, \mathbf{c}$. The environment for an n -ary metavariable \mathbf{m} is specified as $\mathbf{m} \langle t_0 \triangleleft \dots \triangleleft t_{n-1} \rangle$, or just \mathbf{m} if $n = 0$. Some examples are shown below:

```

data _▷_⊢_≈_A_ : ( $\mathfrak{X}$  : MCtx)( $\Gamma$  : Ctx){ $\alpha$  : * $\Gamma$ } → PD  $\mathfrak{X}$   $\alpha$   $\Gamma$  → PD  $\mathfrak{X}$   $\alpha$   $\Gamma$  → Set where
  0U⊕L : [*]          ▷ 0          ⊢ 0 ⊕ a          ≈A a
  0D⊕L : [*] [*] [*] ▷ 0          ⊢ a ⊗ (b ⊕ c)      ≈A (a ⊗ b) ⊕ (a ⊗ c)
  ∂⊕   : [*]          ▷ [*]        ⊢ ∂0 (x0 ⊕ a)    ≈A 1
  ∂⊗   : [*]          ▷ [*]        ⊢ ∂0 (a ⊗ x0)    ≈A a
  ∂C   : [* · * ⊢ *]   ▷ [* · *]   ⊢ ∂1 (∂0 a ⟨ x0 ◀ x1 ⟩) ≈A ∂0 (∂1 a ⟨ x0 ◀ x1 ⟩)
  ∂Ch2 : [* · * ⊢ *] [* ⊢ *] [* ⊢ *] ▷ [*] ⊢
    ∂0 a ⟨ b ⟨ x0 ⟩ ◀ c ⟨ x0 ⟩ ⟩ ≈A (∂ a ⟨ x0 ◀ c ⟨ x1 ⟩ ⟩ | b ⟨ x0 ⟩) ⊗ (∂0 b ⟨ x0 ⟩) ⊕
    (∂ a ⟨ b ⟨ x1 ⟩ ◀ x0 ⟩ | c ⟨ x0 ⟩) ⊗ (∂0 c ⟨ x0 ⟩)

```

The first two constructors correspond to the left unit and distributivity laws, with the right versions automatically derived via the commutativity of \oplus and \otimes (omitted). The four constructors involving differentiation correspond to the axioms of sums and products ($\frac{\partial x \oplus a}{\partial x} = 1$, $\frac{\partial ax}{\partial x} = a$), exchange of derivative operators ($\frac{\partial}{\partial y} \frac{\partial}{\partial x} f(x, y) = \frac{\partial}{\partial x} \frac{\partial}{\partial y} f(x, y)$), and the binary version of the chain rule. Note that these axioms relate open terms with one or two object variables, which must be listed in the context in addition to the metavariables.

We may now feed the axiom relation to the signature-generic equational logic and establish theorems via equational reasoning – be they properties of rings, derivable partial-differentiation laws, or explicit calculations of derivatives. The user can once again make use of some syntactic sugar to streamline the proofs. For a simple example, consider the corollary $\frac{\partial 0}{\partial x} = 0$, derivable from the left annihilation law of 0 and \otimes , and the product differentiation axiom.

$$\begin{aligned} \partial 0 : \boxed{} \triangleright \boxed{} * \boxed{} \vdash \partial_0 0 &\approx 0 \\ \partial 0 = \text{begin } \partial_0 0 &\approx \langle \text{cong} [\text{ax } 0X\otimes^L \text{ with } \langle \langle x_0 \rangle \rangle] \text{ in } \partial_0 \circ^a \rangle_s \\ \partial_0 (0 \otimes x_0) &\approx \langle \text{ax } \partial \otimes \text{ with } \langle \langle 0 \rangle \rangle \rangle \\ 0 &\blacksquare \end{aligned}$$

The $\text{ax } a \text{ with } \langle \langle t_1 \triangleleft \dots \triangleleft t_n \rangle \rangle$ notation associates an axiom a with an instantiation of metavariables in the metavariable context of the axiom, given as a list of terms. Applications of an equation in a subexpression of t is done with the $\text{cong} [e] \text{ in } t \langle \mathfrak{m} \rangle$ combinator, where \mathfrak{m} is a new metavariable (denoted $\circ^{\mathfrak{m}}$ to make its role as a ‘hole’ clear) added to the context to indicate the location in which the equation e is applied. Thus, in the first step, we apply the left annihilation axiom $0 = 0 \otimes x$ instantiated at $x = x_0$ to the subexpression of the ∂_0 operator. As a more involved example, we derive the unary chain rule from the binary axiom (instantiated with $f(x, y) \triangleq f(x)$ and $h(x) = 0$), the $\partial 0$ corollary above, and the unit and annihilation laws for 0 :

$$\begin{aligned} \partial \text{Ch}_1 : \boxed{} * \boxed{} * \boxed{} \vdash \partial_0 \mathfrak{a} \langle \mathfrak{b} \langle x_0 \rangle \rangle &\approx (\partial \mathfrak{a} \langle x_0 \rangle \mid \mathfrak{b} \langle x_0 \rangle) \otimes (\partial_0 \mathfrak{b} \langle x_0 \rangle) \\ \partial \text{Ch}_1 = \text{begin } \partial_0 \mathfrak{a} \langle \mathfrak{b} \langle x_0 \rangle \rangle & \\ \approx \langle \text{ax } \partial \text{Ch}_2 \text{ with } \langle \langle \mathfrak{a} \langle x_0 \rangle \triangleleft \mathfrak{b} \langle x_0 \rangle \triangleleft 0 \rangle \rangle & \\ (\partial \mathfrak{a} \langle x_0 \rangle \mid \mathfrak{b} \langle x_0 \rangle) \otimes (\partial_0 \mathfrak{b} \langle x_0 \rangle) \oplus (\partial \mathfrak{a} \langle \mathfrak{b} \langle x_1 \rangle \rangle \mid 0) \otimes \partial_0 0 & \\ \approx \langle \text{cong} [\text{thm } \partial 0] \text{ in } [\dots] \oplus (\partial \mathfrak{a} \langle \mathfrak{b} \langle x_1 \rangle \rangle \mid 0) \otimes \circ^c & \\ (\partial \mathfrak{a} \langle x_0 \rangle \mid \mathfrak{b} \langle x_0 \rangle) \otimes (\partial_0 \mathfrak{b} \langle x_0 \rangle) \oplus (\partial \mathfrak{a} \langle \mathfrak{b} \langle x_1 \rangle \rangle \mid 0) \otimes 0 & \\ \approx \langle \text{cong} [\text{thm } 0X\otimes^R \text{ with } \langle \langle \partial \mathfrak{a} \langle \mathfrak{b} \langle x_1 \rangle \rangle \mid 0 \rangle \rangle] \text{ in } [\dots] \oplus \circ^c & \\ (\partial \mathfrak{a} \langle x_0 \rangle \mid \mathfrak{b} \langle x_0 \rangle) \otimes (\partial_0 \mathfrak{b} \langle x_0 \rangle) \oplus 0 & \\ \approx \langle \text{ax } 0U\oplus^R \text{ with } \langle \langle (\partial \mathfrak{a} \langle x_0 \rangle \mid \mathfrak{b} \langle x_0 \rangle) \otimes \partial_0 \mathfrak{b} \langle x_0 \rangle \rangle \rangle & \\ (\partial \mathfrak{a} \langle x_0 \rangle \mid \mathfrak{b} \langle x_0 \rangle) \otimes (\partial_0 \mathfrak{b} \langle x_0 \rangle) &\blacksquare \end{aligned}$$

We wrote $[\dots]$ above for the sake of brevity – the context of the congruence needs to be written out explicitly for Agda to be able to evaluate the metasubstitution that instantiates the distinguished hole metavariable. Note also the use of thm , which uses an established (non-axiomatic) equality as a proof step. The precise and sufficiently general definition of metasubstitution ensures that we always have access to the right metavariables and object variables where we need them, making the construction of equational proofs quite intuitive.

6 CONCLUSION

We presented a language-formalisation framework that allows users to produce Agda implementations of second-order languages at the press of a button. The generated term language is explicitly represented as an inductive, intrinsically-encoded data type, and the formalised metatheory can be used as and where required: substitution for operational semantics, compositional interpretations

for denotational semantics, metasubstitution for equational reasoning, etc. All the formalisms featured in the library naturally derive from the mathematical theory of abstract syntax, without the need for ad-hoc definitions or lemmas. The convenient code-generation script allows for rapid prototyping and experimentation, and it is easy to manually extend a formalised signature or incorporate an existing intrinsic term syntax into the framework.

6.1 Related work

The question of formalising and reasoning about abstract syntax was motivated by the development of proof assistants and the realisation that the [Barendregt \[1984\]](#) variable convention – “*rename variables as needed to avoid clashes*” – is difficult to translate into a formal setting. This has led to a host of approaches that address the encoding of variable binding in proof assistants and functional languages, such as: higher-order abstract syntax [[Pfenning and Elliot 1988](#); [Chlipala 2008](#)]; locally nameless representation [[Bird and Paterson 1999](#); [McBride and McKinna 2004](#); [Weirich et al. 2011](#); [Charguéraud 2012](#)]; intrinsically-typed encoding [[Benton et al. 2012](#); [Allais et al. 2021](#); [Érdi 2018](#)]; and others [[Shinwell et al. 2003](#); [Urban and Kaliszky 2011](#); [Copello et al. 2017](#)]. Similarly active is the mathematical study of abstract syntax and variable binding: developments include presheaf models [[Fiore et al. 1999](#); [Hofmann 1999](#)]; nominal sets [[Gabbay and Pitts 1999](#)]; monadic approaches [[Bellegarde and Hook 1994](#); [Altenkirch and Reus 1999](#)]; and others [[Pigozzi and Salibra 1995](#); [Sun 1999](#); [Blanchette et al. 2019](#); [Chen and Roşu 2020](#)].

Benchmarks. The POPLMARK challenge [[Aydemir et al. 2005](#); [Abel et al. 2019](#)] sets out a collection of criteria according to which metatheory-formalisation efforts can be compared. Several submissions use Coq as the target language, in some cases involving code generation from a second-order signature [[Aydemir et al. 2008](#); [Vouillon 2011](#); [Lee et al. 2012](#); [Polonowski 2013](#); [Keuchel et al. 2016](#); [Stark et al. 2019](#)]. Though impressive, the approaches rarely adopt intrinsically-typed encodings of variables and terms, usually opting for numeric de Bruijn indices with all their complicated and error-prone arithmetic. We believe that the nameless, intrinsic representation is hard to surpass in dependently-typed proof assistants thanks to its static guarantees on the typing and scoping of terms. Its drawbacks (boilerplate and types “getting in the way”) are significant and form one of the motivations of our line of research, but they are ultimately not unreasonable: a rigorous pen-and-paper proof of the type-preservation of substitution would involve the same difficulties we encounter in defining it in Agda. Our approach also incorporates generic traversals, and – for the first time, as far as we are aware – equational logic with the aid of parametrised metavariables, all naturally derived from the syntax.

Type- and scope-safe syntax. Our work is closely aligned with that of [AACMM \[2021\]](#): how to simplify the work of a language researcher by automating the boring metatheory. In their discussion of the presheaf model (loc. cit. Section 9.3), the authors suggest that freeing oneself from the formalities of the mathematics enables further progress in the development of the metatheory. We found very much the opposite: without the systematic view of the problem provided by the categorical model, one misses out on powerful principles that simplify proofs and untangle the conceptual labyrinth that is formal abstract syntax. A case in point is our reformulation of AACMM’s core [Semantics](#) record as an instance of initial algebra semantics into the internal hom, which are very useful abstractions that the authors seemingly overlooked despite working from similar foundations. We also inherit the modularity of the categorical model by cleanly separating the mathematical groundwork, abstract metatheory, second-order signatures, and term representation; whereas AACMM’s formalisation is very closely tied to their [Desc](#) data type which users are required to adopt to make use of the library. We do not have an analogue of their proof framework for simulations and fusions; but we suspect that equality-based properties (in addition to the

fusion lemmas, which we already incorporate) can be captured through initiality, without reliance on another complex set of proof techniques.

Presheaf approach. The backdrop of our Agda formalisation is a comprehensive mathematical reformulation of the presheaf model of variable binding [Fiore et al. 1999] with parametrised metavariables [Fiore 2008]. Namely, we shift attention to an indexed model (in the category of families over the set of contexts, equivalently functors from the discrete category $|\mathbb{F}|$ to **Set**) equipped with a pair of canonical adjoint monadic and comonadic modalities ($\diamond \dashv \square$, induced by the inclusion of $|\mathbb{F}|$ in \mathbb{F}) and their respective algebras and coalgebras, which are equivalent to presheaves. Although in this paper we have chosen not to emphasise the abstract categorical development in favour of a discussion geared towards programming-language researchers, we stress that the alternative viewpoint is not a matter of taste but crucial to the practical formal development that we have put forward. For instance, a direct translation of the presheaf model leads to a formalisation that inevitably requires quotients (colimits and coends) and cannot be used for computation (such as pretty printing, because of the lack of canonical representatives).

Our work required a variety of new considerations. Definitions and proofs had to be recast in the setting of (i) skew-closed structure and (ii) initiality; along the way, new notions and techniques had to be developed to bypass quotienting. Concerning (i), moving onto families with skew-monoidal structure is not enough to avoid the need for quotienting (e.g. Borthelle et al. [2020] achieve the freeness proof in a skew-monoidal setting only with the aid of a general lemma of Fiore and Saville [2017, Theorem 4.8] that relies on the presentation of initial algebras as colimits of ω -chains, rather than as inductive data types); while, concerning (ii), the fact that initial algebras in presheaves can be lifted from initial algebras in families had to be given mathematical grounding. As an upshot, using families (indexed types), instead of presheaves, leads to a lightweight practical formalisation that suits the intrinsically-typed setting well.

6.2 Future work

We recognise that the formal systems studied in modern type theory go far beyond second-order ones with algebraic types; indeed, linear, dual-context, polymorphic, dependent, polarised, etc. calculi abound. The presheaf approach has been extended to several of these [Tanaka 2000; Fiore 2006; Hyland and Tasson 2020; Fiore and Hamana 2013; Fiore 2008] and we are of course interested in adapting and/or extending our framework to these and combinations thereof.

As hinted at above, our background work involved a categorical reformulation of the highly abstract presheaf model to the formalisation-friendly realm of sorted families and \square -coalgebras; indeed, much of the Agda code has been directly read off categorical definitions and commutative diagrams. Work on the formalisation, particularly in relation to metasubstitution, is ongoing. We also anticipate interesting applications in the study of parametrised signatures and signature translations; for example, we can encode the second-order equational theory of first-order logic and modularly extend it with the relation and function symbols of any first-order signature. Further experiments with more complex languages and proofs (such as the ones given in the POPMARK challenge) will also inform and motivate the future development of our library.

ACKNOWLEDGMENTS

We thank the anonymous paper and artefact reviewers for their valuable suggestions and support. We are also grateful to Nathanael Arkor and Neel Krishnaswami for helpful discussions about technical and practical aspects of abstract syntax and language formalisation throughout the research.

REFERENCES

- M. Abadi, L. Cardelli, P.-L. Curien, and J.-J. Levy. 1989. Explicit Substitutions. In *Proceedings of the 17th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 1990)*. ACM Press, 31–46. <https://doi.org/10.1145/96709.96712>
- Andreas Abel. 2010. MiniAgda: Integrating Sized and Dependent Types. arXiv:1012.4896 [cs.PL] <https://arxiv.org/abs/1012.4896>
- Andreas Abel, Guillaume Allais, Aliya Hameer, Brigitte Pientka, Alberto Momigliano, Steven Schäfer, and Kathrin Stark. 2019. POPLMARK reloaded: Mechanizing proofs by logical relations. *Journal of Functional Programming* 29 (2019), e19. <https://doi.org/10.1017/S0956796819000170>
- Peter Aczel. 1978. A General Church–Rosser Theorem. (1978). <http://www.ens-lyon.fr/LIP/REWRITING/MISC/AGeneralChurch-RosserTheorem.pdf> Unpublished note.
- Guillaume Allais, Robert Atkey, James Chapman, Conor McBride, and James McKinna. 2021. A type- and scope-safe universe of syntaxes with binding: their semantics and proofs. *Journal of Functional Programming* 31 (2021), e22. <https://doi.org/10.1017/S0956796820000076>
- Guillaume Allais, James Chapman, Conor McBride, and James McKinna. 2017. Type-and-Scope Safe Programs and Their Proofs. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs (CPP 2017)*. ACM Press, 195–207. <https://doi.org/10.1145/3018610.3018613>
- Thorsten Altenkirch, Neil Ghani, Peter Hancock, Conor McBride, and Peter Morris. 2015. Indexed containers. *Journal of Functional Programming* 25 (2015), e5. <https://doi.org/10.1017/S095679681500009X>
- Thorsten Altenkirch and Bernhard Reus. 1999. Monadic Presentations of Lambda Terms Using Generalized Inductive Types. In *Proceedings of the 13th International Workshop on Computer Science Logic (CSL 1999) (Lecture Notes in Computer Science (LNCS), Vol. 1683)*. Springer, 453–468. https://doi.org/10.1007/3-540-48168-0_32
- Nathanael Arkor and Marcelo Fiore. 2020. Algebraic Models of Simple Type Theories: A Polynomial Approach. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2020) (LICS '20)*. ACM Press, 88–101. <https://doi.org/10.1145/3373718.3394771>
- Brian Aydemir, Aaron Bohannon, Matthew Fairbairn, J. Nathan Foster, Benjamin C. Pierce, Peter Sewell, Dimitrios Vytiniotis, Geoffrey Washburn, Stephanie Weirich, and Steve Zdancewic. 2005. Mechanized Metatheory for the Masses: The POPLMARK Challenge. In *Proceedings of the 18th International Conference on Theorem Proving in Higher Order Logics (TPHOLs 2005)*. Springer, 50–65. https://doi.org/10.1007/11541868_4
- Brian Aydemir, Arthur Charguéraud, Benjamin C. Pierce, Randy Pollack, and Stephanie Weirich. 2008. Engineering Formal Metatheory. *ACM SIGPLAN Notices* 43, 1 (2008), 3–15. <https://doi.org/10.1145/1328897.1328443>
- Henk P. Barendregt. 1984. *The Lambda Calculus: Its Syntax and Semantics*. Elsevier. <https://doi.org/10.1016/c2009-0-14341-6>
- Françoise Bellegarde and James Hook. 1994. Substitution: A Formal Methods Case Study Using Monads and Transformations. *Science of Computer Programming* 23, 2-3 (1994), 287–311. [https://doi.org/10.1016/0167-6423\(94\)00022-0](https://doi.org/10.1016/0167-6423(94)00022-0)
- Nick Benton, Chung-Kil Hur, Andrew J Kennedy, and Conor McBride. 2012. Strongly typed term representations in Coq. *Journal of Automated Reasoning* 49, 2 (2012), 141–159. <https://doi.org/10.1007/s10817-011-9219-0>
- Richard S. Bird and Ross Paterson. 1999. De Bruijn Notation as a Nested Datatype. *Journal of Functional Programming* 9, 1 (1999), 77–91. <https://doi.org/10.1017/S0956796899003366>
- Jasmin Christian Blanchette, Lorenzo Gheri, Andrei Popescu, and Dmitriy Traytel. 2019. Bindings as Bounded Natural Functors. *Proceedings of the ACM on Programming Languages* 3, POPL (2019). <https://doi.org/10.1145/3290335>
- Peio Borthelle, Tom Hirschowitz, and Ambroise Lafont. 2020. A Cellular Howe Theorem. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2020)*. ACM Press, 273–286. <https://doi.org/10.1145/3373718.3394738>
- Arthur Charguéraud. 2012. The locally nameless representation. *Journal of Automated Reasoning* 49, 3 (2012), 363–408.
- Xiaohong Chen and Grigore Roşu. 2020. A General Approach to Define Binders Using Matching Logic. *Proceedings of the ACM on Programming Languages* 4, ICFP (2020). <https://doi.org/10.1145/3408970>
- Adam Chlipala. 2008. Parametric Higher-Order Abstract Syntax for Mechanized Semantics. In *Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming (ICFP '08)*. Association for Computing Machinery, 143–156. <https://doi.org/10.1145/1411204.1411226>
- Ernesto Copello, Nora Szasz, and Álvaro Tasistro. 2017. Formal metatheory of the Lambda calculus using Stoughton’s substitution. *Theoretical Computer Science* 685 (2017), 65 – 82. <https://doi.org/10.1016/j.tcs.2016.08.025>
- Brian Day. 1970. On closed categories of functors. In *Reports of the Midwest Category Seminar IV*. Springer, 1–38.
- Gergő Erdi. 2018. Generic description of well-scoped, well-typed syntaxes. arXiv:1804.00119 [cs.PL] <https://arxiv.org/abs/1804.00119>
- Marcelo Fiore. 2006. On the structure of substitution. <https://www.cl.cam.ac.uk/~mpf23/talks/MFPS2006.pdf> Talk at MFPS 2006.

- Marcelo Fiore. 2008. Second-Order and Dependently-Sorted Abstract Syntax. In *Proceedings of the 23rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2008)*. IEEE Computer Society, 57–68. <https://doi.org/10.1109/LICS.2008.38>
- Marcelo Fiore. 2012. Discrete Generalised Polynomial Functors. In *39th International Colloquium on Automata, Languages and Programming (ICALP 2012) (Lecture Notes in Computer Science, Vol. 7392)*. Springer, 214–226.
- Marcelo Fiore and Makoto Hamana. 2013. Multiversal Polymorphic Algebraic Theories: Syntax, Semantics, Translations, and Equational Logic. In *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2013)*. IEEE Computer Society, 520–529. <https://doi.org/10.1109/LICS.2013.59>
- Marcelo Fiore and Chung-Kil Hur. 2010. Second-Order Equational Logic. In *Proceedings of the 24th International Workshop on Computer Science Logic (CSL 2010) (Lecture Notes in Computer Science (LNCS), Vol. 6247)*, Anuj Dawar and Helmut Veith (Eds.). Springer, 320–335. https://doi.org/10.1007/978-3-642-15205-4_26
- Marcelo Fiore and Ola Mahmoud. 2010. Second-Order Algebraic Theories. In *Proceedings of the 35th International Symposium on Mathematical Foundations of Computer Science (MFCS 2010) (Lecture Notes in Computer Science (LNCS), Vol. 6281)*, Petr Hliněný and Antonín Kučera (Eds.). Springer, 368–380. https://doi.org/10.1007/978-3-642-15155-2_33
- Marcelo Fiore, Gordon Plotkin, and Daniele Turi. 1999. Abstract Syntax and Variable Binding. In *Proceedings of the 14th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 1999)*. IEEE Computer Society, 193–202. <https://doi.org/10.1109/LICS.1999.782615>
- Marcelo Fiore and Philip Saville. 2017. List Objects with Algebraic Structure. In *Proceedings of the 2nd International Conference on Formal Structures for Computation and Deduction (FSCD 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 84)*, Dale Miller (Ed.). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 16:1–16:18. <https://doi.org/10.4230/LIPIcs.FSCD.2017.16>
- Murdoch J. Gabbay and Andrew M. Pitts. 1999. A New Approach to Abstract Syntax Involving Binders. In *Proceedings of the 14th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 1999)*. IEEE Computer Society, 214–224. <https://doi.org/10.1109/LICS.1999.782617>
- Joseph A. Goguen, James W. Thatcher, and E. G. Wagner. 1976. An Initial Algebra Approach to the Specification, Correctness and Implementation of Abstract Data Types. *IBM Research Report 6487* (1976).
- Makoto Hamana. 2004. Free Σ -Monoids: A Higher-Order Syntax with Metavariables. In *Proceedings of the Second Asian Symposium on Programming Languages and Systems (APLAS 2004)*, Wei-Ngan Chin (Ed.). Lecture Notes in Computer Science (LNCS), Vol. 3302. Springer, 348–363. https://doi.org/10.1007/978-3-540-30477-7_23
- Martin Hofmann. 1999. Semantical Analysis of Higher-Order Abstract Syntax. *Proceedings of the 14th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 1999)*, 204–213. <https://doi.org/10.1109/lics.1999.782616>
- Jason Hu and Jacques Carette. 2021. Formalizing category theory in Agda. In *Proceedings of the 10th ACM SIGPLAN International Conference on Certified Programs and Proofs*. 327–342.
- Martin Hyland and Christine Tasson. 2020. The linear-non-linear substitution 2-monad. arXiv:2005.09559 [math.CT] <https://arxiv.org/abs/2005.09559>
- Jonas Kaiser, Steven Schäfer, and Kathrin Stark. 2018. Binder Aware Recursion over Well-Scoped de Bruijn Syntax. In *Proceedings of the 7th ACM SIGPLAN Conference on Certified Programs and Proofs (CPP 2018)*. ACM Press, 293–306. <https://doi.org/10.1145/3167098>
- Steven Keuchel, Stephanie Weirich, and Tom Schrijvers. 2016. Needle & Knot: Binder Boilerplate Tied Up. In *Programming Languages and Systems*, Peter Thiemann (Ed.). Springer Berlin Heidelberg, 419–445.
- Anders Kock. 1971. Closed categories generated by commutative monads. *Journal of the Australian Mathematical Society* 12, 4 (1971), 405–424. <https://doi.org/10.1017/S1446788700010272>
- Wen Kokke, Jeremy G. Siek, and Philip Wadler. 2020. Programming language foundations in Agda. *Science of Computer Programming* 194 (2020), 102440. <https://doi.org/10.1016/j.scico.2020.102440>
- Joachim Lambek. 1980. From λ -calculus to cartesian closed categories. *To HB Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism* (1980), 375–402.
- Gyesik Lee, Bruno C. D. S. Oliveira, Sungkeun Cho, and Kwangkeun Yi. 2012. GMeta: A Generic Formal Metatheory Framework for First-Order Representations. In *Programming Languages and Systems*, Helmut Seidl (Ed.). Springer Berlin Heidelberg, 436–455.
- The Coq Development Team. 2004. *The Coq proof assistant reference manual*. <https://coq.inria.fr/distrib/current/files/CoqRefMan.pdf>
- Conor McBride. 2005. Type-preserving renaming and substitution. (2005). <http://strictlypositive.org/ren-sub.pdf> Unpublished note.
- Conor McBride and James McKinna. 2004. Functional pearl: I am not a number – I am a free variable. In *Proceedings of the 2004 ACM SIGPLAN Workshop on Haskell*. 1–9.
- Ulf Norell. 2009. Dependently Typed Programming in Agda. *Proceedings of the 4th International Workshop on Types in Language Design and Implementation (TLDI 2009)* (2009), 1–2. <https://doi.org/10.1145/1481861.1481862>

- Frank Pfenning and Conal Elliot. 1988. Higher-order abstract syntax. *Proceedings of the 9th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 1988)*, 199–208. <https://doi.org/10.1145/960116.54010>
- Matthew Pickering, Gergő Érdi, Simon Peyton Jones, and Richard A. Eisenberg. 2016. Pattern Synonyms. In *Proceedings of the 9th International Symposium on Haskell (Haskell 2016)*. Association for Computing Machinery, 80–91. <https://doi.org/10.1145/2976002.2976013>
- Don Pigozzi and Antonino Salibra. 1995. The abstract variable-binding calculus. *Studia Logica* 55, 1 (1995), 129–179. <https://doi.org/10.1007/bf01053036>
- Andrew M. Pitts. 2019. *Initial algebra for a strictly positive endofunctor constructed using sized types and quotient types*. <https://www.cl.cam.ac.uk/~amp12/agda/initial-T-algebras/>
- Gordon D. Plotkin. 2020. A Complete Equational Axiomatisation of Partial Differentiation. *Electronic Notes in Theoretical Computer Science* 352 (2020), 211 – 232. <https://doi.org/j.entcs.2020.09.011>
- Emmanuel Polonowski. 2013. Automatically Generated Infrastructure for De Bruijn Syntaxes. In *Interactive Theorem Proving*, Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie (Eds.). Springer Berlin Heidelberg, 402–417.
- Mark R. Shinwell, Andrew M. Pitts, and Murdoch J. Gabbay. 2003. FreshML: Programming with Binders Made Simple. In *Proceedings of the 8th ACM SIGPLAN International Conference on Functional Programming (ICFP 2003)*. ACM Press, 263–274. <https://doi.org/10.1145/944705.944729>
- Kathrin Stark, Steven Schäfer, and Jonas Kaiser. 2019. Autosubst 2: reasoning with multi-sorted de Bruijn terms and vector substitutions. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs*. 166–180.
- Ross Street. 2013. Skew-closed categories. *Journal of Pure and Applied Algebra* 217, 6 (2013), 973 – 988. <https://doi.org/10.1016/j.jpaa.2012.09.020>
- Yong Sun. 1999. An algebraic generalization of Frege structures – binding algebras. *Theoretical Computer Science* 211, 1 (1999), 189 – 232. [https://doi.org/10.1016/S0304-3975\(97\)00170-9](https://doi.org/10.1016/S0304-3975(97)00170-9)
- Kornél Szlachányi. 2012. Skew-monoidal categories and bialgebroids. *Advances in Mathematics* 231, 3 (2012), 1694 – 1730. <https://doi.org/10.1016/j.aim.2012.06.027>
- Miki Tanaka. 2000. Abstract Syntax and Variable Binding for Linear Binders. In *Proceedings of the 25th International Symposium on Mathematical Foundations of Computer Science (MFCS 2000) (Lecture Notes in Computer Science (LNCS), Vol. 1893)*, Mogens Nielsen and Branislav Rovan (Eds.). Springer, 670–679. https://doi.org/10.1007/3-540-44612-5_62
- Christian Urban and Cezary Kaliszyk. 2011. General Bindings and Alpha-Equivalence in Nominal Isabelle. In *Programming Languages and Systems*, Gilles Barthe (Ed.). Springer Berlin Heidelberg, 480–500.
- Vladimir Voevodsky, Benedikt Ahrens, Daniel Grayson, et al. 2014. *UniMath – a computer-checked library of univalent mathematics*. <https://github.com/UniMath/UniMath>
- Jérôme Vouillon. 2011. A Solution to the PoplMark Challenge Based on de Bruijn Indices. *Journal of Automated Reasoning* 49 (2011), 327–362.
- Stephanie Weirich, Brent A. Yorgey, and Tim Sheard. 2011. Binders Unbound. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming (ICFP '11)*. Association for Computing Machinery, 333–345. <https://doi.org/10.1145/2034773.2034818>