

GATlab: Modeling and Programming with Generalized Algebraic Theories

Owen Lynch^{a,1} Kris Brown^{a,2} James Fairbanks^{b,3} Evan Patterson^{a,4}

^a *Topos Institute, Berkeley, CA*

^b *University of Florida, Computer & Information Science & Engineering, Gainesville, FL*

Abstract

Categories and categorical structures are increasingly recognized as useful abstractions for modeling in science and engineering. To uniformly implement category-theoretic mathematical models in software, we introduce GATlab, a domain-specific language for algebraic specification embedded in a technical programming language. GATlab is based on generalized algebraic theories (GATs), a logical system extending algebraic theories with dependent types so as to encompass category theory. Using GATlab, the programmer can specify generalized algebraic theories and their models, including both free models, based on symbolic expressions, and computational models, defined by arbitrary code in the host language. Moreover, the programmer can define maps between theories and use them to declaratively migrate models of one theory to models of another. In short, GATlab aims to provide a unified environment for both computer algebra and software interface design with generalized algebraic theories. In this paper, we describe the design, implementation, and applications of GATlab.

Keywords: generalized algebraic theories, GATs, algebraic specification, programming languages, applied category theory

1 Introduction

Category theory has long been recognized as a useful organizing principle for programming. The far-ranging dictionary between category theory and type theory, originating with the equivalence between cartesian closed categories and lambda calculus with product types [21], has led language designers to import all manner of categorical concepts into programming languages. In this usage, category theory acts a *mathematical model of the programming language*, often giving its denotational semantics. But this is not the only role that category theory can play in programming. Applied category theorists have shown how category theory can formalize compositional structure latent in science and engineering, from relational databases to stochastic and quantum processes to differential equations and dynamical systems [37,15,13,24]. Category theory now becomes a *mathematical model of the subject-matter domain*, which (like any model) will be most useful if it can be implemented in software. We thus need a software system that can readily express and compute with categorical structures.

This objective might be served by designing the host language according to category-theoretic ideas but is in principle independent of it. Indeed, there is presently a tension between the two roles, since the

¹ Email: owen@topos.institute

² Email: kris@topos.institute

³ Email: fairbanksj@ufl.edu

⁴ Email: evan@epatters.org

dependently typed languages that most natively express categorical concepts tend to have few facilities for scientific and technical computing, whereas the languages most commonly used by scientists and engineers tend to have minimalistic type systems. The problem can be approached from either direction. In this work, we show how to augment an existing technical computing language with an advanced type system to enable uniform computation over categorical structures.

Rather than trying to retrofit a full dependent type theory onto a loosely typed programming language, we work with generalized algebraic theories, an extension of algebraic theories sufficient to axiomatize categorical structures. A *generalized algebraic theory* (GAT) [8,9], more suggestively called a *dependently typed algebraic theory* [33], is similar to a (typed or sorted) algebraic theory except that its types may depend on terms. The ur-example of a GAT is the theory of categories, in which the type of a morphism depends on a pair of objects, the domain and codomain. Though a relatively simple type theory, GATs suffice to axiomatize essentially any theory consisting of a category equipped with extra structure that is defined algebraically (equationally). Examples of categorical structures axiomatizable as GATs include categories; presheaves and copresheaves; monoidal categories, be they strict or weak, braided or symmetric; categories equipped with chosen finite products or finite limits; and 2-categories, bicategories, and double categories.

1.1 Contributions

In this paper, we describe the design and implementation of GATlab,⁵ a programming framework for generalized algebraic theories built as an embedded domain-specific language in the Julia programming language. GATs have long served as a foundation for Catlab,⁶ a framework for applied category theory focusing on scientific and engineering applications. We have recently rewritten our system for GATs from scratch, while significantly extending its features. Here we describe it in print for the first time.

Major contributions of GATlab are to

- provide an *algebraic specification language* based on a minimal dependent type theory in a technical programming language;
- support a *standard library* of over 90 reusable theories,⁷ ranging from classical algebraic structures like groups and rings to categorical structures like monoidal categories and presheaves;
- enable uniform computing with *models* of GATs, including both *free models*, based on symbolic expressions, and computational models, defined by arbitrary code in the host language;
- declaratively and algebraically *migrate* models of one theory to another, via morphisms of theories.

In short, GATlab aims to be a structural and symbolic foundation for categorically structured modeling tools in science and engineering.

1.2 Related work

Both the mathematics of GATs and the use of such mathematics in software have a long history. In this section we recount some of that history and explain how GATlab relates to it.

The oldest relevant thread of research is the theory of universal algebra. The phrase “universal algebra” dates back at least to Whitehead [39], but the subject was not put on a formal footing until Birkhoff’s 1946 treatment [3]. Later, Lawvere showed in his doctoral thesis how many of the constructions done within universal algebra arise naturally from basic concepts of category theory [23].

Early implementations of universal algebra were in the OBJ and Clear languages [16,5]. The module system of the Clear language influenced both the design of the Standard ML module system [28,29] and later iterations of the OBJ family of languages [17]. Modern incarnations of these universal-algebraic concepts can be found in systems such as Maude [12]. One of our original objectives for GATlab was to build an ML-inspired module system in Julia, in which implementations of a theory were reifiable objects just as

⁵ GATlab is open source software available on GitHub: <https://github.com/AlgebraicJulia/GATlab.jl>

⁶ Catlab is also available on GitHub: <https://github.com/AlgebraicJulia/Catlab.jl>

⁷ The library of GATs, models, and theory maps can be browsed at: <https://forest.localcharts.org/gatlib.xml>

implementations of a signature are in ML. Only later did we realize that ML modules were themselves inspired by universal algebra, so in that sense GATlab draws on a well-established tradition.

GATlab goes beyond that tradition, however, in using generalized algebraic theories, which extend universal algebra to include dependent types, as needed for the theory of categories. The challenge of extending universal algebra with dependent types has been addressed in many ways. One approach is via logical frameworks (LF), first proposed in [20]. Logical frameworks are an algebraic approach to manipulating the syntax of dependent type theory, by embedding it within another dependent type theory. GATlab plays a similar “meta” role as LF, but does not support type constructors which quantify over variables. On the other hand, GATlab offers other features that LF does not, such as the ability to add arbitrary equations to a theory. That was previously thought to be a bad idea, as it breaks the decidability of type checking and equality checking. However, as we will show, interesting and useful things can be done with GATs even in the absence of decidable type checking.

GATlab’s emphasis on theory morphisms as a means of translating between theories and combining theories via colimits is also found in the mathematical knowledge management (MKM) system MMT [34], which supports more general classes of theories but is less focused on computational semantics. Likewise, this approach to MKM is found in MathScheme [6] and related work on theory presentation combinators [7].

Finally, GATlab is a computational implementation of categories and other categorical structures. Related projects in this vein include Rydeheard and Burstalls’s computational category theory [36] and the CAP (Categories, Algorithms, and Programming) computer algebra system [19,2].

2 Background: Generalized algebraic theories and their models

We review the main ideas behind generalized algebraic theories (GATs), focusing on their syntax, in both mathematical and GATlab programming notation, and their standard set-theoretic semantics. GATs were introduced by John Cartmell in his PhD thesis [8] and later in print [9], to which we refer for the many details omitted here. As further references, Pitts [33, §6] and Taylor [38, Chapter VIII] present type theories based closely on Cartmell’s GATs.

2.1 Syntax of GATs

GATs are most easily understood through examples. Figure 1 shows the ur-example, the theory of categories, first in mathematical notion in the style of Pitts [33] and then as a theory in the GATlab programming environment. Most notably, the listings are the same up to cosmetic differences in syntax. Julia’s macro system makes this possible by capturing the code in the body of the `@theory` macro call instead of evaluating it as standard Julia code. Additionally, GATlab allows text identifiers to be aliased by mathematical symbols (via Unicode operators) and axioms to be tagged with names. Note that we break with convention in type theory and put the context on the right so that the name of the constructor is the first thing that one sees on a given line.

In general, a GAT is comprised of the following elements.

Contexts A *context* is a finite (ordered) list of named variables tagged with types, where the type at a given position may depend on variables at previous positions. Examples include the empty context $[]$; the context $[a::0b, b::0b, c::0b]$, or its shorthand $[(a,b,c)::0b]$, having no type dependency; and the context $[a::0b, b::0b, f::\text{Hom}(a,b)]$, where the type of the variable f depends on the previous variables a and b . An *invalid* context is $[f::\text{Hom}(a,b), a::0b, b::0b]$ since the variables a and b are used as type parameters before they are introduced.

Judgments A GAT is defined by a sequence of *judgments*. Every judgment has a context, which we write to the right of the turnstile ($-$). There are three kinds of judgments, summarized in Table 1:

Type constructor New types, possibly dependent on variables in the context, are introduced by *type constructors*.

Term constructor Typed terms are introduced by *term constructors*. Both the term and its type may depend on variables in the context.

Term equality Axioms asserting equations between terms in context are specified by *term equalities*.

The theory of categories, for instance, has two type constructors, for objects and morphisms; two term

$\text{Ob} : \text{type} \dashv$ $\text{Hom}(a, b) : \text{type} \dashv a, b : \text{Ob}$ $f \cdot g : \text{Hom}(a, c) \dashv a, b, c : \text{Ob},$ $f : \text{Hom}(a, b), g : \text{Hom}(b, c)$ $\text{id}(a) : \text{Hom}(a, a) \dashv a : \text{Ob}$ $(f \cdot g) \cdot h = f \cdot (g \cdot h) \dashv a, b, c, d : \text{Ob}, f : \text{Hom}(a, b)$ $g : \text{Hom}(b, c), h : \text{Hom}(c, d)$ $\text{id}(a) \cdot f = f \dashv a, b : \text{Ob}, f : \text{Hom}(a, b)$ $f \cdot \text{id}(b) = f \dashv a, b : \text{Ob}, f : \text{Hom}(a, a)$ <p style="text-align: center;">(a) In mathematical notation</p>	<pre> @theory ThCategory begin Ob :: TYPE \dashv [] @op (\dashv) := Hom Hom(dom :: Ob, codom :: Ob) :: TYPE @op (\cdot) := compose compose(f :: (a \rightarrow b), g :: (b \rightarrow c)) :: (a \rightarrow c) \dashv [(a, b, c) :: Ob] id(a :: Ob) :: (a \rightarrow a) assoc := ((f \cdot g) \cdot h) == (f \cdot (g \cdot h)) \dashv [(a, b, c, d) :: Ob, f :: (a \rightarrow b), g :: (b \rightarrow c), h :: (c \rightarrow d)] idl := id(a) \cdot f == f \dashv [(a, b) :: Ob, f :: (a \rightarrow b)] idr := f \cdot id(b) == f \dashv [(a, b) :: Ob, f :: (a \rightarrow b)] end </pre> <p style="text-align: center;">(b) In GATlab notation</p>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 1. The ur-example of a GAT: the theory of categories

Judgment kind	Math notation	GATlab notation	GATlab examples
Type constructor	$X \text{ type } [\Gamma]$	$\{X\} :: \text{TYPE} \dashv \{\Gamma\}$	$\text{Ob} :: \text{TYPE} \dashv []$ $\text{Hom}(d, c) :: \text{TYPE} \dashv [(d, c) :: \text{Ob}]$
Term constructor	$x : X [\Gamma]$	$\{x\} :: \{X\} \dashv \{\Gamma\}$	$\text{id}(a) :: \text{Hom}(a, a) \dashv [a :: \text{Ob}]$ $(a \otimes b) :: \text{Ob} \dashv [(a, b) :: \text{Ob}]$
Term equality	$x = y : X [\Gamma]$	$\{x\} == \{y\} \dashv \{\Gamma\}$	$(a \otimes b) \otimes c == a \otimes (b \otimes c) \dashv [(a, b, c) :: \text{Ob}]$

 Table 1
 Judgments in a generalized algebraic theory

constructors, for composition and identities; and three term equalities, for the axioms of associativity and left and right unitality.

As a technical aside, we note that, according to Cartmell, GATs admit a fourth kind of judgment, equalities between types. We follow Taylor [38] in excluding type equalities from GATlab. This restriction simplifies the system, even if it does not solve the problem of type checking (since dependent types can still be nontrivially equal due to equations between the terms they depend on). From a category-theoretic perspective, prohibiting explicit type equalities is no loss since equations between types are better handled by isomorphisms anyway.

Specifically, this restriction makes sort-checking and sort-inference quite easy. Sort-inference determines the type constructor used for the type of a term. For instance, it classifies a term as an object or a morphism, without necessarily figuring out the domain or codomain of the morphism. This is a practical and fast check that we perform for most operations in GATlab, and it catches surface-level bugs.

2.2 Semantics of GATs

Under the standard set-theoretic semantics, a *model* of a GAT interprets each type constructor with a set, or family of sets, or family of families of sets, etc., according to the level of type dependency; and similarly each term constructor with a function, or family of functions, etc., according to the level of dependency; in such a way that all of the term equalities are interpreted as true equations. A more precise account is

given by Cartmell [9, §11], but the idea becomes most transparent in examples. A model M of the theory of categories (Figure 1) consists of a set of objects, $M(\text{Ob})$; a family of hom-sets, $(M(\text{Hom})(a, b))_{a, b \in M(\text{Ob})}$; the identities, a family of functions $(M(\text{id})(a) : 1 \rightarrow \text{Hom}(a, a))_{a \in M(\text{Ob})}$; and finally the composition operations, a family of functions

$$M(\cdot)(a, b, c) : M(\text{Hom})(a, b) \times M(\text{Hom})(b, c) \rightarrow M(\text{Hom})(a, c), \quad a, b, c \in M(\text{Ob}),$$

indexed by triples of objects, such that the category axioms are satisfied. Moving up a dimension, a model M of a theory of 2-categories would include a family of families of sets $M(\text{Hom}_2)(f, g)$, indexed by $f, g \in M(\text{Hom})(a, b)$, indexed in turn by $a, b \in M(\text{Ob})$.

GATlab offers the programmer two different ways to work with models of GATs, as we will see. Generalizing their set-theoretic semantics, models of GATs can be taken in any *contextual category* [9, §14]. We do not elaborate on the categorical semantics of GATs as it plays no role in GATlab.

2.3 Alternatives to GATs

Generalized algebraic theories belong to a family of essentially equivalent logics that extend the logic of algebraic theories to encompass the theory of categories and others like it. Besides GATs, the best known of these logics are essentially algebraic theories [1, §3.D], finite limit sketches, and finite limit theories. Cartmell sketched an argument to the effect that, under their set-theoretic semantics, generalized and essentially algebraic theories have equivalent categories of models [9, §6]. Meanwhile, essentially algebraic theories and finite limit sketches are directly intended to be syntactic presentations of finite limit theories, an invariant notion of theory.

Any of these logics could, in principle, play the role that GATs do in GATlab. Our choice of GATs is practical: despite their complicated meta-theory, GATs lead to by far the most readable and intuitive presentations of theories of categorical structures, often closely resembling their textbook forms.

3 Models of GATs in GATlab

In the context of software engineering, GATs play the role of an interface or a formal specification. Models of a GAT, then, are the data structures that implement the interface and satisfy the formal specification.

3.1 Models of algebraic theories

Given a generalized algebraic theory declared with `@theory`, we want to be able to “implement models of that theory in Julia.” What does this mean?

Languages from the ML family have notions of a *signature* and of a *module* that implements the signature. A module comprises a choice of concrete type for each type in the signature and a function for each term constructor. Models are then referenced explicitly when used. For instance, suppose we had a module `IntPlusMonoid` implementing the signature of monoids for integers using addition. Then we would use it as follows:

```
IntPlusMonoid.mappend 1 2 == 3
```

We could have another module `IntTimesMonoid` implementing the signature of monoids for integers using multiplication, and then we’d have

```
IntTimesMonoid.mappend 1 2 == 2
```

The ML module system contrasts with languages like Haskell or Rust, in which type classes or traits have at most a single implementation for a given type, and the implementation is not explicitly referenced. So one must choose between privileging the additive monoidal structure on integers and having `mappend 1 2 == 3` or privileging the multiplicative monoidal structure on integers and having `mappend 1 2 == 2`.

In Julia, multiple dispatch can be used to simulate both module-style and trait-style models of theories.⁸

⁸ The original implementation of GATs in Catlab supported only trait-style models, meaning that a given type or

The trait-style models are most straightforward. For instance, we can implement the free monoid based on string concatenation as follows.

```
@instance ThMonoid{String} begin
  e() = ""
  (x · y) = x * y
end

"a" · "b" == "ab" # true
```

The `@instance` macro is a lightweight wrapper around ordinary Julia function definitions using multiple dispatch that checks that all of the methods in the theory have been given implementations. Note, the type parameter syntax (`{String}`) associates concrete Julia types with the GAT’s ordered list of type constructors.

In order to implement module-style models using Julia’s multiple dispatch, we employ a trick. We associate model implementations with structs (records in Julia), and then use those structs to guide dispatch. In the following code, the `Tuple{Int}` parameter to `Model` tells the system that when `IntPlusMonoid` is used as a model of a theory, the first type constructor in the theory should be interpreted as `Int`. We will discuss why we do this later on; in brief this “exposes” the implementing types to the Julia type system, which makes it easier to create derived models.

```
struct IntPlusMonoid <: Model{Tuple{Int}} end

@instance ThMonoid{Int} [model::IntPlusMonoid] begin
  e() = 0
  (x · y) = x + y
end

e[IntPlusMonoid{Int}]() == 0 # true
```

In the above code, the `@instance` block is macro-expanded to ordinary Julia code resembling

```
ThMonoid.e{::WithModel{IntPlusMonoid}} = 0
ThMonoid.:(·){::WithModel{IntPlusMonoid}, x::Int, y::Int} = x + y
```

The code `e[IntPlusMonoid{Int}]()` invoked above is syntactic sugar for `e(WithModel{IntPlusMonoid{Int}})`.

As a technical detail, the abstract type `Model` takes a tuple type parameter which must match the type parameterization in the `@instance` declaration, and the type `WithModel` is a simple wrapper struct defined as:

```
struct WithModel{M}
  model::M
end
```

We control multiple dispatch by passing the model type as the first argument; however, we use the wrapped `WithModel` type to avoid ambiguities that may arise from cases where the instance Julia types are not `String` or `Int` as above but rather models themselves, as in the case of the `SliceC` model, introduced in [3.2](#).

We also provide a convenience macro `@withmodel` for evaluating expressions in a module-style model. For example, the invocation

```
@withmodel IntPlusMonoid{Int} (e, ·) begin
  2 · e() == 2
end
```

macro-expands to the let block

collection of types could implement a theory at most one way. A important motivation for creating GATlab was to support both styles of models.

```

let
  withmod = WithModel(IntPlusMonoid())
  e = () -> ThMonoid.e(withmod)
  · = (x, y) -> ThMonoid.:(·)(withmod, x, y)
  2 · e() == 2
end # true

```

We are not limited to empty structs when implementing models of theories. For instance, we can implement the monoid $\mathbb{Z}/n\mathbb{Z}$ of integers modulo n as follows

```

struct ModularPlusMonoid <: Model{Tuple{Int}}
  n::Int
end

@instance ThMonoid{Int} [model::ModularPlusMonoid] begin
  e() = 0
  (x · y) = mod(x + y, model.n)
end

```

and use it like so

```

@withmodel ModularPlusMonoid(7) (e, ·) begin
  3 · 4 == e()
end # true

```

We call such models *parametric*. They are similar in spirit to “functors” in ML languages.⁹

3.2 Incorporating dependent types into models

So far we have considered models in GATlab only of algebraic theories. How can we extend these classical notions of model to support dependent types, in light of the fact that Julia doesn’t quite have dependent types?¹⁰ There are two approaches to modeling dependent types within a non-dependent language that we have considered. It is illustrative to consider these two approaches in the example of categories.

The first approach is to embrace the translation from generalized algebraic theories to essentially algebraic theories (see Section 2.3), so that the type constructor $\text{Hom}: \text{Ob} \times \text{Ob} \rightarrow \text{Type}$ becomes instead a span $\text{Ob} \leftarrow \text{Hom} \rightarrow \text{Ob}$ of non-dependent types. Then operations like composition become *partial* functions. We might call this the “fibered” approach to dependent types.

In this style, we could implement the category of finite sets, or rather its skeleton, as follows. First, we define data structures for the objects and morphisms of this category.

```

struct FinSet
  n::Int
end

struct FinFunction
  values::Vector{Int}
  dom::FinSet
  codom::FinSet
end

```

Then we make them into a model of the theory of categories.

⁹ Incidentally, the name “parametric modules” was proposed for functors in ML, but Burstall objected “on the grounds that we don’t call ordinary functions ‘parametric values’” [29]. Since in GATlab they are precisely “models with parameters,” we think the name is appropriate.

¹⁰ The question of whether and to what extent Julia has dependent types has engendered much debate [35]. At the very least, we can say that Julia does not have dependent type sufficient for this purpose, as types cannot depend on, say, arbitrary structs. So our implementation of dependent types does not use Julia’s dependent types.

```

@instance ThCategory{FinSet, FinFunction} begin
  dom(f::FinFunction) = f.dom
  codom(f::FinFunction) = f.codom

  id(x::FinSet) = FinFunction(1:(x.n), x, x)

  function compose(f::FinFunction, g::FinFunction)
    f.codom == g.dom || @fail "domain and codomain do not match"
    FinFunction(g.values[f.values], f.dom, g.codom)
  end
end

```

The `@fail` macro is a special macro that can be used inside `@instance` to throw an error with the given message, but also provide contextual information about theory, model, and method in question.

Storing the domain and codomain along with the data of the morphism was the approach taken in the original Catlab implementation of GATs. Indeed, the domain or codomain often cannot be deduced from the data of the morphism. In the example above, representing the data of a function as an array of values determines its range but not its codomain, so the codomain must be stored separately. While intuitive, this approach has the disadvantage that the data structure for a morphism must carry around extra data. So a *diagram* of morphisms might end up storing many redundant copies of the objects.

In GATlab, we continue to support such models, but we also introduce a new approach that cuts down on the redundant storage of objects. The idea is to translate the type constructor $\text{Hom}: \text{Ob} \times \text{Ob} \rightarrow \text{Type}$ into a predicate $\text{IsHom}: \text{Hom} \times \text{Ob} \times \text{Ob} \rightarrow \mathbf{Bool}$. Now Hom is the union of all of the hom-sets, which need not be a disjoint union. If the previous approach was “fibered,” we might call this one “indexed.” In this style, an implementation of the category of finite sets might look like

```

struct FinSetC <: Model{Tuple{Int, Vector{Int}}} end

@instance ThCategory{Int, Vector{Int}} [model::FinSetC] begin
  Ob(x::Int) = x >= 0 ? x : @fail "expected nonnegative integer"

  function Hom(f::Vector{Int}, n::Int, m::Int)
    length(f) == n ||
      @fail "length of morphism does not match domain: $(length(f)) != $n"
    for i in 1:n
      if !(f[i] in 1:m)
        @fail "index not in codomain: $i"
      end
    end
    return f
  end

  id(x::Int) = collect(1:x)
  compose(f::Vector{Int}, g::Vector{Int}) = g[f]
end

```

Since we can dispatch on the model, we now have the freedom to just use the types `Int` and `Vector{Int}` for our objects and morphisms, instead of having to wrap them in custom structs. Note also that in GATlab, the `Hom` function is not merely a predicate, but rather a coercion function that throws an informative error when the arguments do not type check, and returns a valid morphism when they do.

We can recover the fibered approach from the indexed one by letting $\text{Hom}' = \text{IsHom}^{-1}(\mathbf{true})$, but there is no systematic way to “drop” redundant information, so in a sense the second approach is more primitive than the first. This is related to the fact that the Grothendieck construction only gives an equivalence of categories, not an isomorphism.¹¹

¹¹ One might also think about the indexed approach as a form of the “type erasure” that happens during the

We conclude this section by demonstrating how to create a model for slice categories that is parameterized by the base category. First, we declare a model struct.

```

struct SliceOb{ObT, HomT}
  ob::ObT
  hom::HomT
end

struct SliceC{ObT, HomT, C<:Model{Tuple{ObT, HomT}}}}
  <: Model{Tuple{SliceOb{ObT, HomT}, HomT}}
  cat::C
  over::ObT
end

```

We interpret this as saying that `SliceC(cat, over)` is a model with associated types `SliceOb{ObT, HomT}` and `HomT`. We don't track the fact that it is a model of the theory of categories within the Julia type system because an abstract type in Julia can have only one supertype, but we want to allow for multiple inheritance. We will use the `over` field as the object within the `cat` field that we are slicing over. We then implement the methods of the theory of categories as follows.

```

@instance ThCategory{SliceOb{ObT, HomT}, HomT} [model::SliceC{ObT, HomT, C}]
  where {ObT, HomT, C} begin
    function Ob(x::SliceOb{ObT, HomT})
      try
        Ob[model.cat](x.ob)
      catch e
        @fail ("ob is not valid", e)
      end
      try
        Hom[model.cat](x.hom, x.ob, model.over)
      catch e
        @fail ("hom is not valid", e)
      end
      x
    end

    function Hom(f::HomT, x::SliceOb{ObT, HomT}, y::SliceOb{ObT, HomT})
      try
        Hom[model.cat](f, x.ob, y.ob)
      catch e
        @fail ("morphism is not valid in base category", e)
      end
      compose[model.cat](f, y.hom) == x.hom ||
        @fail "commutativity of triangle does not hold"
      f
    end

    id(x::SliceOb{ObT, HomT}) = id[model.cat](x.ob)

    compose(f::HomT, g::HomT) = compose[model.cat](f, g)
  end

```

compilation of dependently typed languages, where values not relevant to the runtime semantics are erased. That is, the runtime representation of values of type `Hom(a,b)` need not carry around pointers to `a` and `b`. Ideally, this would be an optimization performed by a clever compiler rather than done manually, but we must work within the constraints of the Julia programming language.

The most interesting parts are in the dynamic checks for validity of objects and morphisms. In the function `Ob`, we check that `x.ob` is a valid object inside `model.cat`, and then we check that `x.hom` is a valid morphism from `x.ob` to `model.over`. In the function `Hom`, we check that `f` is a valid morphism from `x.ob` to `y.ob` and that the relevant triangle commutes. The last two methods, `id` and `compose`, are just inherited from the base category `model.cat`.

3.3 Free models of GATs

In addition to models based on arbitrary Julia code, GATlab can construct *free* or *symbolic* models of GATs,¹² via the `@symbolic_model` macro. For instance, the following macro call creates a free model of theory of categories:

```
@symbolic_model FreeCategory{ObExpr, HomExpr} ThCategory begin
  compose(f::Hom, g::Hom) = associate_unit(new(f,g; strict=true), id)
end
```

This code creates a new module called `FreeCategory` in which structs called `Ob` and `Hom` are defined and methods `ThCategory.id` and `ThCategory.compose` implemented for them. It also provides some special logic via an externally-defined function `associate_unit` to enforce that composition is strictly associative and unital, by normalizing to n -ary composition operations whenever `compose` is called. We can use the resulting data structures as follows.

```
A, B = Ob(FreeCategory, :A), Ob(FreeCategory, :B)
f, g = Hom(:f, A, B), Hom(:g, B, A)
```

Essentially, we can think of this model as “the free model of the theory of categories on the set of symbols.” Specifically, this is the free category on the graph whose vertices are the set of symbols and whose edges between any two vertices are also the set of symbols.

One use of free models in Catlab is to schedule and evaluate string diagrams. That is, we can transform a combinatorial presentation of a string diagram into an expression in the theory of symmetric monoidal categories. If we then have a model of the theory of symmetric monoidal categories, we can then interpret this expression in that model.

Within the `@symbolic_model` block, we can specialize the implementation of the operations of the theory to normalize the terms. For instance, in the invocation of `@symbolic_model` above, we normalize with respect to the associative law and the unit law, so that `compose(f, compose(g, h)) == compose(compose(f, g), h)`. Normalization isn’t always possible for more complex theories, but it is convenient when it is possible.

Currently, we have implemented the following normalization strategies:

- `associate`, normalizing a binary operation for associativity
- `associate_unit`, normalizing a binary operation for associativity and unitality
- `associate_unit_inv`, normalizing a binary operation for associativity, unitality, and inverses
- `distribute_unary`, distribute a unary operation over a binary operation
- `involute`, normalize an involutive unary operation
- `normalize_zero`, collapse to zero when the expression contains a zero

However, these are simply the normalization strategies that we have found useful; our approach does not preclude other strategies from being implemented.

4 Morphisms of GATs

While it is always possible to write arbitrary functions that transform a model of one theory into a model of another, the code that one must write is often opaque and not easily verified. In many situations, a solution to this problem is provided by the fact that different GATs are related to each other by *morphisms*,

¹²Strictly speaking, these “free models” are not necessarily models of the theory, because the theory’s equational axioms may hold only up to rewrites of the expressions.

also called *interpretations*. Using morphisms of GAT, we implement a declarative and verifiable syntax for migrating models from one theory to another.

4.1 Syntax of GAT morphisms

Following Cartmell [9, §12], two GATs A and B can be related to each other via an *interpretation*, which assigns to each type (respectively, term) constructor of A a type (respectively, term) in B . This mapping must satisfy some conditions in order for it to be a well-typed and validity-preserving translation of the vocabulary of A into the vocabulary of B .

For example, consider translating statements in the language of monoids into the language of natural number arithmetic:

```

@theory ThMonoid <: ThSet begin
  (x·y) :: default ⊢ [x, y]
  (x·y)·z == (x·(y·z)) ⊢ [x, y, z]
  e() :: default
  e()·x == x ⊢ [x]
  x·e() == x ⊢ [x]
end

@theory ThArith <: ThNat begin
  (x::ℕ + y::ℕ)::ℕ ⊢ [(x, y)::ℕ]
  n + S(m) == S(n+m) ⊢ [(n, m)::ℕ]

  (x::ℕ * y::ℕ)::ℕ
  n*S(m) == ((n*m) + n) ⊢ [(n,m)::ℕ]
end
    
```

Note that `ThSet` is a special theory with one type, `default`, which is the default interpretation of a type if the user does not provide one. We can interpret the theory of monoids either additively or multiplicatively in the theory of arithmetic:

```

@map PlusM(ThMonoid, ThArith) begin
  default => ℕ
  x·y ⊢ [x, y] => x+y
  e() => Z()
end

@map TimesM(ThMonoid, ThArith) begin
  default => ℕ
  x·y ⊢ [x, y] => x*y
  e() => S(Z())
end
    
```

In general, to declare how an interpretation I acts on a type (resp. term) constructor f , on the left side of the arrow we write f 's *generic* type (resp. term) in context, $f(x_1, \dots, x_n) \vdash [arg_1 :: T_1, \dots, arg_n :: T_n]$. On the right side we define $I(f)$ by writing a type (resp. term) implicitly in the context $[x_1 :: I(T_1), \dots, x_n :: I(T_n)]$. This mapping defined on type and term constructors induces a function on all type and term expressions using the operations of A into expressions of B ; we say that this interpretation function “pushes forward types/terms/contexts along I ”. For example, the expression `PlusM(e()·x·e() ⊢ [x])` evaluates to `Z()+x+Z() ⊢ [x::ℕ]`. We call an interpretation $A \rightarrow B$ a *theory map* or *morphism* in the category GAT.

The above examples can be contrasted with user-declared morphisms that are ill-typed or give invalid interpretations:

```

@map Bad1(ThMonoid, ThArith) begin
  default => ℕ
  x·y ⊢ [x,y] => x+y
  # unit axioms invalidated
  e() => S(Z())
end

@map Bad2(ThCategory, ThCategory) begin
  # Ill-typed arguments for Hom
  Ob ⊢ [] => Hom()
  a → b ⊢ [(a,b)::Ob] => a → b
  # Ill-typed argument for id
  id(a) ⊢ [a::Ob] => id(a)
end
    
```

An interpretation is “validity-preserving” when, for each axiom Eq of A , there exists a proof of $I(Eq)$ using the axioms of B . When I is `Bad1` and Eq is the unitality axiom $e() \cdot x = x$, this would mean that we need to find a proof that $1 + x = x$ in `ThArith`, which is impossible. The other possible issue with a user-specified theory morphism is that it is ill-typed, such as `Bad2`, which attempts to map $Ob \mapsto Hom()$, which fails both because `Hom` requires arguments and because `ThCategory` requires that `id` take an `Ob` argument, not a `Hom`.

4.2 Data structures for subclasses of GAT morphisms

Morphisms of GATs are highly expressive, which entails a certain amount of complication. It can be useful to consider GAT morphisms that are more restrictive, yet require less data to specify. These more restricted GAT morphisms also make it easier to compute, for instance, pushouts. GATlab defines no less than four data structures that implement GAT morphisms of varying expressivity:

```

struct IdTheoryMap <: AbsTheoryMap
  gat::GAT
end

struct SimpleTheoryMap <: AbsTheoryMap
  dom::GAT; codom::GAT
  typemap::Dict{Ident, Ident}
  termmap::Dict{Ident, Ident}
end

struct InclTheoryMap <: AbsTheoryMap
  dom::GAT; codom::GAT
end

struct TheoryMap <: AbsTheoryMap
  dom::GAT; codom::GAT
  typemap::Dict{Ident, TypeInCtx}
  termmap::Dict{Ident, TermInCtx}
end

```

Each of these kinds of morphism can be cast to the subsequent kind, thus all can be interpreted as general theory maps, which send type and term constructors respectively to types and terms in a context of the codomain theory. `IdTheoryMap` represents the identity morphism on a GAT, and `InclTheoryMap` represents inclusions of theories. The third kind, `SimpleTheoryMap`, allows for associating type and term constructors of the domain with structurally-identical type and term constructors of the codomain. The fourth and final kind, `TheoryMap`, is a general morphism of theories, as described above.

For example, consider a map `F` which interprets preorders in the language of categories:

```

@theory ThPreorder <: ThSet begin
  Leq(dom, codom)::TYPE ⊢ [dom, codom]
  @op (≤) := Leq
  refl(p)::(p ≤ p) ⊢ [p]
  tran(f::p ≤ q, g::q ≤ r)::(p ≤ r) ⊢ [p,q,r]
  irrev := f == g ⊢ [p, q, (f, g)::(p ≤ q)]
end

@map F(ThCategory, ThPreorder) begin
  Ob => default
  Hom => Leq
  compose => tran
  id => refl
end

```

There is no overlap between the domain and codomain GATs, so this map cannot be represented by an `InclTheoryMap`, but it is a `SimpleTheoryMap`. In addition, `PlusM` and `TimesM` above could both have been expressed using `SimpleTheoryMap` syntax. However, a general `TheoryMap` is needed to represent the following GAT morphism, which takes a monoid to its opposite:

```

@map OpMonoid(ThMonoid, ThMonoid) begin
  default => default; e() => e(); x·y ⊢ [x, y] => y·x
end

```

When `OpMonoid` is applied to $x \cdot (e() \cdot y) \vdash [x, y]$ one obtains $(y \cdot e()) \cdot x \vdash [x, y]$.

4.3 Computation with GAT morphisms

The basic functionalities these data structures must provide in order to play the role of GAT morphisms are composition, pushing forward of terms, pulling back of models, and checking the validity of the morphism. All of these operations are trivial for morphisms of kind `IdTheoryMap` or `InclTheoryMap`. Composition for `SimpleTheoryMap` is simply composition of the underlying functions from identifiers to identifiers, whereas composition for `TheoryMap` is implemented as a composition of substitutions, akin to a Kleisli composition for a term monad.

Pushing forward types and terms along `SimpleTheoryMap` or `TheoryMap` instances requires an implementation of pattern matching and syntactic substitution. The pushforward operation is essential to pulling back models along theory maps: models associate a computational semantics with each expression of a theory B , so by precomposing with the interpretation function of expressions of A to B , we see that GAT morphisms

$A \rightarrow B$ induce a function from models of B to models of A . For example, if `ThArith` were given a model `JuliaArith`, using the native addition and multiplication of Julia, then the model `IntPlusMonoid` defined manually in the previous section could be obtained declaratively by calling GATlab’s `migrate_model` on `PlusM` and `JuliaArith`. As another example, the following theory map allows us to automatically generate models of opposite categories. Once we have checked that the theory map is valid, we have confidence that applying it to correct models will always produce correct models.

```
@map OpCat(ThCategory, ThCategory) begin
  Ob => Ob
  (dom → codom) ⊢ [(dom, codom)::Ob] => (codom → dom)
  id(a) ⊢ [a::Ob] => id(a)
  compose(f, g) ⊢ [(a,b,c)::Ob, f::(a → b), g::(b → c)] => compose(g, f)
end
```

The operation of checking the validity of a theory map is trivial for `IdTheoryMap`. For `InclTheoryMap`, one must check that the identifiers found in `dom` are a strict subset of those found in `codom`. For `SimpleTheoryMap`, one must check that for each type and term constructor of the domain that the assigned type or term in the codomain has the same structure. E.g., to check the map `F`’s assignment `compose => tran`, one looks at the domain’s generic term in context, namely

```
compose(f,g) ⊢ [(a,b,c)::Ob, f::(a → b), g::(b → c)]
```

and checks whether applying the theory map to the term in context is identical (up to renaming variables) to the generic term in context of the codomain theory, which in this case is

```
tran(f, g) ⊢ [(a,b,c)::default, p::(a≤b),q::(b≤c)]
```

For a general theory map, determining validity-preservation requires solving many word problems, and deciding equality is a nontrivial semi-decidable computation in general. For example, one must check that `OpMonoid` preserves the associativity of monoid multiplication.

Furthermore, morphisms can be used to compositionally construct new theories via colimits. This can be done without explicit morphisms due to special support for a multiple inheritance syntax:

```
@theory ThModule begin
  using ThRing: default as Scalar
  using ThAdditiveAbelianGroup: default as Vector
   $\alpha \cdot v$  :: Vector ⊢ [ $\alpha$ ::Scalar, v::Vector]
  # ... axioms
end
```

`ThModule` is an extension of the coproduct of `ThRing` and `ThAdditiveAbelianGroup`.

Alternatively, the same theory can be explicitly constructed as a pushout of maps to suitably renamed versions of `ThRing` and `ThAdditiveAbelianGroup` from the theory of their common overlap (in this case, the empty theory), which can then be extended to include scalar multiplication and axioms. When the span to be pushed out uses only `SimpleTheoryMap` or simpler kinds of morphism, the computation is a simple matter of renaming and taking the union of the component GATs. Pushouts of general theory maps are a much more complicated matter yet to be implemented. Thus, drawing distinctions between different flavors of theory map allows GATlab to distinguish feasible and infeasible computations with GATs.

5 Applications and Extensions

In creating the GATlab package as a replacement for the GAT machinery in Catlab, our first major goal was to explore new design decisions such as

- (i) models of GATs in the ML module style, in addition to the Haskell typeclass style (section 3.1)
- (ii) fibered vs. indexed point of view on dependent types (section 3.2)
- (iii) hygienic substitution via scope tags (appendix A.1)

but retain sufficient backwards-compatibility with Catlab that adoption of these new features could be done incrementally rather than all at once. Having now completed this significant engineering effort, we are beginning to take advantage of GATlab’s new functionality throughout the AlgebraicJulia ecosystem of packages for applied category theory.

One of the motivations behind GATlab was to build functionality for symbolic dynamical systems. In the past, we have built dynamical systems on top of arbitrary Julia functions [24]. However, making the various functions in dynamical systems *symbolic* unlocks a range of new capabilities, such as:

- Systems can be optimized/compiled to produce more efficient code
- Systems can be analyzed symbolically to prove properties about them without simulating them
- Systems can be transferred between different programming languages
- The equations defining a system can be pretty-printed and displayed to the user

There is a blog post written by the first author which explores these directions [25], and we have a preliminary implementation of symbolic resource sharers as an operad algebra [26]. Notably, GATlab allows us to implement these symbolic resource sharers agnostic to the exact collection of “primitive functions” usable in the implementation of a symbolic function. For instance, we could only allow polynomial functions, or only allow polynomial and trigonometric functions, and so on.

To support these systems-theoretic capabilities, another major direction for GATlab is developing richer functionality for computer algebra. We are interested in exploring how types can aid the development of computer algebra for mathematical objects going beyond classical abstract algebra. For instance, we have prototyped a theory in which the primary objects of study are “sections of vector bundles,” where both the vector bundle itself and the space over which the vector bundle sits are arguments to the type constructor. This would enable principled, coordinate-free descriptions of problems in differential geometry, and have applications to partial differential equations via use in Decapodes [32,31].

There are several ways in which we are considering adding more sophisticated rewrite functionality to GATlab for the purposes of computer algebra. We have prototyped our own implementation of e-graphs for GATlab syntax trees, based on a careful reading of the egg source code [40]. However, Julia already has an excellent implementation of e-graphs in the form of Metatheory.jl [10], with which we hope to integrate. In any case, the advantage of e-graphs is that they are applicable to almost any rewriting problem, which is a good fit for the generality of GATs. However, in specialized domains like rings and modules, algorithms that involve Groebner bases can be far more efficient than naive use of e-graphs. We have access to such computer algebra systems in Julia via Symbolics.jl [18] and the CAP project [2], and so another direction would be to integrate these into specific theories within GATlab.

A limitation of GATs is that although we can specify a GAT for categories, GATs only have a 1-category of models rather than a 2-category. So, from the GAT for categories, we can automatically derive the concept of functor, but not the concept of natural transformation. Lambert and Patterson have a notion of “cartesian double theory” [22], the models of which naturally form a 2-category, which could be a more natural framework for doing “computer algebra for categories” than generalized algebraic theories.

Yet another direction for future research is randomized testing. This can be done in two ways. The first is to take an arbitrary equation in a theory, the validity of which is uncertain, and then check if it is true in randomly sampled finite models of a GAT. Of course, this cannot prove that the equation is true, but it can furnish counterexamples in the style of Nitpick [4] and Alloy [30]. The second way is to take a purported model of a GAT, and then sample elements of the types in that model to check if the equations in the GAT actually hold, in the style of QuickCheck [11] and Hypothesis [27]. This would provide a cheap way to get confidence in an implementation, and improve the overall code quality of anything which used GATlab.

Acknowledgement

The authors acknowledge support from the DARPA ASKEM program under grant HR00112220038.

References

- [1] Adamek, J. and J. Rosicky, *Locally Presentable and Accessible Categories*, Cambridge University Press (1994).
<https://doi.org/10.1017/cbo9780511600579>
- [2] Barakat, M., T. Kuhmichel and F. Zickgraf, *CompilerForCAP – a category theory aware compiler, building and compiling towers of categories* (2022).
<https://algebra.mathematik.uni-siegen.de/barakat/CompilerForCAP/CompilerForCAP.pdf>
- [3] Birkhoff, G., *Universal algebra*, in: *Comptes Rendus du Premier Congrès Canadien de Mathématiques*, pages 310–326, University of Toronto Press (1946).
- [4] Blanchette, J. C. and T. Nipkow, *Nitpick: A counterexample generator for higher-order logic based on a relational model finder*, in: M. Kaufmann and L. Paulson, editors, *Interactive Theorem Proving (ITP 2010)*, volume 6172 of *LNCS*, pages 131–146, Springer (2010).
- [5] Burstall, R. M. and J. A. Goguen, *The semantics of clear, a specification language*, pages 292–332, Springer Berlin Heidelberg (1980).
https://doi.org/10.1007/3-540-10007-5_41
- [6] Carette, J., W. M. Farmer and R. O’Connor, *MathScheme: project description*, in: *Intelligent Computer Mathematics: 18th Symposium, Calculemus 2011, and 10th International Conference, MKM 2011, Bertinoro, Italy, July 18-23, 2011. Proceedings 4*, pages 287–288, Springer (2011).
https://doi.org/10.1007/978-3-642-22673-1_23
- [7] Carette, J. and R. O’Connor, *Theory presentation combinators*, in: *International Conference on Intelligent Computer Mathematics*, pages 202–215, Springer (2012).
https://doi.org/10.1007/978-3-642-31374-5_14
- [8] Cartmell, J., *Generalised algebraic theories and contextual categories*, Ph.D. thesis, Oxford University (1978).
- [9] Cartmell, J., *Generalised algebraic theories and contextual categories*, *Annals of Pure and Applied Logic* **32**, pages 209–243 (1986).
[https://doi.org/10.1016/0168-0072\(86\)90053-9](https://doi.org/10.1016/0168-0072(86)90053-9)
- [10] Cheli, A., *Metatheory.jl: Fast and elegant algebraic computation in julia with extensible equality saturation*, *Journal of Open Source Software* **6**, page 3078 (2021).
<https://doi.org/10.21105/joss.03078>
- [11] Claessen, K. and J. Hughes, *Quickcheck: a lightweight tool for random testing of haskell programs*, in: *Proceedings of the fifth ACM SIGPLAN international conference on Functional programming*, ICFP00, ACM (2000).
<https://doi.org/10.1145/351240.351266>
- [12] Clavel, M., F. Durán, S. Eker, S. Escobar, P. Lincoln, N. Martí-Oliet and C. Talcott, *Two Decades of Maude*, pages 232–254, Springer (2015).
https://doi.org/10.1007/978-3-319-23165-5_11
- [13] Coecke, B. and A. Kissinger, *Categorical Quantum Mechanics I: Causal Quantum Processes*, Oxford University Press (2018).
<https://doi.org/10.1093/oso/9780198748991.003.0012>
- [14] Flatt, M., *Bindings as sets of scopes* (2015). (fifth RacketCon).
<https://www.youtube.com/watch?v=ABWLveMNdzg>
- [15] Fong, B., *Causal theories: A categorical perspective on bayesian networks* (2013).
<https://doi.org/10.48550/ARXIV.1301.6201>
- [16] Goguen, J. A., *Abstract errors for abstract data types*, in: E. J. Neuhold, editor, *Formal Description of Programming Concepts: Proceedings of the IFIP Working Conference on Formal Description of Programming Concepts, St. Andrews, NB, Canada, August 1-5, 1977*, pages 491–526, North-Holland (1977).
- [17] Goguen, J. A., T. Winkler, J. Meseguer, K. Futatsugi and J.-P. Jouannaud, *Introducing OBJ*, pages 3–167, Springer (2000).
https://doi.org/10.1007/978-1-4757-6541-0_1
- [18] Gowda, S., Y. Ma, A. Cheli, M. Gwóźdz, V. B. Shah, A. Edelman and C. Rackauckas, *High-performance symbolic-numeric via multiple dispatch*, *ACM Commun. Comput. Algebra* **55**, pages 92–96 (2022), ISSN 1932-2240.
<https://doi.org/10.1145/3511528.3511535>
- [19] Gutsche, S., S. Posur and Ø. Skartæterhagen, *On the syntax and semantics of cap*, in: O. Hasan, M. Pfeiffer and G. D. Reis, editors, *Proceedings of the Workshop Computer Algebra in the Age of Types* (2018).
<https://www.cicm-conference.org/2018/infproc/paper21.pdf>

- [20] Harper, R., F. Honsell and G. Plotkin, *A framework for defining logics*, Journal of the ACM **40**, pages 143–184 (1993).
<https://doi.org/10.1145/138027.138060>
- [21] Lambek, J. and P. J. Scott, *Introduction to higher-order categorical logic*, Cambridge University Press (1986).
- [22] Lambert, M. and E. Patterson, *Cartesian double theories: A double-categorical framework for categorical doctrines*, Advances in Mathematics **444**, page 109630 (2024).
<https://doi.org/10.1016/j.aim.2024.109630>
- [23] Lawvere, F. W., *Functorial Semantics of Algebraic Theories and Some Algebraic Problems in the context of Functorial Semantics of Algebraic Theories*, Ph.D. thesis, Columbia University (1963).
<http://www.tac.mta.ca/tac/reprints/articles/5/tr5abs.html>
- [24] Libkind, S., A. Baas, E. Patterson and J. Fairbanks, *Operadic modeling of dynamical systems: Mathematics and computation*, Electronic Proceedings in Theoretical Computer Science **372**, pages 192–206 (2022), ISSN 2075-2180.
<https://doi.org/10.4204/eptcs.372.14>
- [25] Lynch, O., *Symbolic presentations of dynamical systems* (2023).
<https://blog.algebraicjulia.org/post/2023/05/algebraic-geometry-2/>
- [26] Lynch, O., *Symbolic compositional resource sharers* (2024).
<https://github.com/AlgebraicJulia/GATlab.jl/pull/136>
- [27] MacIver, D., Z. Hatfield-Dodds and M. Contributors, *Hypothesis: A new approach to property-based testing*, Journal of Open Source Software **4**, page 1891 (2019), ISSN 2475-9066.
<https://doi.org/10.21105/joss.01891>
- [28] MacQueen, D., *Modules for standard ML*, in: *Proceedings of the 1984 ACM Symposium on LISP and functional programming - LFP '84*, LFP '84, ACM Press (1984).
<https://doi.org/10.1145/800055.802036>
- [29] MacQueen, D., R. Harper and J. Reppy, *The history of standard ML*, Proceedings of the ACM on Programming Languages **4**, pages 1–100 (2020).
<https://doi.org/10.1145/3386336>
- [30] Milicevic, A., J. P. Near, E. Kang and D. Jackson, *Alloy*: a general-purpose higher-order relational constraint solver*, in: *Proceedings of the 37th International Conference on Software Engineering - Volume 1*, ICSE '15, pages 609–619, IEEE Press (2015), ISBN 9781479919345.
- [31] Morris, L., A. Baas, J. Arias, M. Gatlin, E. Patterson and J. P. Fairbanks, *Decapodes: A diagrammatic tool for representing, composing, and computing spatialized partial differential equations* (2024).
<https://doi.org/10.48550/arXiv.2401.17432>
- [32] Patterson, E., A. Baas, T. Hosgood and J. Fairbanks, *A diagrammatic view of differential equations in physics*, Mathematics in Engineering **5**, pages 1–59 (2022), ISSN 2640-3501.
<https://doi.org/10.3934/mine.2023036>
- [33] Pitts, A. M., *Categorical logic*, in: *Handbook of Logic in Computer Science*, volume 5, pages 39–128 (2001).
<https://doi.org/10.1093/oso/9780198537816.003.0002>
- [34] Rabe, F. and M. Kohlhase, *A scalable module system*, Information and Computation **230**, pages 1–54 (2013), ISSN 0890-5401.
<https://doi.org/10.1016/j.ic.2013.06.001>
- [35] Ramine, A., J. Nash, K. Fischer, J. Sterling, C. T. Schonwald, S. Karpinski, B. McKenna, J. Bezanson, S. Cawtawba, E. Kmett, T. Knopp, R. Valentine, J. Abrahamson, I. Norton, J. M. White and M. Bouchet-Valat, *Julia has no dependent types* (2014).
<https://github.com/JuliaLang/julia/issues/6113>
- [36] Rydeheard, D. E. and R. M. Burstall, *Computational Category Theory*, Prentice Hall (1988).
- [37] Spivak, D. I., *Functorial data migration*, Information and Computation **217**, pages 31–51 (2012), ISSN 0890-5401.
<https://doi.org/10.1016/j.ic.2012.05.001>
- [38] Taylor, P., *Practical foundations of mathematics*, Cambridge University Press (1999).
- [39] Whitehead, A., *A treatise on universal algebra, with applications.*, volume 1, Cambridge University Press (1898).
- [40] Willsey, M., C. Nandi, Y. R. Wang, O. Flatt, Z. Tatlock and P. Panckekha, *egg: Fast and extensible equality saturation*, Proc. ACM Program. Lang. **5** (2021).
<https://doi.org/10.1145/3434304>

A Implementation of GATlab

In this section, we describe certain details of GATlab’s implementation that may be useful to developers looking to create similar systems.

A.1 *Scopes*

In the original Catlab implementation of GATs, variables were identified simply by symbols, which implicitly means that variables were resolved via dynamic scoping. An objective of the GATlab implementation was to provide a lexical scope system, which properly handled issues like shadowing and overloading and helped identify errors in syntax manipulation algorithms. We originally intended to use a system based on deBruijn levels, but the reindexing necessary when moving syntax around while using deBruijn levels proved quite tedious to handle. We eventually settled on a system inspired by Racket’s “scope sets” construct, which is used to ensure hygenicity of the Racket macro system [14]. However, because GATlab does not involve macros that expand to macros, we did not need the full generality of “scope sets.”

Our scope system works in the following way. We call any piece of syntax that introduces new variable bindings a **scope**. For instance, the body of a `@theory` macro is a scope, the right hand side of a turnstile (“ \dashv ”) is a scope, and the collection of arguments to a function is a scope.

At parse time, every scope is assigned a random UUID, which we call a **scope tag**. An **identifier** consists of a scope tag, an index ranging from 1 to the length of the scope, and a name. The scope tag and the index are sufficient to identify the binding that the identifier refers to; the name is just for display purposes.

A context is then an ordered list of scopes. When a symbol is parsed within a context, we produce an identifier with a scope tag that refers to the latest scope within that context that has a binding named with that symbol. However, when later on we want to find the binding associated with an identifier, we look through the context for a scope with a matching scope tag, rather than simply taking the first binding with the right name. The scope mechanism ensures that substitution is *hygenic*: when a variable is substituted into a new expression, in a new environment, it will never accidentally refer to something it was not intended to refer to.

While implementing lexical scopes might seem a minor detail, it was not trivial to settle on this approach to scoping, and so we hope that reporting the details of this solution will help designers of typed computer algebra systems in the future. Using a combination of UUID-based scope identifiers and consecutive-integer indexing is more ergonomic than pure deBruijn indexing because reindexing happens less frequently. More importantly, it is also better at detecting programmer errors. In a computer algebra system, we are frequently moving syntax trees around, so it may happen that we forget to bring the corresponding scope along; having variables reference their scope by UUIDs is an easy way to detect this kind of error.

A.2 *Expanding macros inside macros*

This section is about a trick somewhat specific to Julia as we do not know how many other programming languages support calling `macroexpand` inside a macro. It would be an interesting exercise to make this work in a language like Racket which has strict phase separation.

When one theory in GATlab extends another theory, we need to access the contents of the parent theory *at macro expansion time*. The reason is that the `@theory` macro produces Julia modules and declares functions, so if we had to wait until runtime to get the contents of the parent theory, we would have to call `eval` at runtime to produce those modules. Exactly why this would not be ideal is a complicated story involving the details of how Julia performs compilation, but suffice to say that this is best avoided.

Generally, macros cannot access the *values* referred to by the syntax that is passed to them, only the syntax itself. However, one thing that *can* be done inside a Julia macro is to call `macroexpand`. We exploit this possibility in the following way. A call to `@theory ThMonoid begin ... end` expands into a module called `ThMonoid` as follows.

```
module ThMonoid
...
module Meta
```

```

...
macro theory()
  # (return the data of the theory of monoids)
end
...
end
...
end

```

Then, when macroexpanding an extension of the theory

```

@theory ThGroup <: ThMonoid begin
  ...
end

```

we do something like

```

macro theory(head, body)
  parentname = ...
  parenttheory = macroexpand(__module__, :($parentname.Meta.@theory))
  ...
end

```

We now have access to the data of the theory of monoids in the variable `parenttheory`, and can use it in the rest of the `@theory` macro.