

Categories for PL

Notes from CS7480 (Fall 2025)

Steven Holtzen and John M. Li

*with Shubh Agrawal, Jialu Bao, Bex
Golovanov, Mingtong Lin, Joseph Rotella,
and Michael Zhang*

MAY 26, 2026

Contents

1	<i>Why Categories?</i>	9
2	<i>Categories</i>	15
3	<i>Universal Constructions I: Products</i>	25
4	<i>Universal Constructions II: Exponents</i>	39
5	<i>From Categories to Languages</i>	49
6	<i>Yoneda I: Internal vs. external</i>	63
7	<i>Yoneda II: indexed set theory</i>	71
8	<i>Yoneda III: Representability</i>	79
9	<i>Yoneda IV: universal constructions, elements, and properties</i>	85
10	<i>Functors and natural transformations</i>	93
11	<i>Monoidal Categories</i>	103
12	<i>Limits and colimits</i>	113
13	<i>Adjunctions</i>	123
14	<i>Elementary topoi</i>	133
15	<i>Sheaf semantics</i>	139
16	<i>Monads</i>	147
17	<i>Modeling effects with syntax</i>	153
18	<i>Realizability</i>	157
19	<i>Categorical Noninterference</i>	167
20	<i>Type Refinements as Indexed Inductive Types</i>	191
21	<i>Algebras and automata</i>	231
22	<i>Bibliography</i>	257

Preface

These are notes from a course on category theory for programming languages researchers. I would say that these notes, and this course, are developed in anger: I have been forced to learn and understand a bit about category theory due to its inevitable usefulness. I would not describe myself as a category theorist, and this class is quite unlike any I've ever tried to teach before. My goal in this class is to tease out what I find inevitable about category theory, and present it in the way that makes the most sense to me. Of course, this isn't possible without John, who is the one who really understands everything here.

Many programming languages researchers and students have encountered category theory at one point or another. Probably most see it for the first time when they are learning Haskell, when they see words like “monad” and “functor.” Perhaps others who are more mathematically inclined see it when they are learning about the lambda calculus and hear about Scott, Smyth, and Plotkin's famous developments of the model theory of the untyped lambda calculus [34]. You're here because you probably saw category theory somewhere and wondered why it was needed, or perhaps were skeptical that it was really necessary.

For me, the first time I saw category theory appear in my own research that forced me to stop and understand it was Heunen et al.'s quasi-Borel spaces paper.¹ This was during my second year of grad school, and I have to admit I found it extremely intimidating. This paper was a terrifying combination of (1) seeming very important to my immediate research, and (2) being utterly and completely incomprehensible. In a nutshell, this paper introduced a denotational semantics for a higher-order probabilistic programming language with continuous probability distributions called *quasi-Borel spaces* (QBS). I was – and still am! – a probabilistic programming languages researcher, so this result seemed pivotal: I like functions! However, at the time, it wasn't even obvious to me *why this was a hard thing to do in the first place!* If you can't understand the motivation for a problem, good luck understanding its solution.

This preface is written by Steven.

¹ Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. A convenient category for higher-order probability theory. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12. IEEE, 2017

The core challenge, helpfully provided in the second paragraph of the paper, is:

“Programs in these languages may combine higher-order functions and continuous distributions, or even define a probability distribution on functions. But the standard measure theoretic formalization of probability theory does not handle higher-order functions well, as the category of measurable spaces is not cartesian closed.” [14]

This made *absolutely no sense to me*. The only citation is to a paper from 1961, called “Borel structures for function spaces” [3], which was equally (if not even more!) incomprehensible than the current paper I was trying to understand. I was out of my depth and had to move on, and I did not understand this paper for many years.

At the time the QBS paper seemed like it came like a bolt from the blue and really shook my confidence in my own ability to do research, but over the past 8 years I’ve come to realize that, in some sense, this seemingly tour-de-force alien paper was an almost inevitable combination of simple ideas. It turns out that the idea for quasi-Borel spaces was born not in probability but in a seemingly totally disparate corner of mathematics: differential geometry and topology. Last year I was having a conversation with Max New about QBS, who told me that the setup had occurred to him in grad school, but he wasn’t motivated to fully develop it due to its seemingly straightforward relationship to diffeological spaces: these results are so similar, structurally, that to a trained expert the monstrous QBS paper almost seems *too trivial!* This is our first answer to the question of *why category theory?* – it makes precise the analogies that exist between disparate areas of mathematics, allowing you to solve different problems using very similar machinery.

If this was the only time I encountered category theory I would not have become invested enough in it to attempt to teach a course on it. But, it continues to appear. My subsequent encounters with category theory were driven primarily by my PhD. student John, who is helping me teach this class and write these notes. This story begins with Lilac [22], a probabilistic separation logic that we developed during the first two years of John’s PhD. Before we get into that we need to briefly discuss separation logics – I promise we’ll get to category theory.

Separation logics are logics that describe resources. They are widely used within programming languages to model how resources like memory flow around programs. Concretely, consider the following C program:

```

1 void foo(int* x, int* y) {
2   for(int i = 0; i < 3; i++) {
3     *x += *y;
4   }

```

5 }

One might look at `foo` and assume that it is very inefficiently written: it could be optimized to simply add $3*y$ to x . But, *what if x and y point to the same location in memory* (i.e., they are aliases)? Then, this optimization is no longer valid, since y is also mutating on each iteration of the loop. Since the C compiler must be pessimistic, it cannot perform this optimization.

Separation logics give us a language for describing memory ownership so that we can tell the compiler that these pointers do not alias. If we want this optimization to be valid, then `foo` must have a *precondition* that asserts that the two arguments must not be aliases of each other. In separation logic, this precondition can be written as:

$$[x \mapsto -] * [y \mapsto -]$$

The $*$ is called the *separating conjunction*: it states that the two propositions $[x \mapsto -]$ (read “ x points to some location”) and $[y \mapsto -]$ hold of *disjoint* portions of the heap, and hence x and y cannot be aliases of one another.

However, memory is not the only kind of resource in a program: *randomness* is also a natural kind of resource. In probabilistic programs, we might like to express propositions like:

$$[x \sim \text{Bern } 1/2] * [y \sim \text{Bern } 1/2]$$

This would express that two variables x and y are *independent* Bernoulli random variables. Beyond probability, it’s clear that programs might have many other notions of resources, like network sockets, mutex locks, etc. Must we develop *from-scratch* separation logics for each of them?

This leads us to the second answer to the *why category theory?* question: it makes certain choices that seem ad-hoc *forced*. We will see how there is a categorical description of the separating conjunction that can be *calculated* using something called the “Day convolution” [11, 6]. If your interpretation of $*$ arises from this construction, then you get *for free* that it satisfies many other well-formedness requirements for separating conjunction. This shows how category theory can bring clarity to this proliferation of interactions by highlighting which decisions you’ve made are actually *new* and which are *consequences*: it forces your hand, which is very useful when you don’t know where your hand is supposed to go.

At this point, I’m convinced some familiarity with category theory is a worthwhile investment for anyone doing research in programming languages, especially those working with exotic effects like probability or those working with program logics. I also have some

We are foreshadowing here, but we will return to the Day convolution later on in the course once we’ve worked up to it.

hubris: I believe that it is not so hard to understand, especially if one is sufficiently motivated by some of the applications that we will attempt to highlight as we go. My target audience for this class is past me from 2017 seeing an important paper in their area that feels terribly out of reach because of its fluent application of categorical ideas. I hope that we can present some of these ideas clearly enough for you to see the inevitability, rather than be drowned by the unintelligibility, of these elegant ideas.

1

Why Categories?

One of the most common places one encounters category theory is in semantics. To illustrate why categories appear here and what they are typically used for, we'll consider a concrete example that every programmer has encountered: we want to establish which operations change the behavior of our programs and which are irrelevant.

Consider a tiny programming language `CALC` consisting of let-bindings, pairs, and Booleans:

$M, N ::= \text{let } x \text{ be } M \text{ in } N \mid x \mid \text{true} \mid \text{false} \mid \text{if } M \text{ then } N \text{ else } O \mid \langle M, N \rangle \mid \text{fst } M \mid \text{snd } M$

$A, B ::= A \times B \mid \text{Bool}$

(CALC)

CALC has a simple type system:

$$\frac{}{\Gamma \vdash \text{true} : \text{Bool}} \qquad \frac{}{\Gamma \vdash \text{false} : \text{Bool}}$$

$$\frac{\Gamma \vdash M : \text{Bool} \quad \Gamma \vdash N : A \quad \Gamma \vdash O : A}{\Gamma \vdash \text{if } M \text{ then } N \text{ else } O : A} \qquad \frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \langle M, N \rangle : A \times B}$$

$$\frac{x : A \in \Gamma}{\Gamma \vdash x : A} \qquad \frac{\Gamma \vdash M : A \times B}{\Gamma \vdash \text{fst } M : A} \qquad \frac{\Gamma \vdash M : A \times B}{\Gamma \vdash \text{snd } M : B}$$

$$\frac{\Gamma \vdash M : A \quad \Gamma, x : A \vdash N : B}{\Gamma \vdash \text{let } x \text{ be } M \text{ in } N : B}$$

Here's a simple transformation one might want to justify: two let-bindings can be interchanged if their definitions do not mention each other.

$$\text{let } x \text{ be } M \text{ in } \left(\text{let } y \text{ be } N \text{ in } \langle x, y \rangle \right)$$

$$\stackrel{?}{\equiv}$$

$$\text{let } y \text{ be } N \text{ in } \left(\text{let } x \text{ be } M \text{ in } \langle x, y \rangle \right)$$

Intuitively, this equation holds because `CALC` programs have no side effects. We can justify this intuition by proving that these two programs have the same meaning. For this, we need to give a semantics to `CALC` programs.

1.1 A Denotational Semantics for `CALC`

The first step in designing a semantics is to choose a **metalanguage** for stating it. The metalanguage should be a formal language that is unambiguous to describe and use, because its meaning will be taken for granted when we interpret programs. For a simple language like `CALC`, a natural choice is the language of sets and functions.

Using this metalanguage we can give a **denotational semantics** to `CALC` programs that describes the meaning of syntactic programs using our choice of metalanguage. Types will be interpreted as sets:

$$\llbracket \text{Bool} \rrbracket = \{0, 1\} \quad (1.1)$$

$$\llbracket A \times B \rrbracket = \llbracket A \rrbracket \times \llbracket B \rrbracket \quad (1.2)$$

The Boolean type denotes the two-element set $\{0, 1\}$, and the product type $A \times B$ denotes the Cartesian product of the denotations of A and B . Using this interpretation of types, we can give an interpretation to typing contexts:

$$\llbracket \bullet \rrbracket = \{\star\}, \quad \llbracket \Gamma, x : A \rrbracket = \llbracket \Gamma \rrbracket \times \llbracket A \rrbracket$$

The empty context \bullet is interpreted as the singleton set $\{\star\}$, and context extension by Cartesian product. Taken together, this interpretation defines $\llbracket \Gamma \rrbracket$ to be the set of *substitutions of shape* Γ , whose elements are functions that map variables x in Γ to elements of $\llbracket \Gamma(x) \rrbracket$.

Using this interpretation of types and contexts, we can interpret `CALC` terms $\Gamma \vdash M : A$ as functions $\llbracket \Gamma \vdash M : A \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$:

$$\begin{aligned} \llbracket \text{true} \rrbracket &= \gamma \mapsto 1 \\ \llbracket \text{false} \rrbracket &= \gamma \mapsto 0 \\ \llbracket \text{if } M \text{ then } N \text{ else } O \rrbracket &= \gamma \mapsto \begin{cases} \llbracket N \rrbracket \gamma, & \llbracket M \rrbracket \gamma = 1 \\ \llbracket O \rrbracket \gamma, & \llbracket M \rrbracket \gamma = 0 \end{cases} \\ \llbracket \langle M, N \rangle \rrbracket &= \gamma \mapsto (\llbracket M \rrbracket \gamma, \llbracket N \rrbracket \gamma) \\ \llbracket \text{fst } M \rrbracket &= \gamma \mapsto \pi_1(\llbracket M \rrbracket \gamma) \\ \llbracket \text{snd } M \rrbracket &= \gamma \mapsto \pi_2(\llbracket M \rrbracket \gamma) \\ \llbracket x \rrbracket &= \gamma \mapsto \pi_x(\gamma) \\ \llbracket \text{let } x \text{ be } M \text{ in } N \rrbracket &= \gamma \mapsto \llbracket N \rrbracket \gamma[x \mapsto \llbracket M \rrbracket \gamma] \end{aligned}$$

Notational note: we use “ $\gamma \mapsto v$ ” to mean the function that takes each substitution γ to value v , and π_1 and π_2 for the two projections out of a Cartesian product. As usual, we also use π_x to magically pick out the component of an environment γ that corresponds to a given variable x .

We can use this semantics to validate the reordering from earlier, as both side of the equation denote the same function.

$$\llbracket \text{let } x \text{ be } M \text{ in } (\text{let } y \text{ be } N \text{ in } \langle x, y \rangle) \rrbracket = \gamma \mapsto (\llbracket M \rrbracket \gamma, \llbracket N \rrbracket \gamma) = \llbracket \text{let } y \text{ be } N \text{ in } (\text{let } x \text{ be } M \text{ in } \langle x, y \rangle) \rrbracket$$

But, notice that there are subtleties! We are relying on our metalanguage’s notion of equality of functions here to determine equality.

Two functions $f, g : A \rightarrow B$ are considered equal if they are *extensionally equal*, meaning that for all $a \in A$, it is the case that $f(a) = g(a)$.

If we carefully unpacked these two semantics, we would see that the functions $\llbracket M \rrbracket$ and $\llbracket N \rrbracket$ are composed in different orders. Under the semantics of sets, this order of composition does not matter since it does not change the input-output behavior.

1.2 The Need for Generality

To sum up what just happened: in order to prove a natural program equivalence involving reorder of let-bindings, we gave a denotational semantics to `CALC` in terms of sets and functions. Then, a property of the metalanguage – extensional equality of functions – directly implied that the rewriting was semantics-preserving.

This style of denotational reasoning worked quite well for validating program equivalences, which begs the question: *how well does it generalize* to other kinds of programming languages and analyses?

There are several interesting dimensions along which we may want to generalize this kind of argument:

1. To languages with more interesting kinds of effects, such as randomness, nondeterminism, or mutable state;
2. To languages with more interesting features, such as higher-order functions or more interesting type systems.

We will see that it is actually quite inconvenient to use sets and functions as a metalanguage, which will be a central motivation our development of category theory as an alternative and more flexible metalanguage for giving semantics to programs (and other things, like program logics). But before we can see this, let us first see an example of each of these kinds of generalization.

1.3 Generalization 1: Probabilistic Effects

Let’s see an example of generalizing to languages with an interesting probabilistic effect. Consider the following little *probabilistic programming language* called `TINYPPPL` that extends `CALC` with a primitive `flip` that generates random Booleans:

$$M, N ::= \dots \mid \text{flip} \quad (\text{TINYPPPL})$$

Probabilistic programming languages denote probability distributions. For instance, here is the denotation of `flip`:

$$\begin{aligned} \llbracket \text{flip} \rrbracket &= \llbracket \text{Bool} \rrbracket \rightarrow [0, 1] \\ \llbracket \text{flip} \rrbracket = b &\mapsto \begin{cases} 1/2 & \text{if } b = \text{true} \\ 1/2 & \text{if } b = \text{false} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

We are again interested in proving natural equivalences such as the one we saw earlier where let-bindings can be reordered (recall that N does not refer to x):

$$\begin{aligned} &\text{let } x \text{ be } M \text{ in } (\text{let } y \text{ be } N \text{ in } \langle x, y \rangle) \\ &\quad \stackrel{?}{=} \\ &\text{let } y \text{ be } N \text{ in } (\text{let } x \text{ be } M \text{ in } \langle x, y \rangle) \end{aligned}$$

Intuitively, it seems like this ought to be a semantics-preserving rewrite for a very similar reason to the fact that an identical rewrite is true for `CALC`. But, `TINYPPPL` does not have an identical denotation, and so it seems like it requires a from-scratch proof to establish that this rewriting is sound.

Clearly this reordering is a general phenomenon we would like to study across many languages with different semantic interpretations. Category theory gives us a name for property of a semantic domain: it is called **commutativity**. Moreover, category theory gives us language for describing properties of a wide array of effects: probability is an example of a **monad**, and this monad is commutative and hence such rewritings are valid for *all commutative monads*.¹

1.4 Generalization 2: Higher-order functions

Functional languages have more than let-bindings and pairs. For instance, OCaml has higher-order functions, sum-types, and pattern-matching. These features can be difficult to specify, and category theory gives us a canonical way of designing these specifications.²

We will illustrate this principle by higher-order functions:

$$\begin{aligned} A, B &::= \text{Unit} \mid A \times B \mid A \rightarrow B \\ M, N &::= \langle \rangle \mid \langle M, N \rangle \mid \text{fst } M \mid \text{snd } M \mid \lambda x. M \mid MN \mid x \end{aligned} \quad (\text{STLC})$$

The **function type** $A \rightarrow B$ denotes functions from A to B . Functions are introduced using lambda abstraction rule $\lambda x. M$, defines a function with free variable x . Functions are eliminated with the lambda application rule MN , which calls the function M with argument N .

¹ This commutativity property was studied first by [35]; it follows essentially the order of summation can be interchanged. In the continuous case, this fact is true by Fubini's theorem.

² Indeed, the C in OCaml stands for "categorical" [8].

Following our original setup with `CALC`, we can give a denotational semantics to `STLC` by associating each type with a set and each term with a function between sets. For functions and function applications, this looks like:

$$\begin{aligned} \llbracket \lambda x. M \rrbracket = \gamma \mapsto (v \mapsto \llbracket M \rrbracket (\gamma[x \mapsto v])) \\ \llbracket MN \rrbracket = \gamma \mapsto (\llbracket M \rrbracket \gamma)(\llbracket N \rrbracket \gamma) \end{aligned}$$

You should notice that it is not obvious that this interpretation is valid: it remains to be argued that $(v \mapsto \llbracket M \rrbracket (\gamma[x \mapsto v]))$ is itself representable as an element of a set. This is true because it is possible to represent the set of all functions between two sets as a set: a function can be represented as its *graph* (i.e., a set of input-output pairs), which is itself a set.

Again, following our previous development from `CALC` to `TINYPPPL`, we wonder: is it similarly possible to represent first-class functions in the denotation of a probabilistic programming language? It's not at all obvious, at first, what shape a first-class function in a probabilistic programming language should have and whether it's possible to represent such an object. In discrete probability it is possible to represent first-class functions, but in continuous probability it is not at all so straightforward: this was the problem being studied by [14].

The crucial idea is that category theory will give us a general characterization of what it means for a semantic domain to support a first-class function: these will be called *exponential objects*, and categories that have exponential objects are called *Cartesian closed* and are capable of giving denotations to simply-typed lambda calculi.

1.5 Conclusion

This invites a question: *what sort of structure must a mathematical object have in order to be used to interpret a programming language?* Among many other things, category theory gives us very satisfying answers to this question: we will see that certain kinds of mathematical structure admit properties like “higher order functions”, a meaning of “natural numbers”, “logical truth,” and the “existence of fixed points”. Category theory gives us a formal language for describing what kind of structure a mathematical setting has. But before we get there, we'll see more examples of interpreting programs.

2

Categories

As a metalanguage, sets and functions work best as a model of *purely functional* programming, where a program is thought of as a black box that produces a uniquely determined output value for each input. But real programs are rarely purely functional: programs in the wild regularly manipulate resources (pointers, files, locks), interact with their environment (input and output), and use complex control flow constructs (exceptions, coroutines).

Category theory provides a framework for *swapping out the ambient metalanguage*: it provides techniques for systematically replacing sets and functions with constructs that are more naturally equipped to model realistic programming features. Each category embodies a choice of metalanguage. Sets and functions live in a *category of sets* Set , which we will encounter later on. By swapping out Set with other categories, we gain access to other metalanguages that are better adapted to modelling programming language features:

- Traditional domain theory, of untyped lambda calculus and recursive functional programming such as Haskell, lives in a *category of domains*. As a metalanguage, it supports arbitrary recursive definitions, and recursive types.
- A lot of modern semantics takes place in a category of *step-indexed sets* which as a metalanguage supports *very* recursive types.¹ We will meet this category later, when we discuss presheaves.
- Programs that manipulate resources live in categories of *resourceful sets*, which are used to define modern separation logics. We will meet these categories when we discuss the Day convolution.
- Ordinary continuous probability lives in a *category of measurable spaces*, and its higher-order generalization lives in *quasi-Borel spaces*. We may get to these categories if time permits.

The concept of “category” abstracts over all of these examples. First, a category is defined on a set of **objects**: these interpret types and


¹ For instance, the type $A = \blacktriangleright \wp(A)$ that is isomorphic to its own power set (up to the so-called “later modality” \blacktriangleright).

typing contexts. Objects are related to each other by **morphisms**

$A \xrightarrow{f} B$, or sometimes just f for short; these interpret terms. Each

morphism $A \xrightarrow{f} B$ is associated to a **domain** A and a **codomain** B .

Here is a **picture of a category**² with some objects and morphisms in it:




² A picture of a category shows objects as dots and morphisms as arrows, typically drawn with a box.

Morphisms can be combined via *composition*, which pastes together

morphisms that are incident to each other to form new ones. Composing g after f is denoted $g \circ f$, or sometimes more tersely as gf .

The ordering is a bit confusing; I like to read \circ as “after”. Pictorially, we can paste together edges f and g to get an edge h :



Categories satisfy several requirements that force them to behave nicely:

- Every object has a self-loop called the **identity morphism**, for an object A , we will denote its identity morphism as id_A . This morphism must behave like an identity: for any morphism $A \xrightarrow{f} B$, it must be the case that $\text{id}_B \circ f = f = f \circ \text{id}_A$.
- Composition is **associative**, meaning that for any three morphisms $A \xrightarrow{f} B, B \xrightarrow{g} C, C \xrightarrow{h} D$, it is the case that $(f \circ g) \circ h = f \circ (g \circ h)$.

Summing up,³

Definition 1: Category

A category is a tuple $(O, M, \text{dom}, \text{cod}, \text{id}, \text{comp})$ where

- O is a set whose elements are called “objects”
- M is a set whose elements are called “morphisms”

³ This definition of a category is one of many possible definitions. You may have seen different ones. This definition is used by [9]; different ones are used by [26].

- $\text{dom} : M \rightarrow O$
- $\text{cod} : M \rightarrow O$
- $\text{id} : O \rightarrow M$
- $\text{comp} : \{(f, g) \in M \times M \mid \text{dom}(f) = \text{cod}(g)\} \rightarrow M$

such that

- For all f and g in M such that $\text{dom}(f) = \text{cod}(g)$ it holds that

$$\text{dom}(\text{comp}(f, g)) = \text{dom}(g) \quad \text{and} \quad \text{cod}(\text{comp}(f, g)) = \text{cod}(f)$$

- For all X in O it holds that $\text{dom}(\text{id}_X) = \text{cod}(\text{id}_X) = X$
- For all f, g, h in M such that $\text{dom}(f) = \text{cod}(g)$ and $\text{dom}(g) = \text{cod}(h)$, it holds that

$$\text{comp}(f, \text{comp}(g, h)) = \text{comp}(\text{comp}(f, g), h)$$

- For all f in M it holds that

$$\text{comp}(f, \text{id}_{\text{dom}(f)}) = f = \text{comp}(\text{id}_{\text{cod}(f)}, f)$$

This official definition is a bit unwieldy, so in practice we will use a lot of shorthand. We say two morphisms f, g are *composable* if $\text{dom}(f) = \text{cod}(g)$. The composition $\text{comp}(f, g)$ of two composable morphisms will be written $f \circ g$. If we ever write $f \circ g$, it will be tacitly assumed that f and g are composable. Finally, we sometimes omit the subscript from id_X when X is clear from context. With these abbreviations in mind, the final two laws in the above definition read more comfortably as:

- $f \circ (g \circ h) = (f \circ g) \circ h$
- $f \circ \text{id} = f = \text{id} \circ f$.

2.1 Our first metalanguage

With a precise definition of category in hand, we are now ready to meet our first “metalanguage”.

Definition 2: the category FinSet of finite sets

Let FinSet be the category $(O, M, \text{dom}, \text{cod}, \circ, \text{id})$ where

- O is the set of finite sets.
- M is the set of tuples (A, f, B) where A and B are finite sets and f is a function from A to B .
- dom is the function $\text{dom}(A, f, B) = A$.


Sticklers may raise an eyebrow at the phrase “set of finite sets” in this definition. Strictly speaking, the set of finite sets does not form a set. But it’s not that big of a deal: many workarounds are available. For instance, one could consider just finite subsets of \mathbb{N} , or the set of *hereditarily finite sets* (defined to be the union $\bigcup_{n \in \mathbb{N}} \wp^n(\emptyset)$).

- cod is the function $\text{dom}(A, f, B) = B$.
- \circ is the function $(B, f, C) \circ (A, g, B) = (A, f \circ g, C)$ where $f \circ g$ is the ordinary function composition of f and g , well defined because the domain of f matches the codomain of g .
- id_A is (A, i, A) where i is the identity function on A .

The category laws are satisfied because function composition is associative and composing a function with the identity function does nothing.

Again, shorthand helps to make this definition palatable: we often conflate the tuple (A, f, B) the underlying function f when working with morphisms of FinSet .

Mathematically, FinSet provides a categorical home for finite set theory. As a metalanguage, FinSet embodies *finite purely functional programming*. For instance, the following is a graph of the negation function $f : \{T, F\} \rightarrow \{T, F\}$ in the category FinSet , which could serve as a denotation of a program like $x : \text{Bool} \vdash \text{if } x \text{ then false else true} : \text{Bool}$.



2.2 Gaining intuition for categories

The definition of category is highly abstract, and categories in the wild can look quite different from FinSet . In time, we will see how this level of abstraction allows for clean axiomatizations of the various kinds of “features” that a metalanguage may support, in the form of so-called “universal constructions.” But in order to work at this level of abstraction, it will help to have in mind a few concrete examples. Broadly speaking, there are three different classes of example that are handy to have in the back of one’s mind when thinking about categories.

2.2.1 Categories as graphs

The first perspective, perhaps most comfortable to computer scientists, is that a category is a kind of graph. Objects are vertices and morphisms are edges. On top of this, categories come with a composition operation that pastes two incident edges together to form a third, and identity morphisms id_A that form self-loops centered at each vertex A . The category laws then say that pasting of edges is an

associative operation, and that pasting with an identity morphism does nothing. This provides nice visual intuition, and a large stock of examples. In fact,

Construction 2.2.1. Every directed multigraph can be turned into a category.

Proof. Let G be a directed multigraph, encoded as a tuple $(V, E, \text{src}, \text{trg})$ where V is the set of vertices, E is the set of edges, and src, trg are functions $E \rightarrow V$ that send each edge to its source and target vertices respectively. The graph G can be used to define a category

$$\text{paths}(G) = (O, M, \text{dom}, \text{cod}, \circ, \text{id})$$

where

- $O = V$
- M is the set of paths in G , which are tuples $(v_s, [e_1, \dots, e_n], v_t)$ where v_s, v_t are vertices and e_1, \dots, e_n is a list of edges satisfying $\text{src}(e_1) = v_s$ and $\text{trg}(e_n) = v_t$ and $\text{trg}(e_i) = \text{src}(e_{i+1})$ for all $1 \leq i < n$.
- $\text{dom} : M \rightarrow O$ is the function $\text{dom}(v_s, [e_1, \dots, e_n], v_t) = v_s$
- $\text{cod} : M \rightarrow O$ is the function $\text{cod}(v_s, [e_1, \dots, e_n], v_t) = v_t$
- \circ concatenates paths: $(v_t, [f_1, \dots, f_n], v_u) \circ (v_s, [e_1, \dots, e_m], v_t) = (v_s, [e_1, \dots, e_m, f_1, \dots, f_n], v_u)$. Note that this is well-defined because the two paths \vec{f} and \vec{e} line up tip-to-tail. Note also that \vec{e} comes before \vec{f} in the concatenated path, due to the reversal of order of arguments to \circ .
- id_v is the empty path $(v, [], v)$ starting and ending at v .

The category laws are satisfied because concatenation of paths is associative and concatenating a path with the empty path does nothing. \square

Note 2.2.2. It is essential that morphisms are *paths* in this construction, and not edges! Otherwise, it would be unclear how to define composition and identity.

A key difference between categories and graphs is that two morphisms in a category that come from different “paths” may nonetheless be equal. For example, the following two “paths” through FinSet are actually the same, because negating a Boolean twice is the same as doing nothing:

$$\begin{array}{ccc}
 \{T, F\} & \xrightarrow{\text{not}} & \{T, F\} \\
 & \searrow \text{id}_{\{T, F\}} & \swarrow \text{not} \\
 & \{T, F\} &
 \end{array}$$

This relationship between the paths $\text{not} \circ \text{not}$ and $\text{id}_{\{T,F\}}$ is often conveyed by saying that the triangle above is “filled in”, or “commutes”. The intuition being that one can then “continuously deform” $\text{not} \circ \text{not}$ into $\text{id}_{\{T,F\}}$, by “dragging” it through the filled in triangle.

This idea is a handy visual tool for organizing categorical information and for conducting proofs.

2.2.2 Categories as orders

A special case of a multigraph is a graph, where there is *at most one* arrow between any given pair of vertices. This leads us to the second class of important examples of a category:

Definition 3: preorder


A preorder is a pair (X, \preceq) where X is a set and \preceq is a binary relation on X that is reflexive ($x \preceq x$ for all x in X) and transitive (if $x \preceq y$ and $y \preceq z$ then $x \preceq z$).

Preorders are often fruitfully drawn as *Hasse diagrams*. A simple example is the set \mathbb{N} of natural numbers, ordered by $m \preceq n$ if m is less than or equal to n . We can draw this preorder as follows:



Hasse diagrams are read bottom-up, with the least elements on the bottom and the greatest elements on the top. An undirected edge is drawn between elements with no elements between them (this is why we have no edge from 0 to 2 in the above figure). A more visually interesting preorder on \mathbb{N} is to define $m \preceq n$ if m is a *divisor* of n .

This preorder looks less like a straight line and more like a lattice:



Preorders don't have to be discrete either: the real numbers \mathbb{R} , for instance, form a preorder with $x \preceq y$ if x is less than or equal to y . This is a little harder to draw.

A preorder that arises frequently in PL is the preorder of *run-time environments*. It is common to model these environments as partial functions $\rho : \text{Var} \rightarrow \text{Val}$, assigning to each program variable $x \in \text{Var}$ a concrete value $\rho(x)$. These substitutions form a preorder, with $\rho \preceq \rho'$ if ρ is a subset of ρ' , in the sense that all the variables in ρ are in ρ' and mapped to the same values.

Preorders form a category:

Construction 2.2.3. Every preorder can be turned into a category.

Proof. For any preorder (X, \preceq) , there is a category

$$\text{order}(X, \preceq) = (O, M, \text{dom}, \text{cod}, \circ, \text{id})$$

where

- $O = X$.
- $M = (\preceq)$ (recall that a binary relation like \preceq can be thought of as a set of pairs, with $(x, y) \in (\preceq)$ if and only if $x \preceq y$).
- dom is the function $\text{dom}(x, y) = x$.
- cod is the function $\text{cod}(x, y) = y$.
- \circ is the function defined by $(y, z) \circ (x, y) = (x, z)$. Note that we may assume that the two arguments to \circ share a component y in this definition because of the assumption that (y, z) and (x, y) are composable. Note also that the composite (x, z) is indeed well-defined this is well defined by the transitivity of \preceq .
- id_x is the pair (x, x) , which forms a valid morphism by reflexivity of \preceq .

The category laws are satisfied automatically, as there can only be at most one morphism between any two objects.⁴

□

⁴ Pause: why is this the case?

The intuitive import of Proposition 2.2.3 is that every category can be thought of as a kind of preorder.⁵ where the existence of a morphism $A \xrightarrow{f} B$ expresses that A is in some sense “less than or equal to” B . Thinking of categories in this way leads to interesting questions like:

- Does a category have a “smallest” or “largest” object? For instance, the smallest runtime environment is the empty environment \emptyset with no variables in it.
- Does a category have “ascending sequences”? For example,

$$\emptyset \subseteq \{x \mapsto A\} \subseteq \{x \mapsto A, y \mapsto B\} \subseteq \dots \quad (2.1)$$

is an ascending sequence in the preorder of runtime environments.

- Does a category have “limit” objects, like suprema or infima? For instance, you probably know from calculus that there limits in \mathbb{R} : for example, the sequence $\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16} \dots$ approaches 0 in the limit. A very interesting question is: which categories “have limits”, in the sense that such sequences always converge to some object in the category?

An example of a situation where there is no limit is the sequence in Eq. (2.1). The set of environments grows and grows, and since there is no notion of a “greatest environment”, there is no limit to this sequence.

- Given two categories, is there a notion of “monotone map” between them?

Each of the above observations is foreshadowing a notion of a “universal construction” in a category, which we will explore in the following lectures.

2.2.3 Categories as algebras

Whereas preorders are like multigraphs where there is at most one *edge* between any two vertices, the opposite extreme is a multigraph with exactly one *vertex*, so that all edges are self-loops:



Categories of this shape form our final class of important example: a special kind of algebraic structure called a *monoid*.

⁵ *Reminder*: not all categories directly correspond to some preorder construction, since there can be multiple morphisms between objects in a general category. A category with at most 1 morphism between any 2 objects is called a *thin category*, and it does correspond to some preorder.

Definition 4: monoid

A monoid is a tuple (X, \bullet, e) where \bullet is a binary operation on X (a function $X \times X \rightarrow X$) that is associative ($x \bullet (y \bullet z) = (x \bullet y) \bullet z$) and has e as a unit ($x \bullet e = x$ and $e \bullet x = x$).

A simple example of a monoid is $(\mathbb{N}, +, 0)$, the monoid of natural numbers under addition. This forms a monoid because addition of natural numbers is associative and has zero as a unit.

An example that is closer to PL comes from the semantics of imperative programs. It is common to interpret an imperative program as a function $S \rightarrow S$, where S is a set of memory states. For example, the program

$$x := x + 1;$$

can be interpreted as a function that sends a memory state s to a new memory state s' in which the value of x has been incremented by 1. The functions $S \rightarrow S$ form a monoid $(S \rightarrow S, \bullet, \text{id})$, where the monoid operation is given by $f \bullet g = g \circ f$ (note the reversal!), corresponding to the sequencing of the programs f and g , and the identity element is the identity function on S , corresponding to the empty program that does nothing.

Algebraic identities capture program rewrites for this imperative language. For example, if f is the interpretation of $x := 0$ and g is the interpretation of $y := 1$, and x and y are distinct program variables, then the program equation

$$\begin{pmatrix} x := 0 \\ y := 1 \end{pmatrix} \equiv \begin{pmatrix} y := 1 \\ x := 0 \end{pmatrix} \quad (2.3)$$

can be expressed as the algebraic identity $g \circ f = f \circ g$.

Construction 2.2.4. Every monoid can be turned into a category.

Proof. For any monoid (X, \bullet, e) there is a category

$$\text{mon}(X, \bullet, e) = (O, M, \text{dom}, \text{cod}, \circ, \text{id})$$

where

- $O = \{\star\}$, a singleton set with a single dedicated element \star .
- $M = X$
- $\text{dom}(M) = \star$
- $\text{cod}(M) = \star$
- Composition is defined by $x \circ y = x \bullet y$

- $\text{id}_* = e$

The category laws follow from the associativity of (\bullet) and that it has e as a unit. \square

Note 2.2.5. It is very important that the set of objects O is a singleton set $\{*\}$. A natural thought one might have is that the set of objects should be the set X of elements of the monoid. But this in fact does not work: what would be the morphisms?

There is actually a good reason why the set of objects is the seemingly-arbitrary singleton set $\{*\}$ in this construction. It comes from the fact that a monoid is what is called a *single-sorted* algebraic structure; hence $\{*\}$ is a set with a single element.

The intuitive import of Proposition 2.2.4 is that every category can be thought of as a kind of algebraic structure, where composition is like “multiplication” and identity morphisms are like “multiplicative units”.

The algebraic perspective leads to questions like: given an equation between morphisms containing some unknowns, are there zero, one, or infinitely many solutions? When can I “cancel” a morphism from both sides of an equation (e.g., from $fg = fh$ deduce $g = h$, or from $fh = gh$ deduce $f = g$)? When does a morphism have an “inverse”? When can a morphism f be “factored” into two pieces, $f = gh$?

3

Universal Constructions I: Products

We've been hinting at the utility of knowing that categories have certain special objects in them. For example, if we want to give a category that lets us validate equations of programming languages, that category will need to be able to represent the many features that programs might have: products, sums, functions, etc. Here we will introduce the essential idea of a **universal construction** that characterizes when categories contain these special objects.

3.1 Terminal objects & the unit type

Recall that when designing categorical semantics for programming languages we associate types and typing contexts with objects in the category and terms with morphisms. Concretely, if we have a well-typed term $\Gamma \vdash e : A$, its interpretation $\llbracket \Gamma \vdash e : A \rrbracket$ is a morphism $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket e \rrbracket} \llbracket A \rrbracket$. Previously, we showed how `FinSet` can be used to give an interpretation to `CALC` programs by associating types like the unit type `Unit` with special sets, like a set containing only a single element $\{\star\}$. But, what if we want to see if categories other than `FinSet` can give a reasonable interpretation for `Unit`?

To proceed, let's pick a particular object $\llbracket \text{Unit} \rrbracket$ in a category \mathcal{C} and explore what properties it needs to have in order to be "Unit-like":

- **Fact 1, Introduction rules:** It must be possible in any typing context to produce a value of type `Unit`. This tells us that for any context Γ , there must exist a \mathcal{C} -morphism $\llbracket \Gamma \rrbracket \rightarrow \llbracket \text{Unit} \rrbracket$. For example, it must be the case that $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket \langle \rangle \rrbracket} \llbracket \text{Unit} \rrbracket$.
- **Fact 2, Equational laws:** Since `CALC` has no effects, it is the case that if $\Gamma \vdash e : \text{Unit}$, then it is indistinguishable from the term $\langle \rangle$. This is called **eta equality**, and we write this as $e \equiv \langle \rangle$.¹ Eta equality is an equational law that characterizes how a `Unit` term must behave. Now for an insight: *these equational laws are translated*

¹ The symbol " \equiv " means "is equationally equal to".

into morphism equalities. It must be that, for any morphism $[[\Gamma]] \xrightarrow{[[e]]} [[\text{Unit}]]$, it is the case that $[[e]] = [[\langle \rangle]]$. In other words, the morphism $[[\Gamma]] \xrightarrow{[[\langle \rangle]]} \text{Unit}$ must be unique.

These two properties can be packaged up into a nice concise definition:

Definition 5: Terminal object

Let T be an object in a category \mathcal{C} . An object T in \mathcal{C} is called a terminal object if

- for every object Γ in \mathcal{C} there is a morphism $\Gamma \xrightarrow{\langle \rangle} T$;
- for every morphism $\Gamma \xrightarrow{f} T$ it holds that $f = \langle \rangle$.

Let's pause to remark on the three categorical interpretations of terminal objects T :

- Graphically, an object is terminal if (1) there is an arrow from every other object to it, and (2) it "thins all parallel arrows", meaning that if we have a situation like the following for f and g morphisms $X \xrightarrow{f} T, X \xrightarrow{g} T$:

$$\begin{array}{c} T \\ f \nearrow \quad \searrow g \\ X \end{array},$$

then we know that $f = g$. In other words, there is a unique morphism $\Gamma \rightarrow T$ for every Γ :

$$A \xrightarrow{\exists!} T$$

- The order-theoretic perspective is that the terminal object is the "greatest object" in the category.
- The algebraic perspective says that, for every object Γ , the equation $f = f$ has a unique solution in the unknown $\Gamma \xrightarrow{f} T$.

We say a category **has terminal objects** if there is an object in the category that is a terminal object. This is our first example of a **universal construction**, which are special objects in categories characterized by their morphisms into them. We will see how universal constructions can be used to give interpretations to all the type formers we have seen so far. But, we can already see their utility: whether or not a particular category admits certain universal constructions will immediately inform us of that category's suitability for modeling certain language properties.

3.1.1 FinSet has terminal objects, isomorphisms

We saw in previous lectures that we could use the singleton set $\{*\}$ to give an interpretation to **Unit** in FinSet. We can easily prove that $\{*\}$ is a terminal object:

Theorem 3.1.1. The singleton set $\{*\}$ is terminal in FinSet.

Proof. • Existence: The map $_ \mapsto \{*\}$ has the required type.

- Uniqueness: Suppose we have two morphisms $A \xrightarrow{f} \{*\}$ and $A \xrightarrow{g} \{*\}$. Then by function extensionality we immediately have that $f = g$ as set-functions, and so they are equal as FinSet morphisms. □

Now you can observe the above proof and see that it immediately works for *any* choice of singleton set: does this mean that FinSet has *many terminal objects*? Yes, and no. *Yes*, in the sense that there are indeed many valid choices of terminal object; but *no*, in the sense that the choice itself does not (*should* not!) actually matter. We capture this by the notion of isomorphism:

Definition 6: Isomorphism of objects
Two objects A and B in a category \mathcal{C} are isomorphic if there exist morphisms $A \xrightarrow{f} B$ and $B \xrightarrow{g} A$ such that (1) $g \circ f = \text{id}_A$ and (2) $f \circ g = \text{id}_B$.

We can immediately see that any two singleton sets in FinSet are isomorphic: this proof is straightforward. What might be surprising is that this fact is true for all terminal objects:

Theorem 3.1.2. Terminal objects are unique up to isomorphism.

Proof. Let X and Y be terminal objects in some category \mathcal{C} . Then there exists a pair of unique morphisms between them:

$$\begin{array}{ccc}
 & \xrightarrow{!f} & \\
 X & & Y \\
 & \xleftarrow{!g} &
 \end{array}$$

By the fact that X and Y are terminal, they must also have unique identity morphisms id_X and id_Y . Then by the uniqueness of identity we have that $g \circ f = \text{id}_X$ and $f \circ g = \text{id}_Y$. □

This provides some evidence that the choice of terminal object does not matter. But are we done? Suppose we had two terminal objects T and T' . By the above argument, we have that these two objects are isomorphic. But now another potential choice arises: what

if there are multiple different *isomorphisms* between T and T' ? Then it could be that the choice to use T instead of T' may implicitly come with a choice of specific isomorphism as well.

Luckily, terminal objects satisfy a stronger property which ensures that one's choice of terminal object is truly irrelevant.

Proposition 3.1.3. Terminal objects are unique up to *unique* isomorphism. That is, if T and T' are both terminal, then there is exactly one isomorphism $(f : T \rightarrow T', g : T' \rightarrow T)$ between T and T' .

Proof. Both f and g , as constructed in the proof of Theorem 3.1.2, are the only possible morphisms $T \rightarrow T'$ and $T' \rightarrow T$ respectively. Thus they together form the only isomorphism between T and T' . \square

Later on, when we have the language for it, we will show that all universal constructions are unique up to unique isomorphism in this sense, and hence free of non-canonical choices.

3.2 Products

Next we want to interpret product types $A \times B$. As before with the unit type, the typing rules for products forces the existence of certain morphisms, and the equational theory of products forces equalities of certain morphisms. The introduction rule that describes how to create new products is:

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \langle M, N \rangle : A \times B} \text{ (T-PAIR)}$$

The elimination rules that break apart products are:

$$\frac{\Gamma \vdash M : A \times B}{\Gamma \vdash \mathbf{fst} M : A} \text{ (T-FST)} \quad \frac{\Gamma \vdash M : A \times B}{\Gamma \vdash \mathbf{snd} M : B} \text{ (T-SND)}$$

These rules imply the existence of several morphisms that characterize what it means for a special object $\llbracket A \times B \rrbracket$ to be a product. Reading off from the rules:

- The T-PAIR rule says that if there are two morphisms $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket M \rrbracket} \llbracket A \rrbracket$ and $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket N \rrbracket} \llbracket B \rrbracket$, then there is a morphism $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket \langle M, N \rangle \rrbracket} \llbracket A \times B \rrbracket$.
- The T-FST rule says that if there is a morphism $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket M \rrbracket} \llbracket A \times B \rrbracket$, then there is a morphism $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket \mathbf{fst} - \rrbracket} \llbracket A \rrbracket$.²
- The T-SND rule says that if there is a morphism $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket M \rrbracket} \llbracket A \times B \rrbracket$, then there is a morphism $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket \mathbf{snd} - \rrbracket} \llbracket B \rrbracket$.

² Note: We are naming these morphisms *suggestively*. The typing rules only enforce the *existence* of certain morphisms: they don't tell us anything about how they behave.

We can package these morphisms up into the following diagram:

$$\begin{array}{c}
 \llbracket \Gamma \rrbracket \\
 \swarrow \quad \searrow \\
 \llbracket M \rrbracket \quad \llbracket N \rrbracket \\
 \downarrow \llbracket \langle M, N \rangle \rrbracket \\
 \llbracket A \times B \rrbracket \\
 \swarrow \quad \searrow \\
 \llbracket A \rrbracket \quad \llbracket B \rrbracket \\
 \leftarrow \llbracket \text{fst } - \rrbracket \quad \llbracket \text{snd } - \rrbracket \rightarrow
 \end{array} \quad (3.1)$$

Next, the equational theory of products will assert certain equalities of morphisms. We begin with the β -rules that give the equational behavior of elimination after introduction:

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \text{fst } \langle M, N \rangle \equiv M : A} (\beta_1^\times) \quad \frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \text{snd } \langle M, N \rangle \equiv N : B} (\beta_2^\times)$$

Translating these equations into morphism equalities:

- β_1 enforces that $\llbracket M \rrbracket = \llbracket \text{fst } - \rrbracket \circ \llbracket \langle M, N \rangle \rrbracket$ in (3.1)
- β_2 enforces that $\llbracket N \rrbracket = \llbracket \text{snd } - \rrbracket \circ \llbracket \langle M, N \rangle \rrbracket$ in (3.1)

Hence, the β -rules together enforce that the diagram in 3.1 commutes. Next, we consider the second part of the equational theory of products that describes the behavior of introduction after elimination, sometimes called the η -rules:

$$\frac{\Gamma \vdash M : A \times B}{\Gamma \vdash M \equiv \langle \text{fst } M, \text{snd } M \rangle : A \times B} (\eta^\times)$$

This law says asserts the extensional equality of any two pair-formation operations. Translated into morphisms, it says that any morphism of the form $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket M \rrbracket} \llbracket A \times B \rrbracket$ is equal to a canonical morphism $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket \langle \text{fst } M, \text{snd } M \rangle \rrbracket} \llbracket A \times B \rrbracket$. Put another way: there is a unique morphism $\llbracket \Gamma \rrbracket \rightarrow \llbracket A \times B \rrbracket$ making the above diagram **commute**.³

3.2.1 Packaging up intuition into a definition

Now, as in the case of terminal objects, we can give a more generic description of products that isn't so specific to the case of equational theories of programs:

Definition 7: Product

Let A and B be two objects of a category \mathcal{C} . A *product of A and B* is a tuple (P, π_A, π_B) where

³ Formally, a diagram **commutes** if every face of the diagram is “filled in”, meaning that all paths beginning and ending in the same place are equal as morphisms.

- P is an object of \mathcal{C} .
- π_A is a morphism from P to A , called the *projection onto A* .
- π_B is a morphism from P to B , called the *projection onto B* .

such that

- For any object Γ and any two morphisms $f : \Gamma \rightarrow A$ and $g : \Gamma \rightarrow B$, there exists a morphism $\langle f, g \rangle : \Gamma \rightarrow P$ such that the following equations hold:

$$\pi_A \circ \langle f, g \rangle = f \quad (3.2)$$

$$\pi_B \circ \langle f, g \rangle = g \quad (3.3)$$

- The following equation holds for any morphism $h : \Gamma \rightarrow P$.

$$h = \langle \pi_A \circ h, \pi_B \circ h \rangle \quad (3.4)$$

Proposition 3.2.1. Let A and B be finite sets. The tuple $(A \times B, \pi_A, \pi_B)$ is a product in FinSet , where π_A is the morphism $(A \times B, \pi_A, A)$ and π_B is the morphism $(A \times B, \pi_B, B)$.

Proposition 3.2.2. Let A and B be two objects of a category \mathcal{C} . A tuple (P, π_A, π_B) is a product of A and B if and only if the following **universal property** holds: for any object Γ and any two morphisms $f : \Gamma \rightarrow A$ and $g : \Gamma \rightarrow B$, there exists a unique morphism $h : \Gamma \rightarrow P$ such that $\pi_A \circ h = f$ and $\pi_B \circ h = g$.

Proof. Suppose (P, π_A, π_B) is a product of A and B . By definition of product, we have for any Γ and morphisms $f : \Gamma \rightarrow A$ and $g : \Gamma \rightarrow B$ that the morphism $\langle f, g \rangle : \Gamma \rightarrow P$ satisfies $\pi_A \circ \langle f, g \rangle = f$ and $\pi_B \circ \langle f, g \rangle = g$. This morphism is the unique such, because if any morphism $h : \Gamma \rightarrow P$ satisfies $\pi_A \circ h = f$ and $\pi_B \circ h = g$ then it must be that

$$h = \langle \pi_A \circ h, \pi_B \circ h \rangle = \langle f, g \rangle. \quad (3.5)$$


Conversely, if (P, π_A, π_B) is any tuple satisfying the above universal property, then one can define $\langle -, - \rangle$ as the function that sends each pair of morphisms $f : \Gamma \rightarrow A$ and $g : \Gamma \rightarrow B$ to the unique h such that $\pi_A \circ h = f$ and $\pi_B \circ h = g$. \square

Algebraically, the universal property can be formulated as saying the following system of equations has exactly one solution in the unknown $h : \Gamma \rightarrow P$.


$$\pi_A \circ h = f \quad (3.6)$$

$$\pi_B \circ h = g \quad (3.7)$$

Diagrammatically, it can be formulated as saying there exists a unique dashed arrow that fills in all the triangles in the following picture.



Order-theoretically, a product of A and B is the “greatest diagram” of the following shape, where the question marks are filled in with concrete objects and morphisms of \mathcal{C} :




We will make this perspective more precise in Section 3.6.

Finally, as is the case with terminal objects, we say a category **has products** if for any two objects A and B in a category there exists an object $A \times B$ satisfying the universal property of objects.

3.3 Products in a category that is a preorder


Universal properties serve as a powerful tool for identifying analogies between mathematical objects. The essence of a product is made clear by (3.1): it is a means for placing two pieces of data “side by side,” along with ways of extracting those two pieces of data after they’ve been packaged up without losing any information about those pieces of data. When viewed through the lens of universal properties, we start to see how general this idea of taking a product of two pieces of data can be.

Let’s see an example of a category that has an interesting and somewhat surprising example of products: the category of natural numbers ordered by divisibility. Natural numbers ordered by divisibility forms a preorder visualized by the following Hasse diagram:



The set (\mathbb{N}, \preceq) forms a preorder. So, what do products look like in (\mathbb{N}, \preceq) considered as a category? Let's consider the concrete instance of the product of 12 and 16. Let's consider the three possible interpretations of the product of 12 and 16 – let's call it X – to understand what it means here:

- The graph-theoretic picture places the product in a diagram. Suppose there is some natural number Y such that $Y \preceq 12$ and $Y \preceq 16$. Then, the graphical depiction of the product X asserts the existence of a unique morphism $Y \rightarrow X$:



- The order theoretic interpretation tells us that X is the greatest element that is less than both 12 and 16 in the partial order – i.e., it is the greatest lower bound. In the case of this particular partial order, this object has a special name: it is the greatest common divisor.

Finally, similar to terminal objects, we say that a category \mathcal{C} **has products** if for any two objects A and B in \mathcal{C} there exists an object $A \times B$ satisfying the universal property of products.

3.4 Getting comfortable working with products

Though we arrived at the definition of product through the β and η laws for product types, the universal property is taken as primary as it is the one that generalizes to other type formers. This property has three crucial parts. Let Γ, A, B be objects and $f : \Gamma \rightarrow A$ and $g : \Gamma \rightarrow B$. Then the following three facts hold:

1. There exists a morphism $h : \Gamma \rightarrow P$.
2. The morphism h satisfies $\pi_1 \circ h = f$ and $\pi_2 \circ h = g$.
3. The morphism h is the *unique* morphism satisfying property (2).
This is made precise as follows: if any morphism $h' : \Gamma \rightarrow P$ satisfies $\pi_1 \circ h' = f$ and $\pi_2 \circ h' = g$, then $h' = h$.

Taken together, these three bullet points say that the special morphism h is *defined by* the property that $\pi_1 \circ h = f$ and $\pi_2 \circ h = g$. A lot of proofs involving this universal property perform equational reasoning on the special morphisms h . Here is an example:

In what way is this behaving like an intersection?

Proposition 3.4.1. The following holds for all $p : X \rightarrow Y$ and $f : Y \rightarrow Z$ and $g : Y \rightarrow W$ and products (P, π_1, π_2) of Z and W .

$$\begin{aligned} & (\text{the unique } h \text{ such that } \pi_1 \circ h = f \text{ and } \pi_2 \circ h = g) \circ p \\ &= (\text{the unique } h \text{ such that } \pi_1 \circ h = f \circ p \text{ and } \pi_2 \circ h = g \circ p) \end{aligned}$$

Another example:

Proposition 3.4.2. Let (P, π_1, π_2) be a product of A and B in a category \mathcal{C} . Then

$$(\text{the unique } h \text{ such that } \pi_1 \circ h = \pi_1 \text{ and } \pi_2 \circ h = \pi_2) = \text{id}_P$$

It gets a little tiresome to write “the unique h such that $\pi_1 \circ h = f$ and $\pi_2 \circ h = g$ ” everywhere. So category theorists use a convenient notational abbreviation for this: the angle-bracketed expression $\langle f, g \rangle$.⁴ With this notation, the above propositions read as the following algebraic identities:


$$\langle \pi_1, \pi_2 \rangle = \text{id} \tag{3.9}$$

$$\langle f, g \rangle \circ p = \langle f \circ p, g \circ p \rangle \tag{3.10}$$


These algebraic identities can then be used to prove more elaborate facts about products.

Proposition 3.4.3. Let \mathcal{C} be a category with products. Then $A \times B \cong B \times A$ for any objects A, B of \mathcal{C} .

Proof. Since this is our first interesting proof we will give a fully diagrammatic argument. To show an isomorphism of objects we must find two morphisms:




such that $g \circ f = \text{id}_{A \times B}$, $f \circ g = \text{id}_{B \times A}$. We only have two ways of showing that two morphisms are equal to one another: by reading equations off commuting diagrams, or by showing that two morphisms satisfy the same universal property. Let’s first show that $\text{id}_{A \times B}$ satisfies the same universal property as two morphisms $g \circ f$. We will do this very slowly. Let’s start with the universal property of the product $A \times B$:




⁴ Note that this notational convention hides some data: in particular, it hides the codomain of $\langle f, g \rangle$. In practice this notational convention is quite useful because, as we will show later (in Theorem 3.6.1), products are unique up to unique isomorphism, so hiding which particular choice of product object is utilized is prudent. Throughout category theory this sort of careful notation can be used (and confused!) to simplify arguments, but it’s always important to know why a notational choice is valid.

The universal property for products says that, for any object X of the category, if it has morphisms $X \rightarrow A$ and $X \rightarrow B$, then there is a unique morphism $X \xrightarrow{h} A \times B$ such that the diagram commutes. Our proof will proceed by carefully picking different objects for X . First, we pick $X = A \times B$:




But, there is another morphism that satisfies this same universal property: the identity morphism!



Note how $\text{id}_{A \times B}$ also makes this diagram commute. This establishes that $\text{id}_{A \times B} = \langle \pi_1, \pi_2 \rangle$ (i.e., it is a graphical proof of Eq. (3.9))

Now, let's find morphisms $g \circ f$ that satisfy the same universal property as $\text{id}_{A \times B}$. Now we can do this by constructing a chain of products:



Look closely: in the above diagram, $\langle \pi_2, \pi_1 \rangle \circ \langle \pi_1, \pi_2 \rangle$ satisfies the same universal property as $\text{id}_{A \times B}$ in Eq. (3.11). So, by the universal property of products, we have that $\langle \pi_2, \pi_1 \rangle \circ \langle \pi_1, \pi_2 \rangle = \text{id}_{A \times B}$. So, we've found our f and g . The other half of the isomorphism can be shown by picking the other ordering of products in the sequence of diagrams. □

Next, here is a theorem with a familiar shape, if one thinks of products like the familiar notion of product of integers and terminal objects as behaving like the product unit:

Proposition 3.4.4. Let \mathcal{C} be a category with products and a terminal object. Then $1 \times A \cong A$ for any object A of \mathcal{C} .

Definition 8: Ternary product

Let A, B, C be objects of a category \mathcal{C} . A *ternary product* of A, B, C is a tuple (P, π_A, π_B, π_C) where $\pi_A : P \rightarrow A$ and $\pi_B : P \rightarrow B$ and $\pi_C : P \rightarrow C$ and the following universal property holds: for any object Γ and morphisms $f : \Gamma \rightarrow A$ and $g : \Gamma \rightarrow B$ and $h : \Gamma \rightarrow C$, there exists a unique morphism $k : \Gamma \rightarrow P$ such that $\pi_A \circ k = f$ and $\pi_B \circ k = g$ and $\pi_C \circ k = h$.

Proposition 3.4.5. Let A, B, C be objects in a category \mathcal{C} with products. The tuple $(A \times (B \times C), \pi_A, \pi_B \circ \pi_{B \times C}, \pi_C \circ \pi_{B \times C})$ is a ternary product of A, B, C .

Definition 9: Finite products

A category has **finite products** if it has products (Def 7) and a terminal object (Def 5).

3.5 *Initial objects & duality*

The categorical perspective can push you towards new ideas you may not have thought of. Here’s an example of this: given a category \mathcal{C} , thought of as a graph, one can do a natural thing: *flip the direction of all arrows in the opposite direction*. This new category is called the **opposite category**, and is written \mathcal{C}^{op} . Relationships between properties of \mathcal{C} and properties of its opposite category \mathcal{C}^{op} are broadly referred to as **dualities**.

Dualities can be a source of free ideas. For instance, let’s consider the dual notion of a terminal object where we flip the direction of the arrows:

Definition 10: Initial object

Let I be an object in a category \mathcal{C} . An object I in \mathcal{C} is called an **initial object** if, for every other object A in \mathcal{C} there exists a unique morphism $I \rightarrow A$. Graphically,

$$I \xrightarrow{\exists!} A$$

Note that an object that is terminal in \mathcal{C} must be initial in \mathcal{C}^{op} and vice versa: this is what characterizes the terminal object as dual to the initial object.

Now that we have a “free” universal construction, it might be natural to ask: *is there an interesting type that this special initial object corresponds to?* Indeed there is: the initial object corresponds to the **Void** type that has no inhabitants. It has the following typing rules:

$$\frac{M : \mathbf{Void}}{\Gamma \vdash \text{absurd } M : A}$$

Intuitively, the **Void** type represents logical absurdity or “invalid program state”. Some programming languages have a construct for working with void types and absurdity: Haskell has the `void`, which is eliminated with `absurd`.

3.6 Products as Terminal Objects (*)


The diagrammatic form suggests an order-based intuition for products: the product (P, p, q) is in some sense the “largest” tuple of that shape, with its universal property showing that any other tuple of the same shape is “less than or equal” to it by virtue of the existence of the dashed morphism.

Sections marked with * are optional and won’t be covered in class.

Definition 11: Category of rooted spans


Let A and B be objects of a category \mathcal{C} . The category of *spans in \mathcal{C} rooted at A and B* , written $\text{Span}_{\mathcal{C}}(A, B)$, is the category whose

- Objects are tuples (X, f, g) where X is an object of \mathcal{C} , $f : X \rightarrow A$, and $g : X \rightarrow B$. Pictorially:



- A morphism is a tuple $((X, f, g), \alpha, (X', f', g'))$ where (X, f, g) and (X', f', g') are objects and $\alpha : X \rightarrow X'$ is a morphism of \mathcal{C}

such that $f' \circ \alpha = f$ and $g' \circ \alpha = g$. Pictorially:




- Composition is defined by

$$((X', f', g'), \alpha', (X'', f'', g'')) \circ ((X, f, g), \alpha, (X', f', g')) \quad (3.13)$$

$$= ((X, f, g), \alpha' \circ \alpha, (X'', f'', g'')). \quad (3.14)$$

Pictorially:



- The identity morphism at (X, f, g) is $((X, f, g), \text{id}_X, (X, f, g))$.

The associativity and identity laws in $\text{Span}_{\mathcal{C}}(A, B)$ are inherited from the corresponding laws for \mathcal{C} .

This perhaps tortured-looking category serves a valuable purpose: it makes precise the sense in which a product (P, p, q) is the “largest” among such tuples.

Proposition 3.6.1. Let A and B be objects of a category \mathcal{C} . A tuple (P, p, q) is a product of A and B if and only if it is a terminal object of $\text{Span}_{\mathcal{C}}(A, B)$.

To unpack what this is saying, from the perspective of categories as metalanguages: a product for A and B in the metalanguage embodied by \mathcal{C} is a unit type for the metalanguage embodied by $\text{Span}_{\mathcal{C}}(A, B)$. Some of the power of category theory can be seen here:

- Constructions on categories can be used to quickly build nontrivial “metalanguages” from old. (What would $\text{Span}_{\mathcal{C}}(A, B)$ even look like as a traditional language?)
- Category theory “eats itself”: the categorical notion of terminal object sheds light on the categorical notion of product, if one knows to look at the right category.

To illustrate the benefits of this perspective, we get a proof that products are suitably unique just like terminal objects are, simply because they *are terminal objects*:

Proposition 3.6.2. Products are unique up to unique isomorphism.

Proof. Products are terminal objects, and terminal objects are unique up to unique isomorphism.⁵ □

⁵ Phew! What a nice short proof!

4

Universal Constructions II: Exponents

Now, let's move on to the next major addition to our library of language features: first-class functions and function types. Let's recall the rules for a language with function types. First, we have the rule for making terms of function type:

$$\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x. M : A \rightarrow B} \text{ (T-LAM)}$$

Next, we have the elimination rule for using terms of function type:

$$\frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash MN : B} \text{ (T-APP)}$$

Finally, we have the β and η laws, which state that applying a function corresponds to substitution and that every term of function type is equivalent to a λ -expression.

$$\frac{\Gamma, x : A \vdash M : B \quad \Gamma \vdash N : A}{\Gamma \vdash (\lambda x. M) N \equiv M[N/x] : B} (\beta \rightarrow)$$

$$\frac{\Gamma \vdash M : A \rightarrow B}{\Gamma \vdash M \equiv \lambda x. Mx : A \rightarrow B} (\eta \rightarrow)$$

Our development here will mirror the development so far of product types and unit types: we will identify a special object in a category, characterized by some morphism equalities and a universal property, such that it validates the equational theory of λ -expressions.

Let's designate a special object in the category to denote the interpretation of function types. This object will be called the **exponential object**:

$$\llbracket A \rightarrow B \rrbracket = \llbracket B \rrbracket^{\llbracket A \rrbracket} \quad (4.1)$$

Intuitively, this exponential object must summarize all the morphisms from A to B . Like product and unit types, this special object will be uniquely characterized by a universal property. Let's work up to what that universal property is. We can start by filling in our semantics for λ -formation, which we know must look like:

$$\llbracket \Gamma \vdash \lambda x : A. M : A \rightarrow B \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket B \rrbracket^{\llbracket A \rrbracket} \quad (4.2)$$

We are gathering clues: the above tells us that there must exist some morphism from $\llbracket \Gamma \rrbracket$ into the special object $\llbracket B \rrbracket^{\llbracket A \rrbracket}$. Let's suggestively call this morphism $\llbracket \Gamma \rrbracket \xrightarrow{\lambda \llbracket M \rrbracket} \llbracket B \rrbracket^{\llbracket A \rrbracket}$:

$$\llbracket \Gamma \vdash \lambda x : A. M : B \rrbracket = \llbracket \Gamma \rrbracket \xrightarrow{\lambda \llbracket M \rrbracket} \llbracket B \rrbracket^{\llbracket A \rrbracket} \quad (4.3)$$

The purpose of the the morphism $\lambda \llbracket M \rrbracket$ is to compute the closure of M : it must package up the body of the lambda M into some representation for which it can later be called. Now, to interpret application, we need another special morphism called app that denotes applying a closure to its argument:

$$\llbracket B \rrbracket^{\llbracket A \rrbracket} \times \llbracket A \rrbracket \xrightarrow{\text{app}} \llbracket B \rrbracket \quad (4.4)$$

With this special morphism we can give an interpretation to application:

$$\left[\frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash MN : B} \right] = \llbracket \Gamma \rrbracket \xrightarrow{\langle \lambda \llbracket M \rrbracket, \text{id} \rangle} \llbracket B \rrbracket^{\llbracket A \rrbracket} \times \llbracket \Gamma \rrbracket \xrightarrow{\langle \pi_1, \llbracket N \rrbracket \circ \pi_2 \rangle} \llbracket B \rrbracket^{\llbracket A \rrbracket} \times \llbracket A \rrbracket \xrightarrow{\text{app}} \llbracket B \rrbracket \quad (4.5)$$

Now, the closure morphism $\lambda \llbracket M \rrbracket$ and the application morphism app need to satisfy the equational laws. We can package up these laws into a diagram that must commute:

$$\begin{array}{ccc} \llbracket \Gamma \rrbracket \times \llbracket A \rrbracket & & \\ \langle \lambda \llbracket M \rrbracket \circ \pi_1, \pi_2 \rangle \downarrow & \searrow \llbracket M \rrbracket & \\ \llbracket B \rrbracket^{\llbracket A \rrbracket} \times \llbracket A \rrbracket & \xrightarrow{\text{app}} & \llbracket B \rrbracket \end{array}$$

Let's unpack this diagram. First, the morphism $\llbracket \Gamma \rrbracket \times \llbracket A \rrbracket \xrightarrow{\llbracket M \rrbracket} \llbracket B \rrbracket$ is the morphism that gives the semantics of the open term $\Gamma, x : A \vdash M : B$. The diagram says that $\llbracket M \rrbracket$ should commute with first forming a closure and then applying the closure to the same argument. Finally, the η laws assert the uniqueness of the closure-formation morphism $\lambda \llbracket M \rrbracket$. Now we can package up all this intuition into a definition of a universal property:

Definition 12: Universal property of exponential object

Let A and B be objects of a category \mathcal{C} with products. An *exponential of B by A* is a pair (B^A, app) consisting of an object B^A of \mathcal{C} and a morphism app from $B^A \times A$ to B such that the following holds: for any object Γ and any morphism $f : \Gamma \times A \rightarrow B$, there exists a unique morphism $\lambda f : \Gamma \rightarrow B^A$ such that $\text{app} \circ \langle \lambda f \circ \pi_1, \pi_2 \rangle = f$.

A category **has exponents** if for any two objects A and B there is an exponential of A by B . Categories that have exponents get a special name: they are called **Cartesian-closed categories**.

4.1 Examples of exponential objects

4.1.1 Exponents in FinSet

Let's begin as usual by seeing examples of universal properties in categories that we're familiar with. We'll begin with the category of finite sets FinSet . For two finite sets A and B , the exponential B^A is an object that represents the set of all functions from A to B . For instance, suppose $A = \{a, b\}$, $B = \{c, d\}$, and let $\Gamma = \{\star\}$. Then, there are 4 possible functions between these two sets:

$$\begin{aligned} f_1(\gamma, x) &= c \\ f_2(\gamma, x) &= \begin{cases} c & \text{if } x = a \\ d & \text{if } x = b \end{cases} \\ f_3(\gamma, x) &= \begin{cases} d & \text{if } x = a \\ c & \text{if } x = b \end{cases} \\ f_4(\gamma, x) &= d \end{aligned}$$

The crucial idea is that this collection of 4 morphisms can itself be represented as a finite set. Recall that the *graph* of a function $f : A \rightarrow B$ is the set of pairs (a, b) where a is the input and b is the output of f . Then, the exponential object B^A is the *finite set of all graphs of functions from A to B* . So, for instance, the set $\{c, d\}^{\{a, b\}}$ is:

$$\left\{ \{(a, c), (b, c)\}, \{(a, d), (b, d)\}, \{(a, c), (b, d)\}, \{(a, d), (b, c)\} \right\}$$

Using this, we can define the closure λf as mapping a function to its graph. For example:

Note that in this example, since $\Gamma = \{\star\}$, the argument γ is ignored in this definition. If Γ were more interesting, it may be the case that different graphs are returned for different contexts.

$$\begin{aligned}\lambda f_1(\gamma) &= \{(a, c), (b, c)\} \\ \lambda f_2(\gamma) &= \{(a, c), (b, d)\} \\ \lambda f_3(\gamma) &= \{(a, d), (b, c)\} \\ \lambda f_3(\gamma) &= \{(a, d), (b, d)\}\end{aligned}$$

Next, an exponential object is paired with a morphism `app` that calls the closure on an argument. We can implement the function `app` by picking the entry of the graph that corresponds to the argument. For example, $\text{app} \circ \langle \lambda f_1, a \rangle = c$.

The exponential object must satisfy the property that $\text{app} \circ \langle \lambda f, A \rangle = f$. This is surely true, because applying the graph of a function to an argument a must agree with calling f on a . Finally, uniqueness follows by function extensionality: any λf that satisfies this property must behave identically to the graph of f . Summing up:

Proposition 4.1.1. `FinSet` is Cartesian-closed.

Proof. We have already shown that `FinSet` has products (the Cartesian product of finite sets). Let A, B be finite sets. Then, we need to show that for any finite set Γ with morphism $\Gamma \times A \xrightarrow{f} B$, there is a unique morphism λf such that the following diagram commutes:

$$\begin{array}{ccc} \Gamma \times A & & \\ \langle \lambda f \circ \pi_1, \text{id} \circ \pi_2 \rangle \downarrow & \searrow f & \\ B^A \times A & \xrightarrow{\text{app}} & B \end{array}$$

To satisfy this, we will define B^A to be set of all graphs of functions $A \rightarrow B$. Then, we define $\text{app}(f, a)$ to pick the component of the graph f that corresponds to $a \in A$.

Now let's show this choice validates the universal property. Let Γ be any set and f be a function $\Gamma \times A \rightarrow B$. Then, there exists a function $\lambda f : \Gamma \rightarrow B^A$ such that the diagram commutes: it's the function that sends each $\gamma \in \Gamma$ to the graph $\{(a, b) \mid f(\gamma, a) = b\}$ representing the "partial application" of f to γ . Observe that, by function extensionality, λf is the unique function making the diagram commute, completing the proof. \square

4.1.2 Exponents in Formula

Recall from Assignment 1 the category `Formula` whose objects are logical formulae whose morphisms denote formula entailment. Formulae are described by the following grammar:

$$\varphi, \psi ::= x \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \neg \psi \tag{4.6}$$

Fix a global finite set of variable names Ω , and assume $x \in \Omega$. Then, we can give a semantics to Boolean formulae in terms of substitutions. A substitution γ is a function $\Omega \rightarrow \{\top, \perp\}$ that maps each variable to a truth value. Then, we can define the semantics of formulae inductively:


$$\begin{aligned} \llbracket x \rrbracket(\gamma) &= \gamma(x) \\ \llbracket \varphi \wedge \psi \rrbracket(\gamma) &= \begin{cases} \top & \text{if } \llbracket \varphi \rrbracket(\gamma) = \top, \llbracket \psi \rrbracket(\gamma) = \top \\ \perp & \text{otherwise} \end{cases} \\ \llbracket \varphi \vee \psi \rrbracket(\gamma) &= \begin{cases} \perp & \text{if } \llbracket \varphi \rrbracket(\gamma) = \perp, \llbracket \psi \rrbracket(\gamma) = \perp \\ \top & \text{otherwise} \end{cases} \\ \llbracket \neg \varphi \rrbracket(\gamma) &= \begin{cases} \top & \text{if } \llbracket \varphi \rrbracket(\gamma) = \perp \\ \perp & \text{if } \llbracket \varphi \rrbracket(\gamma) = \top \end{cases} \end{aligned}$$

The *set of models*, written $\text{Mods}(\varphi)$, is the set of all substitutions that evaluate to \top , i.e. $\text{Mods}(\varphi) = \{\gamma \mid \llbracket \varphi \rrbracket(\gamma) = \top\}$. We say a formula φ *semantically entails* a formula ψ , written $\varphi \models \psi$, if $\text{Mods}(\varphi) \subseteq \text{Mods}(\psi)$. The category **Formula** is the category whose objects are Boolean formula and whose morphisms denote formula entailment:

- Let the set of objects O be the set of all syntactic Boolean formulae generated by the above grammar.
- The set of morphisms M consists of pairs (φ, ψ) where $\varphi \models \psi$.
- Define $\text{dom}(\varphi, \psi) = \varphi, \text{cod}(\varphi, \psi) = \psi$
- Identity is defined by $\text{id}(\varphi) = (\varphi, \varphi)$
- Composition is defined by $\text{comp}((\psi, \xi), (\varphi, \psi)) = (\varphi, \xi)$


In the assignment you showed that this category has products, and that products denoted logical conjunction of formulae. Now, let's investigate what exponential objects mean in this category.

Suppose α and β are logical formulae, and let γ be some other logical formula such that $\gamma \wedge \alpha \models \beta$. Pictorially, we can draw this relationship using a Venn diagram:



Let $\llbracket \alpha \rrbracket$ denote the set of all models of α . To show that $\alpha \wedge \gamma \models \beta$, it must be the case that the $\llbracket \alpha \rrbracket \cap \llbracket \gamma \rrbracket \subseteq \llbracket \beta \rrbracket$. This relationship clearly holds in the above figure.

Viewed from an order-theoretic perspective, the exponential object β^α is the *greatest subset* of Ω such that $\llbracket \alpha \rrbracket \cap \llbracket \beta^\alpha \rrbracket \subseteq \llbracket \beta \rrbracket$. We can draw this subset in green:



Now the question: can this set of models be represented as a logical formula? Indeed it can: it corresponds with the logical formula $\neg\alpha \vee \beta$, also known as the (classical) logical implication $\alpha \Rightarrow \beta$.

So, we have a candidate exponential object. Now, let's check that it satisfies the universal property: we must show that for any formulae α and β , for any formula γ with morphism $\alpha \times \gamma \rightarrow \beta$, there exists an object $\neg\alpha \vee \beta$ such that:

1. There is a morphism $\gamma \rightarrow \neg\alpha \vee \beta$. This morphism represents *logical weakening* of γ .
2. There is a morphism $(\neg\alpha \vee \beta) \times \alpha \rightarrow \beta$. This morphism represents *modus ponens*.

First we prove (1):

Proof. We need to show that if $\alpha \wedge \gamma \models \beta$, then $\gamma \models \neg\alpha \vee \beta$. Suppose by contradiction that $\llbracket \gamma \rrbracket \not\subseteq \llbracket \neg\alpha \vee \beta \rrbracket$. Then, there exists an $x \in \Omega$ such that $x \in \llbracket \gamma \rrbracket$ and $x \notin \llbracket \neg\alpha \vee \beta \rrbracket$. By De Morgan's law we have that $x \in \llbracket \alpha \wedge \neg\beta \rrbracket$. This shows that $x \in \llbracket \gamma \wedge \alpha \rrbracket$ and $x \notin \llbracket \beta \rrbracket$, a contradiction. \square

Point (2) is a well-known fact about propositional logic called modus ponens, which we won't prove here. Since this is a thin category it is not necessary to check uniqueness. This establishes that Formula has exponential objects and hence is a Cartesian-closed category.

4.1.3 An example from topology

Category theory is a bridge between many branches of mathematics, and we can gain significant insight into categorical definitions by viewing them through the lens of different mathematical subfields. Many important developments in category theory stemmed from topology, so we explore some of the basics of topology here.

In the previous section we drew a picture of the exponential as a Venn diagram. The exponential was then the largest shaded region C satisfying the constraint $A \cap C \subseteq B$. This spatial intuition can be made formal.

Definition 13: Open set

A subset U of the real numbers \mathbb{R} is **open** if it is a (possibly infinite) union of open intervals, where an open interval (a, b) is the collection of real numbers between a and b excluding a and b , i.e. $\{r \mid a < r < b\}$.

An equivalent way to define an open set is a subset U of the reals satisfying that, for any $x \in U$, an ϵ -ball around x is also contained in U . I.e., there exists an $\epsilon > 0$ such that $(x - \epsilon, x + \epsilon) \subseteq U$.

The set of all open subsets of \mathbb{R} , written $\mathcal{O}(\mathbb{R})$, forms a partial order under subset inclusion, and hence also a thin category (Construction 2.2.3). As a category, $\mathcal{O}(\mathbb{R})$ has finite products: the terminal object is the largest open subset of \mathbb{R} , namely the whole space \mathbb{R} , and the product of two opens subsets U and V is their intersection $U \cap V$.

Proposition 4.1.2. $\mathcal{O}(\mathbb{R})$ is Cartesian closed.

Proof. Following the intuition that the exponential of two objects U and V is the largest object W such that $U \cap W \subseteq V$, we define

$$V^U = \text{the largest open subset } W \text{ of } \mathbb{R} \text{ such that } U \cap W \subseteq V \quad (4.7)$$

$$= \text{the largest open subset } W \text{ of } \mathbb{R} \text{ contained in } (\mathbb{R} \setminus U) \cup V. \quad (4.8)$$

This largest open subset has a name: it is called the **interior** of $(\mathbb{R} \setminus U) \cup V$, written $\text{int}((\mathbb{R} \setminus U) \cup V)$. Intuitively, taking the interior of a set is a way of “forcing” it to be open, by deleting all of its boundary points. For example, the interior of a closed interval is the corresponding open interval: $\text{int}[0, 1] = (0, 1)$. From this perspective, the subset $\text{int}((\mathbb{R} \setminus U) \cup V)$ defining the exponential V^U is the largest open contained in $(\mathbb{R} \setminus U) \cup V$, which corresponds to the Venn diagram (4.1.2). of $(\mathbb{R} \setminus W) \cup V$. \square

As an example, let us compute the exponential where $U = (-\infty, 0)$ and $V = \emptyset$. In this case,

$$V^U = \text{int}([0, \infty) \cup \emptyset) = \text{int}[0, \infty) = (0, \infty). \quad (4.9)$$

Note how, in the last step, the use of the interior operator “rounds out” the closed interval $[0, \infty)$ into the open interval $(0, \infty)$.

Thinking of the exponential \emptyset^U as a “complement” or “negation” of U , this says that the complement of $(-\infty, 0)$ in $\mathcal{O}(\mathbb{R})$ is $(0, \infty)$.

Note that \mathbb{R} is *not* the union of U and its complement:

$$U \cup \emptyset^U = (-\infty, 0) \cup (0, \infty) = \mathbb{R} \setminus \{0\} \neq \mathbb{R}$$

What this demonstrates is that the law of excluded middle fails in $\mathcal{O}(\mathbb{R})$: the open subsets of \mathbb{R} form a model of intuitionistic propositional logic that refutes $A \vee \neg A = \top$. This is a somewhat surprising bridge between geometry and logic, which we will explore later if we have time.

4.1.4 A non-example: natural numbers and divisibility

It’s worth exploring examples of categories that *do not* have exponential objects. A good example of this is the category (\mathbb{N}, \preceq) , the category of natural numbers with divisibility relations as morphisms.

Let’s see why this is the case. Concretely, we will show that 2 and 3 do not have an exponential.

Suppose that there *was* an exponential k . Then we would have two things.

1. A morphism $k \times 2 \xrightarrow{\text{app}} 3$.
2. For any k' and any morphism $k' \times 2 \rightarrow 3$, there exists a unique morphism $k' \rightarrow k$ such that a triangle commutes.

Unpacking point (1), we have that $\text{gcd}(k, 2)$ divides 3 (because product in this category is greatest common divisor and morphisms are divisibility relations). Unpacking point (2), we have that for any k' such that $\text{gcd}(k', 2)$ divides 3, it holds that k' divides k . The fact that the morphism $k' \rightarrow k$ is the unique one making a given triangle

commute does not mean much here because our category is thin: the moment we know a morphism $k' \rightarrow k$ exists, we also know it must be the unique such morphism.

What do these two conditions tell us about the supposed exponential object k ? Intuitively, they say that k is the “largest number k such that $\gcd(k, 2)$ divides 3”. It is suspicious that there would be a largest such number: $\gcd(k, 2) = \gcd(3k, 2)$ for any k , so if $\gcd(k, 2)$ divides 3 then so does $\gcd(3k, 2)$. In short, we can always keep increasing k by multiplying it by 3, so there cannot be a largest such.

We can turn this intuitive argument into a rigorous one as follows. We will exhibit a counterexample to point (2) above. Set $k' = 3k$. We have that $\gcd(3k, 2) = \gcd(k, 2)$, and $\gcd(k, 2)$ divides 3, so $\gcd(3k, 2)$ divides 3 too. This fulfills the precondition of point (2). The conclusion of point (2) then tells us that $3k$ divides k . But this is impossible.

Since we derived a contradiction starting from an arbitrary choice of k satisfying (1) and (2), this argument shows there can be no k satisfying the universal property of the exponential 3^2 , and hence that (\mathbb{N}, \leq) is not Cartesian closed.

5

From Categories to Languages

Consider a variant of `CALC` from Chapter 1, with a unit type and parameterized by an collection of base types b and constants c .

$$\begin{aligned} M, N ::= & \text{let } x \text{ be } M \text{ in } N \mid x \mid \langle \rangle \mid \langle M, N \rangle \mid \text{fst } M \mid \text{snd } M \\ A, B ::= & A \times B \mid \text{Unit} \mid b \end{aligned} \quad (\text{CALC})$$

The only goal of this chapter is to prove the following.

Theorem 5.0.1. `CALC` can be interpreted in any category with finite products (Definition 9), provided suitable interpretations of each base type b as an object and each constant c as a morphism.

Proof. Let \mathcal{C} be a category with products and a terminal object. We will build an interpretation function $\llbracket - \rrbracket$ that maps each of the judgments of `CALC` into \mathcal{C} :

- Types A will become objects $\llbracket A \rrbracket$ of \mathcal{C} .
- Contexts Γ will become objects $\llbracket \Gamma \rrbracket$ of \mathcal{C} .
- Derivations of $\Gamma \vdash M : A$ will become morphisms $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket M \rrbracket} \llbracket A \rrbracket$.
- Derivations of $\Gamma \vdash M \equiv N : A$ will become equalities:

$$\begin{array}{ccc} & \llbracket M \rrbracket & \\ & \curvearrowright & \\ \llbracket \Gamma \rrbracket & \parallel & \llbracket A \rrbracket \\ & \curvearrowleft & \\ & \llbracket N \rrbracket & \end{array}$$

Each of these will be defined by induction.¹

First, let us interpret `CALC` types. We assume given the interpretation of base types, with each base type b denoting some chosen object $\llbracket b \rrbracket$ of \mathcal{C} . This interpretation is bootstrapped into one for all types by induction:

$$\llbracket \text{Unit} \rrbracket = 1 \quad (5.1)$$

$$\llbracket A \times B \rrbracket = \llbracket A \rrbracket \times \llbracket B \rrbracket \quad (5.2)$$

¹ Warning: this is going to take a while.

These equations look just like (1.1), the interpretation of types given in Section 1.1. But their meaning is very different. Whereas before the operator \times was used to denote the Cartesian product of finite sets, here it denotes the categorical product of the objects $\llbracket A \rrbracket$ and $\llbracket B \rrbracket$ of an arbitrary category \mathcal{C} . Similarly, 1 denotes the terminal object in \mathcal{C} , not the singleton set $\{\star\}$.

The interpretation of `CALC` contexts is similarly straightforward. The empty context is interpreted as the terminal object and context extension is interpreted via product—again, we are using the same symbols for the abstract categorical definitions rather than the set-theoretic ones.

$$\llbracket \bullet \rrbracket = 1 \quad (5.3)$$

$$\llbracket \Gamma, x : A \rrbracket = \llbracket \Gamma \rrbracket \times \llbracket A \rrbracket \quad (5.4)$$

Next up we have the interpretation of typing derivations. This is where things start getting interesting. We have one case per typing rule. As a warm-up, here is the interpretation of the typing rule for `⟨⟩`.

$$\left\llbracket \frac{}{\Gamma \vdash \langle \rangle : \mathbf{Unit}} \right\rrbracket = \langle \rangle_{\llbracket \Gamma \rrbracket} \quad (5.5)$$

In words: the interpretation of `⟨⟩` with respect to typing context Γ is the unique morphism $\langle \rangle_{\llbracket \Gamma \rrbracket}$ from $\llbracket \Gamma \rrbracket$ to $\llbracket \langle \rangle \rrbracket$, guaranteed to exist because $\llbracket \langle \rangle \rrbracket$ is a terminal object of \mathcal{C} .

The interpretation of `⟨⟩` is special because its typing rule has no premises. More generally, the interpretations of the other typing rules will rely inductively on interpretations of premises. For instance, here is the interpretation of the typing rule for `fst M`.

$$\left\llbracket \frac{\begin{array}{c} \vdots \\ \Gamma \vdash M : A \times B \end{array}}{\Gamma \vdash \mathbf{fst} M : A} \right\rrbracket = \pi_1 \circ \left\llbracket \frac{\begin{array}{c} \vdots \\ \Gamma \vdash M : A \times B \end{array}}{} \right\rrbracket \quad (5.6)$$

² In words: to interpret a typing derivation for `fst M` as shown, first interpret the subderivation establishing $\Gamma \vdash M : A \times B$ to obtain a morphism $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket M \rrbracket} \llbracket A \times B \rrbracket$, and then compose this morphism with the projection $\llbracket A \times B \rrbracket \xrightarrow{\pi_1} \llbracket A \rrbracket$, guaranteed to exist because $\llbracket A \times B \rrbracket$ is a product of $\llbracket A \rrbracket$ and $\llbracket B \rrbracket$.

The typing rule for `snd M` is interpreted analogously.

$$\left\llbracket \frac{\begin{array}{c} \vdots \\ \Gamma \vdash M : A \times B \end{array}}{\Gamma \vdash \mathbf{snd} M : B} \right\rrbracket = \pi_2 \circ \left\llbracket \frac{\begin{array}{c} \vdots \\ \Gamma \vdash M : A \times B \end{array}}{} \right\rrbracket \quad (5.7)$$

² **todo:** Really wish I knew how to fix the spacing here.

The pair formation rule is interpreted using the universal property of product.

$$\left[\frac{\frac{\vdots}{\Gamma \vdash M : A} \quad \frac{\vdots}{\Gamma \vdash N : B}}{\Gamma \vdash \langle M, N \rangle : A \times B} \right] = \left\langle \left[\frac{\vdots}{\Gamma \vdash M : A} \right], \left[\frac{\vdots}{\Gamma \vdash N : B} \right] \right\rangle \quad (5.8)$$

Finally we have the typing rules for variables and let-bindings. We have saved these rules for last because they highlight some interesting points regarding the categorical interpretation of operations on the typing context Γ .

First let us see how to interpret variables. Because typing contexts Γ are interpreted as iterated products, variables are interpreted as projections:

$$\left[\frac{(x : A) \in \Gamma}{\Gamma \vdash x : A} \right] = \pi_x \quad (5.9)$$

In words: π_x is the canonical projection $[\Gamma] \rightarrow [A]$ that extracts the “ x th component” out of the iterated product used to define $[\Gamma]$.

You may find this interpretation of variables unsatisfying: what exactly is this “canonical” projection, and what does it mean to take the “ x th component”? Here is one way to make this precise. The trick is to reformulate the typing rule for variables in terms of two more elementary rules:

$$\frac{}{\Gamma, x : A \vdash x : A} \text{ HIT} \qquad \frac{\Gamma \vdash x : A}{\Gamma, y : B \vdash x : A} \text{ MISS}$$

Together these rules can be thought of as defining a procedure for checking whether a given binding $(x : A)$ lives in a given typing context Γ . The rule HIT covers the case where $(x : A)$ matches the variable right at the end of the typing context. The rule MISS covers the case where it doesn’t and one has to keep looking further into Γ .

This reformulation of the variable rule allows for a precise definition of the handwavy π_x .

$$\left[\frac{}{\Gamma, x : A \vdash x : A} \text{ (HIT)} \right] = \pi_2 \quad (5.10)$$

$$\left[\frac{\frac{\vdots}{\Gamma \vdash x : A}}{\Gamma, y : B \vdash x : A} \text{ (MISS)} \right] = \left[\frac{\vdots}{\Gamma \vdash x : A} \right] \circ \pi_1 \quad (5.11)$$

Using these two rules, we can see that π_x is actually a composite consisting of a string of π_1 s followed by a π_2 , with the length of this string depending on where x appears in Γ . For instance,

$$\left[\frac{\vdots}{x : A, y : B, z : C \vdash x : A} \right] \quad (5.12)$$

$$= \left[\frac{\vdots}{x : A, y : B \vdash x : A} \right] \circ \pi_1 \quad (5.13)$$

$$= \left[\frac{\vdots}{x : A \vdash x : A} \right] \circ \pi_1 \circ \pi_1 \quad (5.14)$$

$$= \pi_2 \circ \pi_1 \circ \pi_1. \quad (5.15)$$

Note that there are two π_1 s in this composite. This corresponds to the fact that the de Bruijn index of x in the context $x : A, y : B, z : C$ is two.³

With variables out of the way, let us now turn to let-bindings.

These also illustrate an interesting interaction between operations on the typing context and categorical products.

³ You are now a third of the way through the proof of Theorem 5.0.1.

$$\left[\frac{\frac{\vdots}{\Gamma \vdash M : A} \quad \frac{\vdots}{\Gamma, x : A \vdash N : B}}{\Gamma \vdash \text{let } x \text{ be } M \text{ in } N : B} \right] \\ = \left[\frac{\vdots}{\Gamma, x : A \vdash N : B} \right] \circ \left\langle \text{id}_{[\Gamma]}, \left[\frac{\vdots}{\Gamma \vdash M : A} \right] \right\rangle$$

Let's break this down a bit. By induction, we have that M denotes a morphism $[\Gamma] \rightarrow [A]$ and N denotes a morphism $[\Gamma, x : A] = [\Gamma] \times [A] \rightarrow [B]$. The interpretation of `let x be M in N` composes these together. The twist is that, because the variables in Γ are shared between M and N , the object $[\Gamma]$ must be plumbed around as well. This is accomplished by the morphism $\langle \text{id}_{[\Gamma]}, [\Gamma \vdash M : A] \rangle$ guaranteed to exist by the universal property of products.

In practice, it's fairly laborious to write down the interpretation function explicitly as a function on derivations as we have done

above. So people tend to shorten the more verbose $\left[\frac{\vdots}{\Gamma \vdash M : A} \right]$ to just $[[M]]$. Written this way, the foregoing discussion can be sum-

marized in terms of the following equations.

$$\begin{aligned}
\llbracket \langle \rangle \rrbracket &= \langle \rangle \\
\llbracket \text{fst } M \rrbracket &= \pi_1 \circ \llbracket M \rrbracket \\
\llbracket \text{snd } M \rrbracket &= \pi_2 \circ \llbracket M \rrbracket \\
\llbracket \langle M, N \rangle \rrbracket &= \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle \\
\llbracket x \rrbracket &= \pi_x \\
\llbracket \text{let } x \text{ be } M \text{ in } N \rrbracket &= \llbracket N \rrbracket \circ \langle \text{id}_{\llbracket \Gamma \rrbracket}, \llbracket M \rrbracket \rangle
\end{aligned}$$

This kind of equational definition is what you will typically encounter in a paper that uses categorical semantics.⁴

We have now seen how to interpret the types, typing contexts, and typing derivations of CALC. Only one task remains: interpret derivations of $\Gamma \vdash M \equiv N : A$. There are five groups of rules to interpret.

1. First there are the rules stating that \equiv is an equivalence relation on well-typed terms.

$$\begin{array}{c}
\frac{\Gamma \vdash M : A}{\Gamma \vdash M \equiv M : A} \text{REFL} \qquad \frac{\Gamma \vdash M \equiv N : A}{\Gamma \vdash N \equiv M : A} \text{SYM} \\
\\
\frac{\Gamma \vdash M \equiv N : A \quad \Gamma \vdash N \equiv O : A}{\Gamma \vdash M \equiv O : A} \text{TRANS}
\end{array}$$

2. Next there are the *congruence rules*, so called because they state that \equiv is a congruence for each of the CALC term formers.⁵

$$\begin{array}{c}
\frac{\Gamma \vdash M \equiv M' : A \quad \Gamma \vdash N \equiv N' : B}{\Gamma \vdash \langle M, N \rangle \equiv \langle M', N' \rangle : A \times B} \text{CONG-PAIR} \\
\\
\frac{\Gamma \vdash M \equiv M' : A \times B}{\Gamma \vdash \text{fst } M \equiv \text{fst } M' : A} \text{CONG-FST} \\
\\
\frac{\Gamma \vdash M \equiv M' : A \times B}{\Gamma \vdash \text{snd } M \equiv \text{snd } M' : A} \text{CONG-SND} \\
\\
\frac{\Gamma \vdash M \equiv M' : A \quad \Gamma, x : A \vdash N \equiv N' : B}{\Gamma \vdash (\text{let } x \text{ be } M \text{ in } N) \equiv (\text{let } x \text{ be } M' \text{ in } N') : B} \text{CONG-LET}
\end{array}$$

3. Then there are the β laws that describe what happens when one applies an introduction form for a given type former followed by an elimination form.

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \text{fst } \langle M, N \rangle \equiv M : A} \beta_1^\times \qquad \frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \text{fst } \langle M, N \rangle \equiv N : B} \beta_2^\times$$

⁴ In some cases it is desirable to have an interpretation function that operates directly on terms, rather than on derivations, as derivations can contain extraneous information not present in a term. In these cases a tricky issue comes up: for $\llbracket - \rrbracket$ to be a function on terms, it must be the case that the denotation of a term does not depend on the precise way in which it is proved to be well-typed. Checking this *coherence condition* involves showing that, for any term M , the denotations of all typing derivations for M are equal [33]. In our case CALC satisfies the property that every term has at most one typing derivation when each bound variable is annotated with its type, so we won't worry too much about this.

⁵ In general, an equivalence relation \sim on a set X is a *congruence* with respect to a function $f : X \rightarrow X$ if $x \sim x'$ implies $f(x) \sim f(x')$.

4. Then there are the η laws that describe what happens when one applies an elimination form followed by an introduction form.

$$\frac{\Gamma \vdash M : A \times B}{\Gamma \vdash \langle \text{fst } M, \text{snd } M \rangle \equiv M : A \times B} \eta^\times \quad \frac{\Gamma \vdash M : \text{Unit}}{\Gamma \vdash \langle \rangle \equiv M : \text{Unit}} \eta^{\text{Unit}}$$

5. Last but not least, there is the law that let-binding is equivalent to substitution.

$$\frac{\Gamma \vdash M : A \quad \Gamma, x : A \vdash N : B}{\Gamma \vdash (\text{let } x \text{ be } M \text{ in } N) \equiv N[M/x] : B} \text{LET-SUB}$$

Let us tackle each of these groups one by one.

The rules in group (1) are easy to interpret: they correspond directly to the reflexivity, symmetry, and transitivity of equality of morphisms. For instance, the `SYM` rule holds because, if $\llbracket M \rrbracket = \llbracket N \rrbracket$ for some well-typed terms M and N , then also $\llbracket N \rrbracket = \llbracket M \rrbracket$.⁶

The rules in group (2) are also relatively easy. In each case the proof boils down to showing that a given operation on morphisms respects equality. For instance, the rule `CONG-PAIR` boils down to the following true statement: if $\llbracket M \rrbracket = \llbracket M' \rrbracket$ and $\llbracket N \rrbracket = \llbracket N' \rrbracket$ then $\langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle = \langle \llbracket M' \rrbracket, \llbracket N' \rrbracket \rangle$. The other cases are similarly straightforward.

The rules in group (3) follow from the definition of categorical product. The law β_1^\times boils down to showing that $\pi_1 \circ \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle = \llbracket M \rrbracket$. The law β_2^\times is analogous.

The rules in group (4) follow from the definitions of categorical product and terminal object. First, the law η^{Unit} follows from the fact that any two morphisms into a terminal object are equal. The law η^\times boils down to showing that, for any morphism $\llbracket M \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket \times \llbracket B \rrbracket$, it is the case that $\llbracket M \rrbracket = \langle \pi_1 \circ \llbracket M \rrbracket, \pi_2 \circ \llbracket M \rrbracket \rangle$. This follows from a uniqueness argument. By definition, $\langle \pi_1 \circ \llbracket M \rrbracket, \pi_2 \circ \llbracket M \rrbracket \rangle$ is the unique morphism f satisfying $\pi_1 \circ f = \pi_1 \circ \llbracket M \rrbracket$ and $\pi_2 \circ f = \pi_2 \circ \llbracket M \rrbracket$. But $\llbracket M \rrbracket$ satisfies this same property. Hence $\llbracket M \rrbracket = f = \langle \pi_1 \circ \llbracket M \rrbracket, \pi_2 \circ \llbracket M \rrbracket \rangle$.

All that's left is `LET-SUB`—the rule stating that let-binding is equivalent to substitution. Validating this rule boils down to showing the following:

Lemma 5.0.2 (“Let is substitution”). For all contexts Γ , types A , B and terms M and N such that $\Gamma \vdash M : A$ and $\Gamma, x : A \vdash N : B$, it holds that $\llbracket N[M/x] \rrbracket = \llbracket N \rrbracket \circ \langle \text{id}_{\llbracket \Gamma \rrbracket}, \llbracket M \rrbracket \rangle$.

This lemma essentially relates the syntactic operation of substituting of M into N with the semantic operation of composing the morphism $\llbracket N \rrbracket$ with $\langle \text{id}_{\llbracket \Gamma \rrbracket}, \llbracket M \rrbracket \rangle$.

⁶ The stickler may wonder how we know that M and N are well-typed here, since all we have in the premise of `SYM` is that $M \equiv N$. It turns out that \equiv satisfies the following property: if $\Gamma \vdash M \equiv N : A$, then $\Gamma \vdash M : A$ and $\Gamma \vdash N : A$. This fact can be established by induction on derivations of $\Gamma \vdash M \equiv N : A$, and we won't belabor it. A slicker way of maintaining this well-typedness invariant throughout the definition of \equiv is to make well-typedness a *presupposition* of the judgment $\Gamma \vdash M \equiv N : A$ [27].

A natural proof strategy to try here is to proceed by induction on (a typing derivation for) N , as the substitution operation $N[M/x]$ is defined recursively on it. But this approach runs into a problem: the induction hypothesis one gets is not strong enough. The problematic case is when $N = (\text{let } x' \text{ be } N_1 \text{ in } N_2)$. In this case, we have that

1. $\Gamma, x : A \vdash N_1 : A'$ for some type A'
2. $\Gamma, x : A, x' : A' \vdash N_2 : B$

and the goal is to show that

$$\llbracket (\text{let } x' \text{ be } N_1 \text{ in } N_2)[M/x] \rrbracket = \llbracket \text{let } x' \text{ be } N_1 \text{ in } N_2 \rrbracket \circ \langle \text{id}_{\llbracket \Gamma \rrbracket}, \llbracket M \rrbracket \rangle. \quad (5.16)$$

By the definition of substitution, we have that

$$(\text{let } x' \text{ be } N_1 \text{ in } N_2)[M/x] = (\text{let } x' \text{ be } N_1[M/x] \text{ in } N_2[M/x]).$$

At first glance, this equation looks awfully provable from the above assumptions,⁷ by applying the induction hypothesis to the terms $N_1[M/x]$ and $N_2[M/x]$ and then performing some algebraic manipulations. But a closer look will reveal that $N_2[M/x]$ actually does not fit the shape required for the induction hypothesis to apply. This is because our induction hypothesis only applies to substitutions of M into terms well-typed in context $\Gamma, x : A$. On the other hand, $N_2[M/x]$ is a substitution of $\Gamma \vdash M : A$ into $\Gamma, x : A, x' : A' \vdash N_2 : B$. This extra red entry in the typing context does not match the shape of our induction hypothesis.

⁷ You are now two thirds of the way through the proof of Theorem 5.0.1.

What has happened here is that the statement of Lemma 5.0.2 applies only to terms of the form $N[M/x]$ where x is the right-most variable in the typing context used to type N . But while performing the substitution $N[M/x]$, we will recursively substitute M into subterms where x is *not* the right-most variable anymore, due to the presence of let-bindings. Thus there is a mismatch between the recursive structure of the substitution function and the shape of our inductive argument.

We can remedy this by strengthening Lemma 5.0.2 to the following statement, which applies to terms $N[M/x]$ where x may occur anywhere in the context.

Lemma 5.0.3 (Strengthening 5.0.2). For all contexts Γ and Δ , types A , and terms M and N such that $\Gamma \vdash M : A$ and $\Gamma, x : A, \Delta \vdash N : B$, it holds that $\llbracket N[M/x] \rrbracket = \llbracket N \rrbracket \circ \langle \pi_{\llbracket \Gamma \rrbracket}, \llbracket M \rrbracket \circ \pi_{\llbracket \Gamma \rrbracket}, \pi_{\llbracket \Delta \rrbracket} \rangle$.

The statement of Lemma 5.0.3 requires some clarification.

First, it involves morphisms $\llbracket \Gamma, \Delta \rrbracket \xrightarrow{\pi_{\llbracket \Gamma \rrbracket}} \llbracket \Gamma \rrbracket$ and $\llbracket \Gamma, \Delta \rrbracket \xrightarrow{\pi_{\llbracket \Delta \rrbracket}} \llbracket \Delta \rrbracket$. These morphisms are defined by the following fact, which can be

proved by induction on Δ : for all Γ and Δ it holds that $[[\Gamma, \Delta]]$ is a product of $[[\Gamma]]$ and $[[\Delta]]$. The morphisms $\pi_{[[\Gamma]]}$ and $\pi_{[[\Delta]]}$ are the projections out of this product.

Second, we have used a ternary angle-bracket operator to form the morphism $\langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle$. This notation is justified by the following fact, which also can be proved by induction on Δ : for all Γ and A and Δ it holds that $[[\Gamma, x : A, \Delta]]$ is a ternary product of $[[\Gamma]]$, $[[A]]$, and $[[\Delta]]$. The ternary angle-bracket then denotes the unique morphism given by the universal property of this ternary product.

We are now ready to prove Lemma 5.0.3.

Proof of Lemma 5.0.3. By induction on the derivation of $\Gamma, x : A, \Delta \vdash N : A$. We start with the trickiest case:

- Suppose $N = (\text{let } x' \text{ be } N_1 \text{ in } N_2)$ for some $\Gamma, x : A, \Delta \vdash N_1 : A'$ and some $\Gamma, x : A, \Delta, x' : A' \vdash N_2 : B$. Then both the left- and right-hand sides of the desired equation simplify to the same expression. On the one hand we have

$$\begin{aligned}
& [[(\text{let } x' \text{ be } N_1 \text{ in } N_2)[M/x]] \\
&= [[\text{let } x' \text{ be } N_1[M/x] \text{ in } N_2[M/x]]] \\
&= [[N_2[M/x]]] \circ \langle \pi_{[[\Gamma, x : A, \Delta]]}, [[N_1[M/x]]] \rangle \\
&\stackrel{\text{IH}}{=} ([[N_2]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta, x' : A']]} \rangle) \circ \langle \text{id}_{[[\Gamma, x : A, \Delta]]}, [[N_1[M/x]]] \rangle \\
&\stackrel{\text{IH}}{=} ([[N_2]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta, x' : A']]} \rangle) \circ p \\
&\quad \text{where } p = \langle \text{id}_{[[\Gamma, x : A, \Delta]]}, [[N_1]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle \rangle \\
&= ([[N_2]] \circ \langle \pi_{[[\Gamma]]} \circ p, [[M]] \circ \pi_{[[\Gamma]]} \circ p, \pi_{[[\Delta, x' : A']]} \circ p \rangle) \\
&= ([[N_2]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta, x' : A']]} \circ p \rangle) \\
&= ([[N_2]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \langle \pi_{[[\Delta]]}, [[N_1]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle \rangle \rangle)
\end{aligned}$$

and on the other hand we have

$$\begin{aligned}
& [[\text{let } x' \text{ be } N_1 \text{ in } N_2]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle \\
&= [[N_2]] \circ \langle \text{id}_{[[\Gamma, x : A, \Delta]]}, [[N_1]] \rangle \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle \\
&= [[N_2]] \circ \langle \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle, [[N_1]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle \rangle.
\end{aligned}$$

The equality of these two morphisms boils down to the equality of

$$\langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \langle \pi_{[[\Delta]]}, [[N_1]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle \rangle \rangle$$

and

$$\langle \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle, [[N_1]] \circ \langle \pi_{[[\Gamma]]}, [[M]] \circ \pi_{[[\Gamma]]}, \pi_{[[\Delta]]} \rangle \rangle.$$

This equality follows from a uniqueness argument: both of these morphisms are the unique $f : [[\Gamma, \Delta]] \rightarrow [[\Gamma, x : A, \Delta, x' : A']]$ satisfying the following four properties:

1. $\pi_{[\Gamma]} \circ f = \pi_{[\Gamma]}$
2. $\pi_{[A]} \circ f = \llbracket M \rrbracket \circ \pi_{[\Gamma]}$
3. $\pi_{[\Delta]} \circ f = \llbracket M \rrbracket \circ \pi_{[\Delta]}$
4. $\pi_{[A']} \circ f = \llbracket N_1 \rrbracket \circ \langle \pi_{[\Gamma]}, \llbracket M \rrbracket \circ \pi_{[\Gamma]}, \pi_{[\Delta]} \rangle$

This completes the proof of the case $N = (\text{let } x' \text{ be } N_1 \text{ in } N_2)$.

The other cases are much easier.

- Suppose $N = \langle \rangle$. Then

$$\begin{aligned} \llbracket \langle \rangle [M/x] \rrbracket &= \llbracket \langle \rangle \rrbracket = \langle \rangle = \langle \rangle \circ \langle \pi_{[\Gamma]}, \llbracket M \rrbracket \circ \pi_{[\Gamma]}, \pi_{[\Delta]} \rangle \\ &= \llbracket \langle \rangle \rrbracket \circ \langle \pi_{[\Gamma]}, \llbracket M \rrbracket \circ \pi_{[\Gamma]}, \pi_{[\Delta]} \rangle. \end{aligned}$$

- Suppose $N = \text{fst } N'$ for some $\Gamma, x : A, \Delta \vdash \text{fst } N'$. Then

$$\begin{aligned} \llbracket (\text{fst } N') [M/x] \rrbracket &= \llbracket \text{fst } (N' [M/x]) \rrbracket = \pi_1 \circ \llbracket N' [M/x] \rrbracket \\ &\stackrel{\text{IH}}{=} \pi_1 \circ \llbracket N' \rrbracket \circ \langle \pi_{[\Gamma]}, \llbracket M \rrbracket \circ \pi_{[\Gamma]}, \pi_{[\Delta]} \rangle \\ &= \llbracket \text{fst } N' \rrbracket \circ \langle \pi_{[\Gamma]}, \llbracket M \rrbracket \circ \pi_{[\Gamma]}, \pi_{[\Delta]} \rangle. \end{aligned}$$

- Case $N = \text{snd } N'$ is similar.
- Suppose $N = \langle N_1, N_2 \rangle$ for some B, N_1, N_2 such that $B = B_1 \times B_2$ and $\Gamma, x : A, \Delta \vdash N_1 : B_1$ and $\Gamma, x : A, \Delta \vdash N_2 : B_2$. Then

$$\begin{aligned} \llbracket \langle N_1, N_2 \rangle [M/x] \rrbracket &= \llbracket \langle N_1 [M/x], N_2 [M/x] \rangle \rrbracket \\ &= \langle \llbracket N_1 [M/x] \rrbracket, \llbracket N_2 [M/x] \rrbracket \rangle \\ &\stackrel{\text{IH}}{=} \langle \llbracket N_1 \rrbracket \circ p, \llbracket N_2 \rrbracket \circ p \rangle \text{ where } p = \langle \pi_{[\Gamma]}, \llbracket \Gamma \vdash M : A \rrbracket \circ \pi_{[\Gamma]}, \pi_{[\Delta]} \rangle \\ &= \langle \llbracket N_1 \rrbracket, \llbracket N_2 \rrbracket \rangle \circ p \\ &= \llbracket \langle N_1, N_2 \rangle \rrbracket \circ p \end{aligned}$$

- Suppose $N = x$. For this case, it will be helpful to disambiguate the ambient typing context under which a term is interpreted:

$$\begin{aligned} &\llbracket \Gamma, \Delta \vdash x[M/x] : A \rrbracket \\ &= \llbracket \Gamma, \Delta \vdash M : A \rrbracket \\ &\stackrel{(*)}{=} \llbracket \Gamma \vdash M : A \rrbracket \circ \pi_{[\Gamma]} \\ &= \pi_x \circ \langle \pi_{[\Gamma]}, \llbracket M \rrbracket \circ \pi_{[\Gamma]}, \pi_{[\Delta]} \rangle \end{aligned}$$

The step marked $(*)$ can be proved as a separate lemma, by induction on typing derivations.

- Suppose $N = x'$ for some $x' \neq x$ such that $\Gamma, x : A, \Delta \vdash x' : A'$.
Then

$$\begin{aligned} & \llbracket \Gamma, \Delta \vdash x' [M/x] : A' \rrbracket \\ &= \llbracket \Gamma, \Delta \vdash x' : A' \rrbracket \\ &= \pi_{x'} \\ &= \llbracket \Gamma, x : A, \Delta \vdash x' : A' \rrbracket \circ \langle \pi_{\llbracket \Gamma \rrbracket}, \llbracket M \rrbracket \circ \pi_{\llbracket \Gamma \rrbracket}, \pi_{\llbracket \Delta \rrbracket} \rangle. \end{aligned}$$

This closes the induction, completing the proof of Lemma 5.0.3. \square

Lemma 5.0.2 now follows as a corollary, by setting $\Delta = \bullet$. This in turn establishes the validity of the equational rule LET-SUB, which completes the interpretation of all of the equational axioms of CALC and hence concludes the proof of Theorem 5.0.1.⁸ \square

⁸ You heard that right—Theorem 5.0.1 is proved. You are free! Well done if you made it this far.

5.1 Working in internal languages

Whew! That was quite a long proof. You may be wondering: what have we gained by spending so much energy on proving Theorem 5.0.1? From the PL perspective, Theorem 5.0.1 is not super exciting—CALC is a toy programming language. But Theorem 5.0.1 is extremely helpful for doing category theory. This is because it provides a convenient language for working with an arbitrary category that has finite products.

Let's see an example of this. Recall the divisor lattice of natural numbers (\mathbb{N}, \leq) extended with a terminal object (the number 0). We showed previously that this category has products, and that these products can be interpreted as greatest common divisors. So, Theorem 5.0.1 applies to this category, and we can use it to prove the following seemingly non-trivial fact about greatest common divisors:

Proposition 5.1.1. Let X, Y, Z be natural numbers. Then, $\gcd(\gcd(X, Y), Z)$ divides $\gcd(Y, Z)$.

Proof. Recall that in this category, the $\gcd(X, Y)$ is the categorical product $X \times Y$, and divisibility corresponds to the existence of a morphism. So, the proof strategy is to use Theorem 5.0.1 to prove the existence of a morphism of the correct type.

First, we need a way in CALC of discussing arbitrary objects in a category. We do this by extending our grammar of CALC types with type formers for three arbitrary objects in the category:

$$A ::= \dots \mid X \mid Y \mid Z$$

Then, we define the denotation of these three new types to be $\llbracket X \rrbracket = X, \llbracket Y \rrbracket = Y, \llbracket Z \rrbracket = Z$. Now, the proof amounts to showing that

there exists a morphism $(X \times Y) \times Z \rightarrow (Y, Z)$ using Theorem 5.0.1, which means we have to give a term of type $(X \times Y) \times Z \rightarrow Y \times Z$. Such a term is quite easy to find if you've ever programmed in a functional programming language:

$$p : (X \times Y) \times Z \vdash \langle \text{fst}(\text{snd } p), \text{snd } p \rangle : Y \times Z$$

This completes the proof. \square

Intuitively, Theorem 5.0.1 gives us license to reason about categorical products like how we are familiar reasoning about them in functional programs: we can project out of them, compose them with other projects, etc. A more elaborate application of this theorem is to establish equalities of certain morphisms using the equational theory of `CALC`:

Proposition 5.1.2. Let \mathcal{C} be a category with products. Then $X \times Y \cong Y \times X$ for any objects X, Y of \mathcal{C} .

Proof. The following argument will hold for any choice of objects X and Y in the category \mathcal{C} . For each such pair, extend the grammar of `CALC` types with new base types:

$$A ::= \dots \mid X \mid Y.$$

Interpret these new types as the objects X and Y of the category \mathcal{C} :

$$\begin{aligned} \llbracket X \rrbracket &= X \\ \llbracket Y \rrbracket &= Y \end{aligned}$$

Now define $p : X \times Y \vdash M : Y \times X$ by $M := \langle \text{snd } p, \text{fst } p \rangle$ and similarly define $q : Y \times X \vdash N : X \times Y$ by $N := \langle \text{snd } q, \text{fst } q \rangle$. By Theorem 5.0.1, we have morphisms $\llbracket M \rrbracket : X \times Y \rightarrow Y \times X$ and $\llbracket N \rrbracket : Y \times X \rightarrow X \times Y$. Simple equational reasoning then shows that

$$p : X \times Y \vdash N[M/q] \equiv p : X \times Y \quad (5.17)$$

$$q : Y \times X \vdash M[N/p] \equiv q : Y \times X \quad (5.18)$$

which by Theorem 5.0.1 implies $\llbracket N \rrbracket \circ \llbracket M \rrbracket = \text{id}_{X \times Y}$ and $\llbracket M \rrbracket \circ \llbracket N \rrbracket = \text{id}_{Y \times X}$ as needed to show $X \times Y \cong Y \times X$.⁹ \square

This style of proof is broadly referred to as *working in the internal language of a category*. You may have heard of other internal languages for categories; perhaps the most well-known are string diagrams, which we may return to later.

In the proof of Proposition 5.1.2, we extended `CALC` with new base types X and Y to model the objects X and Y of the category \mathcal{C} . We could give `CALC` access to other objects in \mathcal{C} by extending it further.

⁹ Those that are familiar with the Curry-Howard correspondence may notice something here: equalities in the internal language are proofs of equalities in the model category, so proving equalities in the model category feels like manipulating proof terms.

Taking this idea to its extreme, we could extend CALC with *all* objects in \mathcal{C} , and also all morphisms. This extension to CALC is known as the **internal language** of \mathcal{C} , and programming in it is known as “working internal to \mathcal{C} ”.

Definition 14: Internal language for finite products

Let \mathcal{C} be a category with finite products. The **internal language** of \mathcal{C} is the syntactic extension to CALC that adds the following:

- One primitive type X for each object X of \mathcal{C} .
- One term former $f(M)$ for each morphism $X \xrightarrow{f} Y$, with typing rule:

$$\frac{\Gamma \vdash M : X}{\Gamma \vdash f(M) : Y}$$

- For each object X of \mathcal{C} , an equational law:

$$\frac{\Gamma \vdash M : X}{\Gamma \vdash \text{id}(M) \equiv M : X}$$

Here id denotes the term former corresponding to id_X .

- For all morphisms $f : Y \rightarrow Z$ and $g : X \rightarrow Y$ and $h : X \rightarrow Z$ in \mathcal{C} such that $f \circ g = h$, an equational law:

$$\frac{\Gamma \vdash M : X}{\Gamma \vdash f(g(M)) \equiv h(M) : Z}$$

Proposition 5.1.3. Any category \mathcal{C} with finite products is a model of its own internal language in the evident way.

Internal languages are what formally make categories behave like “metalanguages” for PL. Metalanguage features correspond to additional structure: as we progress through the course, we will encounter universal constructions that model richer languages with features such as records, pattern matching and structural recursion, higher-order functions, and refinement types.

5.2 CCCs are models of the simply-typed lambda calculus

Cartesian closed categories give us all the structure that we need to interpret simply-typed lambda calculus programs. Recall the lan-

guage STLC:

$$\begin{aligned}
 A, B &::= \mathbf{Unit} \mid A \times B \mid A \rightarrow B \\
 M, N &::= \langle \rangle \mid \langle M, N \rangle \mid \mathbf{fst} M \mid \mathbf{snd} M \mid \lambda x. M \mid M N \mid x
 \end{aligned}
 \tag{STLC}$$

We give the following interpretation of STLC in a Cartesian closed category \mathcal{C} . First we interpret types:

$$\begin{aligned}
 \llbracket \mathbf{Unit} \rrbracket &= \mathbf{1} \\
 \llbracket A \times B \rrbracket &= \llbracket A \rrbracket \times \llbracket B \rrbracket \\
 \llbracket A \rightarrow B \rrbracket &= \llbracket B \rrbracket^{\llbracket A \rrbracket}
 \end{aligned}$$

Then we interpret terms as morphisms:

$$\begin{aligned}
 \llbracket \langle \rangle \rrbracket &= \langle \rangle \\
 \llbracket \mathbf{fst} M \rrbracket &= \pi_1 \circ \llbracket M \rrbracket \\
 \llbracket \mathbf{snd} M \rrbracket &= \pi_2 \circ \llbracket M \rrbracket \\
 \llbracket \langle M, N \rangle \rrbracket &= \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle \\
 \llbracket x \rrbracket &= \pi_x \\
 \llbracket \lambda x. M \rrbracket &= \lambda \llbracket M \rrbracket \\
 \llbracket M N \rrbracket &= \mathbf{app} \circ \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle
 \end{aligned}$$

These morphisms satisfy the equational theory of the simply-typed lambda calculus:

Theorem 5.2.1. Every Cartesian closed category is a model of the simply-typed lambda calculus.

We'll spare you the details; if curious, check out some of the standard references [20, 9].

6

Yoneda I: Internal vs. external

This section marks the start of a long journey that ends in the statement and proof of the Yoneda lemma. As with any important concept in category theory, the Yoneda lemma can be understood through a variety of perspectives. In these notes we will explore a particular perspective that feels most relevant to PL: the Yoneda lemma is a means of connecting the “internal” to the “external”.

Consider, for example, the case of product. In the preceding sections, we have given you the definition of product that is most commonly found in standard textbooks on category theory. This definition has an “internal” character, because it defines product in a category \mathcal{C} purely in terms of other objects and morphisms *inside* of \mathcal{C} .

But there is another way to define product. Think back to Theorem 5.0.1, which showed how to interpret a simple programming language in an arbitrary category \mathcal{C} with finite products. The proof of Theorem 5.0.1 gave a recipe for interpreting each of the judgments of programming language in \mathcal{C} . Following this recipe, each of the typing rules became *functions on morphisms*. For example, the typing rule

$$\frac{\Gamma \vdash M : A \times B}{\Gamma \vdash \text{fst } M : A} \text{T-FST}$$

for projecting out the first component of a pair was interpreted as a function

$$\{\text{morphisms } \llbracket \Gamma \rrbracket \text{ to } \llbracket A \times B \rrbracket\} \longrightarrow \{\text{morphisms } \llbracket \Gamma \rrbracket \text{ to } \llbracket A \rrbracket\},$$

namely the function that sends interpretations of $\llbracket M \rrbracket$ to interpretations of $\llbracket \text{fst } M \rrbracket$.

It turns out that we can promote this recipe into a *definition* for product. In contrast to the textbook definition in terms of objects and morphisms of \mathcal{C} , this definition characterizes the product in terms of *functions between sets* of morphisms. This gives it an “external”

character, because it defines the product in terms of set-theoretic objects (sets and functions) that live *outside* of \mathcal{C} .

First, we must make precise what we mean by this “external” definition of product that defines product in terms of functions on morphisms. To formalize this idea of “maps between morphisms”, it’s useful to have the following definition that we’ve put off for a long time.

Definition 15: Hom-set

Let X and Y be objects of a category \mathcal{C} . The *hom-set* of \mathcal{C} at X and Y , written $\mathcal{C}(X, Y)$, is the set of morphisms from X to Y :

$$\mathcal{C}(X, Y) := \{f \mid \text{dom}(f) = X, \text{cod}(f) = Y\}.$$


For example,

- In the category of finite sets and set functions FinSet , the hom-set $\text{FinSet}(\{a, b\}, \{1, 2, 3\})$ is the set of all functions from $\{a, b\}$ to $\{1, 2, 3\}$.
- Hom-sets can be empty. In the divisibility category (\mathbb{N}, \preceq) , the hom-set $\mathbb{N}(3, 4) = \emptyset$, while $\mathbb{N}(3, 6) = \{\top\}$.

This definition lets us state formally what it means for an inference rule like T-FST to denote a function from morphisms to morphisms. We shall say that the interpretation of T-FST is a family of functions $(\text{Fst}_\Gamma)_{\Gamma \in \mathcal{C}}$, indexed by the ambient context Γ , of the following type:

$$\text{Fst}_\Gamma : \mathcal{C}(\Gamma, A \times B) \rightarrow \mathcal{C}(\Gamma, A)$$

We can visualize this family of functions as follows:



As an example of working with this function, we have that $\text{Fst}_{\Gamma_1}(f) = f'$.

In addition to the basic type shown above, we will also require that Fst be *natural* in Γ . Intuitively, what this means is that Fst is in a sense “polymorphic in Γ ”. Just as polymorphic functions in System F respect “change of representation” for the type variables that they are

The word “hom” is an abbreviation of “homomorphism”, and the name hom-set comes from the fact many of the original motivating examples of categories had morphisms that were homomorphisms between algebraic structures.

polymorphic over [32], we will require that the family of functions $(\text{Fst}_\Gamma)_{\Gamma \in \text{Ob}(\mathcal{C})}$ respect “change of Γ ”. Concretely, this means that if any input f as shown in the picture above is changed into $f \circ s$ via some morphism $s : \Gamma' \rightarrow \Gamma$, it must be the case that $\text{Fst}_\Gamma(f)$ is changed into $\text{Fst}_{\Gamma'}(f \circ s)$ via s too. Formally, the following *naturality condition* must hold for any such f and s :

$$\text{Fst}_\Gamma(f) \circ s = \text{Fst}_{\Gamma'}(f \circ s)$$

Another intuition for this condition comes from the requirement that the interpretation of a programming language be *stable under substitution*. For instance, we know that $\text{fst } M$ is invariant under substitution for some (possibly multi-way) substitution s :

$$(\text{fst } M)[s] \equiv \text{fst } (M[s])$$

This program equation corresponds directly to the naturality condition given above.

Applying the above analysis to the typing rule T-SND gives an analogous function for extracting the second component of a pair:

$$\frac{\Gamma \vdash M : A \times B}{\Gamma \vdash \text{snd } M : B} \rightsquigarrow \text{Snd}_\Gamma : \mathcal{C}(\Gamma, A \times B) \rightarrow \mathcal{C}(\Gamma, B)$$

such that $\forall s : \Gamma' \rightarrow \Gamma, \text{Snd}_\Gamma(f) \circ s = \text{Snd}_{\Gamma'}(f \circ s)$

The rule T-PAIR for forming pairs becomes a function that takes in two morphisms and bundles them together into a morphism into the product:

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \langle M, N \rangle : A \times B} \rightsquigarrow \text{Pair}_\Gamma : \mathcal{C}(\Gamma, A) \times \mathcal{C}(\Gamma, B) \rightarrow \mathcal{C}(\Gamma, A \times B)$$

such that $\forall s : \Gamma' \rightarrow \Gamma, \text{Pair}_\Gamma(f, g) \circ s = \text{Pair}_{\Gamma'}(f \circ s, g \circ s)$

Finally, the β and η laws become equations between nested applications of the functions Fst , Snd , and Pair .

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \text{fst } \langle M, N \rangle \equiv M} \rightsquigarrow \text{Fst}_\Gamma(\text{Pair}_\Gamma(f, g)) = f$$

$$\frac{\Gamma \vdash M : A \quad \Gamma \vdash N : B}{\Gamma \vdash \text{snd } \langle M, N \rangle \equiv N} \rightsquigarrow \text{Snd}_\Gamma(\text{Pair}_\Gamma(f, g)) = g$$

$$\frac{\Gamma \vdash M : A \times B}{\Gamma \vdash M \equiv \langle \text{fst } M, \text{snd } M \rangle \equiv N} \rightsquigarrow \text{Pair}_\Gamma(\text{Fst}_\Gamma(f), \text{Snd}_\Gamma(f)) = f$$

Packaging up all of these constraints into a definition,¹

We will save a formal abstract definition of naturality for later on.

¹ We have chosen to name this concept “PL-product” as it was derived from the typing rules for product types in a programming language.

Definition 16: PL-product

Let A and B be objects of a category \mathcal{C} . A *PL-product* of A and B consists of:

- An object $A \times B$
- \mathcal{C} -indexed functions:

$$\begin{aligned} \text{Fst}_\Gamma &: \mathcal{C}(\Gamma, A \times B) \rightarrow \mathcal{C}(\Gamma, A) \\ \text{Snd}_\Gamma &: \mathcal{C}(\Gamma, A \times B) \rightarrow \mathcal{C}(\Gamma, B) \\ \text{Pair}_\Gamma &: \mathcal{C}(\Gamma, A) \times \mathcal{C}(\Gamma, B) \rightarrow \mathcal{C}(\Gamma, A \times B) \end{aligned}$$

satisfying the necessary properties to be \mathcal{C} -indexed, i.e.:

- Fst , Snd , and Pair “respect substitution”: for all $\Gamma' \xrightarrow{s} \Gamma$,

$$\begin{aligned} \text{Fst}_\Gamma(f) \circ s &= \text{Fst}_{\Gamma'}(f \circ s) && \text{for all } \Gamma \xrightarrow{f} A \times B \\ \text{Snd}_\Gamma(f) \circ s &= \text{Snd}_{\Gamma'}(f \circ s) && \text{for all } \Gamma \xrightarrow{f} A \times B \\ \text{Pair}_\Gamma(f, g) \circ s &= \text{Pair}_{\Gamma'}(f \circ s, g \circ s) && \text{for all } \Gamma \xrightarrow{f} A \text{ and } \Gamma \xrightarrow{g} B \end{aligned}$$

- Analogs of the β and η laws hold:

$$\begin{aligned} \text{Fst}_\Gamma(\text{Pair}_\Gamma(f, g)) &= f && \text{for all } \Gamma \xrightarrow{f} A \text{ and } \Gamma \xrightarrow{g} B \\ \text{Snd}_\Gamma(\text{Pair}_\Gamma(f, g)) &= g && \text{for all } \Gamma \xrightarrow{f} A \text{ and } \Gamma \xrightarrow{g} B \\ \text{Pair}_\Gamma(\text{Fst}_\Gamma(f), \text{Snd}_\Gamma(f)) &= f && \text{for all } \Gamma \xrightarrow{f} A \times B \end{aligned}$$

Now we can state an important theorem.

Theorem 6.0.1. PL-products are equivalent to ordinary products.

This theorem breaks down into multiple components.


1. Every PL-product can be turned into a categorical product.
2. Every categorical product can be turned into a PL product.
3. Constructions (1) and (2) round-trip to the identity.

The proof tackles each piece separately.

Construction 6.0.2. Every PL product of A and B can be turned into a categorical product.²

Proof. Suppose you have $(A \times B, \text{Fst}_\Gamma, \text{Snd}_\Gamma, \text{Pair}_\Gamma)$. Remember, Fst_Γ (and all these other functions) behave like polymorphic functions, and we need to find a diagram like:

² In general, we will use the word “construction” to signal a lemma whose proof details are going to be relevant later. In this case, we will need the precise definition of how PL products are turned into categorical ones later on, to prove that the two constructions we have done round-trip to the identity.



How do we get a morphism `fst`? The key is that `Fst` is polymorphic in Γ , so we can plug in $A \times B$ there to get something of the required type:

$$\text{Fst}_{A \times B} : \mathcal{C}(A \times B, A \times B) \rightarrow \mathcal{C}(A \times B, A)$$

Now we need to get a morphism of the required type of `fst`. From here, we can play “type tetris”: we call `FstA×B` with a morphism that meets the type constraint, the identity!


$$\text{Fst}_{A \times B}(\text{id}_{A \times B}) =: \text{fst}$$

We can do the same thing to get the morphism `snd`.

$$\text{Snd}_{A \times B}(\text{id}_{A \times B}) =: \text{snd}$$

Now the question: does this satisfy the universal property for products?

First, let’s show that there exists an h such that the diagram com-



Where can we find h ? We can let $h = \text{Pair}_C(f, g)$.

We need to show that the diagram commutes. The left triangle commutes by the following equational argument:

$$\begin{aligned} \text{fst} \circ h &= \text{Fst}_{A \times B}(\text{id}_{A \times B}) \circ \text{Pair}_C(f, g) && \text{expanding definitions} \\ &= \text{Fst}_C(\text{id}_{A \times B} \circ \text{Pair}_C(f, g)) && \text{naturality of Fst} \\ &= \text{Fst}_C(\text{Pair}_C(f, g)) && \text{identity law} \\ &= f && \beta \end{aligned}$$

The right triangle commutes similarly.

Finally, we need to show that `PairC(f, g)` is the unique h making the above diagram commute. For this we use the η law. Suppose h' satisfies `fst` \circ $h' = f$ and `snd` \circ $h' = g$. Then the following calculation

uses the η law to establish that $h' = h$.

$$\begin{aligned}
h' &= \text{Pair}_\Gamma(\text{Fst}_\Gamma(h'), \text{Snd}_\Gamma(h')) && \eta \\
&= \text{Pair}_\Gamma(\text{Fst}_{A \times B}(\text{id}_{A \times B}) \circ h', \text{Snd}_{A \times B}(\text{id}_{A \times B}) \circ h') && \text{substitution} \\
&= \text{Pair}_\Gamma(\text{fst} \circ h', \text{snd} \circ h') \\
&= \text{Pair}_\Gamma(f, g) && \text{assumption} \\
&= \text{Pair}_\Gamma(\text{fst} \circ h, \text{snd} \circ h) && \text{shown above} \\
&= \text{Pair}_\Gamma(\text{Fst}_{A \times B}(\text{id}_{A \times B}) \circ h, \text{Snd}_{A \times B}(\text{id}_{A \times B}) \circ h) \\
&= \text{Pair}_\Gamma(\text{Fst}_\Gamma(h), \text{Snd}_\Gamma(h)) && \text{substitution} \\
&= h && \eta
\end{aligned}$$

□

We can also go in the other direction: every product in the ordinary categorical sense can be turned into a PL-product.

Construction 6.0.3. Any product can be turned into a PL-product.

Proof. The previous construction provides a strong hint as to how to do this one. Given a product $(A \times B, \text{fst}, \text{snd})$, define

$$\begin{aligned}
\text{Fst}_\Gamma(f) &= \text{fst} \circ f \\
\text{Snd}_\Gamma(f) &= \text{snd} \circ f
\end{aligned}$$

for all morphisms $\Gamma \xrightarrow{f} A \times B$, and

$$\text{Pair}_\Gamma(f, g) = \langle f, g \rangle$$

for all morphisms $\Gamma \xrightarrow{f} A$ and $\Gamma \xrightarrow{g} B$. The fact that Fst and Snd respect substitution boil down to the following equations:

$$\begin{aligned}
(\text{fst} \circ f) \circ s &= \text{fst} \circ (f \circ s) \\
(\text{snd} \circ f) \circ s &= \text{snd} \circ (f \circ s)
\end{aligned}$$

The fact that Pair respects substitution boils down to Proposition 3.4.1.

The β law follows from the equations $\text{fst} \circ \langle f, g \rangle = f$ and $\text{snd} \circ \langle f, g \rangle = g$, and the η law from the uniqueness property of $\langle f, g \rangle$. □

Finally, we show that the two constructions just described roundtrip to the identity.

Proposition 6.0.4. Constructions 6.0.2 and 6.0.3 are mutually inverse.

Proof. We show both roundtrips are the identity. First, suppose one has a categorical product of A and B , hence a tuple $(A \times B, \text{fst}, \text{snd})$

satisfying the universal property of product. Applying Theorem 6.0.3 yields a PL-product $(A \times B, \text{Fst}, \text{Snd}, \text{Pair})$ where

$$\begin{aligned}\text{Fst}_\Gamma(f) &= \text{fst} \circ f \\ \text{Snd}_\Gamma(f) &= \text{snd} \circ f \\ \text{Pair}_\Gamma(f, g) &= \langle f, g \rangle\end{aligned}$$

Next, applying Theorem 6.0.2 to this PL-product yields a tuple $(A \times B, \text{fst}', \text{snd}')$ where

$$\begin{aligned}\text{fst}' &= \text{Fst}_{A \times B}(\text{id}_{A \times B}) \\ \text{snd}' &= \text{Snd}_{A \times B}(\text{id}_{A \times B})\end{aligned}$$

The goal is now to show that $\text{fst} = \text{fst}'$ and $\text{snd} = \text{snd}'$. This follows from a straightforward calculation. We show only the proof of $\text{fst} = \text{fst}'$; the case of snd is analogous.

$$\begin{aligned}\text{fst}' &= \text{Fst}_{A \times B}(\text{id}_{A \times B}) \\ &= \text{fst}_{A \times B} \circ \text{id}_{A \times B} \\ &= \text{fst}_{A \times B}\end{aligned}$$

Now for the other roundtrip. Suppose one has a PL-product $(A \times B, \text{Fst}, \text{Snd}, \text{Pair})$. Applying Theorem 6.0.2 to this PL-product yields the tuple product $(A \times B, \text{fst}, \text{snd})$ where

$$\begin{aligned}\text{fst} &= \text{Fst}_{A \times B}(\text{id}_{A \times B}) \\ \text{snd} &= \text{Snd}_{A \times B}(\text{id}_{A \times B})\end{aligned}$$

Then, applying Theorem 6.0.3 to this product yields $(A \times B, \text{Fst}', \text{Snd}', \text{Pair}')$ where

$$\begin{aligned}\text{Fst}_\Gamma(f) &= \text{fst} \circ f \\ \text{Snd}_\Gamma(f) &= \text{snd} \circ f \\ \text{Pair}_\Gamma(f, g) &= \langle f, g \rangle\end{aligned}$$

Note that in the final equation, $\langle f, g \rangle$ is guaranteed to exist by the universal property of $(A \times B, \text{fst}, \text{snd})$ established in Theorem 6.0.2.

The goal is now to show that $\text{Fst} = \text{Fst}'$ and $\text{Snd} = \text{Snd}'$ and $\text{Pair} = \text{Pair}'$. By function extensionality, it's enough to show these functions are equal for any object Γ and morphisms out of Γ . First, we have

$$\text{Fst}'_\Gamma(f) = \text{fst} \circ f = \text{Fst}_{A \times B}(\text{id}_{A \times B}) \circ f = \text{Fst}_\Gamma(f).$$

The case of Snd is analogous. Then, we have

$$\text{Pair}'_\Gamma(f, g) = \langle f, g \rangle \stackrel{(!)}{=} \text{Pair}_\Gamma(f, g).$$

Note that the equation marked (!) follows from the fact that, in the proof of Theorem 6.0.2, the unique mediating map h making the relevant product diagram involving f and g commute was built using Pair.³ □

³ It is also possible to establish (!) without rifling through the proof of Theorem 6.0.2, by appealing to the uniqueness property of the morphism $\langle f, g \rangle$.

7

Yoneda II: indexed set theory

In the last chapter we saw an important example of functions indexed by objects in a category, which formed the components of the PL-pair. This idea of sets and functions indexed by objects in a category is pervasive in programming languages and in category theory, and in this chapter we will spend some time exploring these constructions in their own right.

Many of the notions that you are familiar with from set theory will have \mathcal{C} -indexed analogues:

Ordinary	\mathcal{C} -Indexed	Intuition
Sets	\mathcal{C} -indexed sets	A copy of \mathcal{C} inside of sets
Functions	\mathcal{C} -indexed functions	Polymorphic functions
Cartesian products	\mathcal{C} -indexed product	Cartesian product that respects \mathcal{C} -indexing
...		

We will make use of these analogies extensively. But before we do, here is the formal definition of a \mathcal{C} -indexed set:

Definition 17: \mathcal{C} -indexed set

For a category \mathcal{C} , a \mathcal{C} -indexed set F is:

- A family of sets $F(X)$ for each object X in \mathcal{C}
- For each morphism $X \xrightarrow{f} Y$ in \mathcal{C} , a function $F(f) : F(Y) \rightarrow F(X)$ satisfying the following two *functoriality properties*:
 - preserves identity: $F(\text{id}_X) = \text{id}_{F(X)}$
 - preserves composition: $F(g \circ f) = F(f) \circ F(g)$ for all $X \xrightarrow{f} Y \xrightarrow{g} Z$.


A \mathcal{C} -indexed set is also known as a **presheaf**. Let's see some examples of \mathcal{C} -indexed sets:

- Consider a category with 1 object • and an identity morphism id :

Use your typechecker! Some of these compositions \circ are composition of morphisms, and some are composition of functions.

We will see soon that a \mathcal{C} -indexed set is an instance of a general phenomenon called a *functor*, which can be thought of as a structure-preserving map between categories.

At this point it might be a good exercise to consider a few more cases of \mathcal{C} -indexed sets. In particular, try to draw a balloon diagram that shows an example of a \mathcal{C} -indexed set for the following 3-object category:



7.1 \mathcal{C} -indexed functions

Next, let's define the notion of a \mathcal{C} -indexed function, which should behave like a "polymorphic function":¹

¹ We will see soon that a \mathcal{C} -indexed function is an instance of a general phenomenon called a *natural transformation* between functors F and G .

Definition 18: \mathcal{C} -indexed function

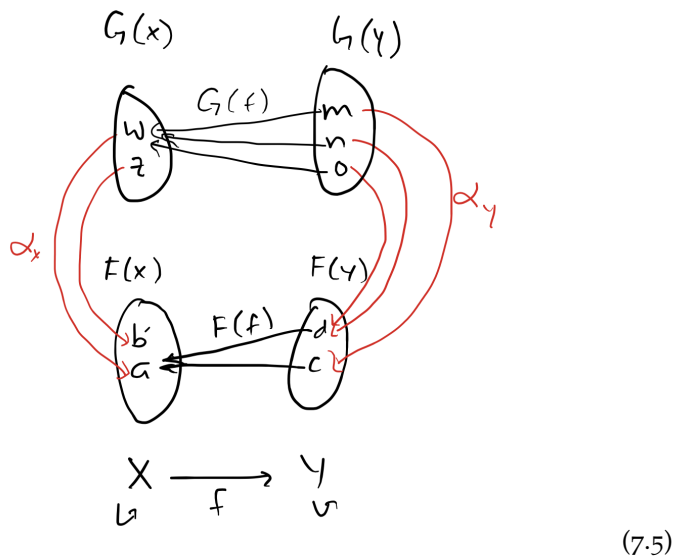
Given a category \mathcal{C} and two \mathcal{C} -indexed sets F and G , a \mathcal{C} -indexed function $\alpha : F \Rightarrow G$ is a family of functions $\alpha_X : F(X) \rightarrow G(X)$ for objects X of \mathcal{C} satisfying the following *naturality condition*:
 $\alpha_X \circ F(f) = G(f) \circ \alpha_Y$ for any \mathcal{C} morphism $X \xrightarrow{f} Y$. This naturality condition can be visualized as a commutative square:

$$\begin{array}{ccc}
 F(Y) & \xrightarrow{\alpha_Y} & G(Y) \\
 \downarrow F(f) & & \downarrow G(f) \\
 F(X) & \xrightarrow{\alpha_X} & G(X)
 \end{array}$$

Once again let's work through some examples of \mathcal{C} -indexed functions to get a sense for how they behave.

Returning to the one-object category in (7.1), \mathcal{C} -indexed functions are simply functions between sets. Concretely, suppose $F(\bullet) = \{a, b\}$ and $G(\bullet) = \{c, d\}$. Then, there is a \mathcal{C} -indexed function α where $\alpha_\bullet : F(\bullet) \rightarrow G(\bullet)$ is a function. The naturality constraint is somewhat trivially satisfied, but is worth checking as an exercise. Note that there are 4 possible choices of α here, one for each possible function between the sets $F(\bullet)$ and $G(\bullet)$.

A much more interesting case is the 2-object category with a morphism between them in (7.3). We can visualize the data of a particular \mathcal{C} -indexed function α in this category by drawing a "stacked balloon diagram":



In this picture, we have $G(X) = \{w, z\}$, $G(Y) = \{m, n, o\}$, $F(X) = \{b, a\}$, and $F(Y) = \{c, d\}$. Then, a particular \mathcal{C} -indexed function $\alpha : G \Rightarrow F$ is a family of functions $\alpha_X : G(X) \rightarrow F(X)$ and $\alpha_Y : G(Y) \rightarrow F(Y)$; an example of one such choice of α is shown in red. To check that this is a valid choice of α , we need to check that it satisfies the naturality condition, i.e. that $F(f) \circ \alpha_Y = \alpha_X \circ G(f)$. To check this, we can play the “two finger game”: put your finger on each particular element of $G(Y)$ and follow the two possible paths that it can take, and make sure your fingers end up in the same spot in $F(X)$. For instance, following the element m , we have $\alpha_Y(m) = c$, then $F(f)(c) = a$; in the other path, we have $G(f)(m) = w$, and $\alpha_X(w) = a$. So, these two paths agree on what to do with m , and so do all the other choices of starting element in $G(Y)$.

At this point, it’s a good exercise to try to find all possible $\alpha : G \Rightarrow F$ for the \mathcal{C} -indexed sets F and G shown in (7.5).

7.2 The category of \mathcal{C} -indexed sets (presheaf categories)

At this point, you might be smelling something categorical: might it be the case that, for any category \mathcal{C} , one can form a category whose objects are all possible \mathcal{C} -indexed sets and whose morphisms are all possible \mathcal{C} -indexed functions? Indeed one can, but one requires the following to show this:

- Just as there is an identity function in set theory, there is an identity \mathcal{C} -indexed function in \mathcal{C} -indexed set theory. For each \mathcal{C} -indexed set P , there is a \mathcal{C} -indexed function $\text{id}_P : P \Rightarrow P$ defined by $\text{id}_{P,X} = \left(P(X) \xrightarrow{\text{id}_{P(X)}} P(X) \right)$. The naturality condition amounts to

commutativity of the following diagram.

$$\begin{array}{ccc} F(Y) & \xrightarrow{\text{id}_{F(Y)}} & F(Y) \\ F(f) \downarrow & & \downarrow F(f) \\ F(X) & \xrightarrow{\text{id}_{F(X)}} & F(X) \end{array}$$

- \mathcal{C} -indexed functions can also be composed: the composition of $\alpha : Q \Rightarrow R$ and $\beta : P \Rightarrow Q$, written $\alpha \circ \beta : P \Rightarrow R$, is defined by $(\alpha \circ \beta)_X = \alpha_X \circ \beta_X$. The naturality condition amounts to the commutativity of the following diagram.

$$\begin{array}{ccccc} F(Y) & \xrightarrow{\beta_Y} & G(Y) & \xrightarrow{\alpha_Y} & H(Y) \\ \downarrow F(f) & & \downarrow G(f) & & \downarrow H(f) \\ F(X) & \xrightarrow{\beta_X} & G(X) & \xrightarrow{\alpha_X} & H(X) \end{array}$$

From this we can conclude the following definition is well-formed:²

Definition 19: Category of \mathcal{C} -indexed sets

Let \mathcal{C} be a category. Then, the *category of \mathcal{C} -indexed sets* is the category whose objects are \mathcal{C} -indexed sets and whose morphisms are \mathcal{C} -indexed functions, where identities and composition of morphisms are defined as above.

This category is also called the **category of presheaves on \mathcal{C}** . In general, \mathcal{C} -indexed set theory looks kind of like set theory “with shape \mathcal{C} ”. The category of \mathcal{C} -indexed sets has fantastically rich structure: in fact, it is something called a *topos*, an idea that we will return to later. For now, the most exciting thing we can say about this category is that it is Cartesian-closed, meaning that it has products. Let’s show this fact:

Theorem 7.2.1. Let \mathcal{C} be a category. Then, the category of presheaves on \mathcal{C} has products.

Proof. We prove this using the familiar internal argument involving the universal property of products. We begin by identifying a candidate \mathcal{C} -indexed set, and then we prove that it satisfies the universal property of products in the category of presheaves on \mathcal{C} :

Definition 20: \mathcal{C} -indexed product

Given a category \mathcal{C} and two \mathcal{C} -indexed sets F and G , the \mathcal{C} -indexed product of F and G is the \mathcal{C} -indexed set $F \times G$ defined

² This definition technically runs afoul of some cardinality limitations in our original definition of a category (Definition 1). Do not worry; the original definition of category can be adjusted to account for this fact without invalidating anything we have done so far.

by:

- On objects $X \in \mathcal{C}$, $(F \times G)(X) = F(X) \times G(X)$
- On morphisms $X \xrightarrow{f} Y$, $(F \times G)(f)$ is the function $F(Y) \times G(Y) \rightarrow F(X) \times G(X)$ defined by $(F \times G)(f) = (x, y) \mapsto (F(f)(x), G(f)(y))$

Alongside $F \times G$ there are \mathcal{C} -indexed functions $\pi_1 : F \times G \Rightarrow F$ and $\pi_2 : F \times G \Rightarrow G$ defined:

$$\begin{aligned} \pi_{1,X} : (F \times G)(X) &\rightarrow F(X) \\ &= (x, y) \mapsto x \end{aligned}$$

and

$$\begin{aligned} \pi_{2,X} : (F \times G)(X) &\rightarrow G(X) \\ &= (x, y) \mapsto y \end{aligned}$$

First we need to show $F \times G$ is a \mathcal{C} -indexed set:

- *Preserves identity:*

$$\begin{aligned} (F \times G)(\text{id}_X) &= (a, b) \mapsto (\text{id}_{F(X)}(a), \text{id}_{G(X)}(b)) \\ &= (a, b) \mapsto \text{id}_{F(X) \times G(X)}(a, b) \\ &= \text{id}_{F(X) \times G(X)}. \end{aligned}$$

- *Preserves composition:* For morphisms $X \xrightarrow{f} Y \xrightarrow{g} Z$:

$$\begin{aligned} (F \times G)(g \circ f) &= (a, b) \mapsto (F(g \circ f)(a), G(g \circ f)(b)) && \text{Def} \\ &= (a, b) \mapsto ((F(f) \circ F(g))(a), (G(f) \circ G(g))(b)) && \text{Functoriality of } F \text{ and } G \\ &= (a, b) \mapsto [(F \times G)(f) \circ (F \times G)(g)](a, b) && \text{Def} \\ &= (F \times G)(f) \circ (F \times G)(g). \end{aligned}$$

We also need to show that π_1 and π_2 are \mathcal{C} -indexed functions. Let $X \xrightarrow{f} Y$ be a morphism in \mathcal{C} . Then, we must show the naturality diagram commutes:

$$\begin{array}{ccc} (F \times G)(Y) & \xrightarrow{\pi_{1,Y}} & F(Y) \\ \downarrow (F \times G)(f) & & \downarrow F(f) \\ (F \times G)(X) & \xrightarrow{\pi_{1,X}} & F(X) \end{array}$$

We can show this by simple calculation. Let $(a, b) \in (F \times G)(Y)$. Then, $[F(f) \circ \pi_{1,Y}](a, b) = F(f)(a)$. For the other path, we have

that $[(F \times G)(f)]$ applied to (a, b) is the pair $(F(f)(a), F(f)(b))$, so $\pi_{1,X} \circ [(F \times G)(f)]$ applied to (a, b) is the first component $F(f)(a)$. An identical argument holds for $\pi_{2,X}$.

Next, we need to show that $(F \times G, \pi_1, \pi_2)$ satisfies the universal property for products. Suppose there is a \mathcal{C} -indexed set Γ with \mathcal{C} -indexed functions $f : \Gamma \Rightarrow F$ and $g : \Gamma \Rightarrow G$. Then, we must show there exists a unique \mathcal{C} -indexed function $\langle f, g \rangle : \Gamma \Rightarrow F \times G$ such that the relevant diagram commutes. Let's define $\langle f, g \rangle : \Gamma \Rightarrow F \times G$ as:

$$\begin{aligned} \langle f, g \rangle_X : \Gamma(X) &\rightarrow (F \times G)(X) \\ &= a \mapsto (f_X(a), g_X(a)) \end{aligned}$$

The naturality of this definition follows from the naturality of f and g : for any morphism $Y \xrightarrow{s} X$ of \mathcal{C} , we have that

$$(F \times G)(s) \circ \langle f, g \rangle_X = \langle f, g \rangle_Y \circ \Gamma(s)$$

because, for any input $a \in \Gamma(X)$, it holds that

$$\begin{aligned} \text{LHS}(a) &= [(F \times G)(s)](f_X(a), g_X(a)) && \text{Def. of } \langle f, g \rangle_X \\ &= (F(s)(f_X(a)), G(s)(g_X(a))) && \text{Def. of } (F \times G)(s) \\ &= (f_Y(\Gamma(s)(a)), g_Y(\Gamma(s)(a))) && \text{Naturality of } f \text{ and } g \\ &= \langle f, g \rangle_Y(\Gamma(s)(a)) && \text{Def. of } \langle f, g \rangle_Y \\ &= \text{RHS}(a). \end{aligned}$$

Next we have the following for any indexing object X , establishing that the relevant product diagram commutes.

$$\pi_{1,X} \circ \langle f, g \rangle_X = f_X, \quad \pi_{2,X} \circ \langle f, g \rangle_X = g_X.$$

Finally, all that remains is to show uniqueness of $\langle f, g \rangle$. This follows immediately from function extensionality: if $\langle f, g \rangle$ is to satisfy $\pi_{1,X} \circ \langle f, g \rangle_X = f_X$ and $\pi_{2,X} \circ \langle f, g \rangle_X = g_X$ for all X , then it must be defined as above. \square

8

Yoneda III: Representability

The crucial idea of the Yoneda lemma is that a copy of \mathcal{C} lives inside of its corresponding presheaf category. This lets ideas, definitions, and intuitions from the \mathcal{C} -indexed category – which, as we’ve already described, has fantastically rich structure – including products, exponentials, etc. – be reflected back into the indexing category \mathcal{C} .

8.1 Representables


Our first step towards understanding this picture is to come up with a way of injecting a category \mathcal{C} into its presheaf category. Here is an inscrutable definition that explains how to do this:

Definition 21: The representable \mathcal{C} -indexed set at X
<p>The <i>representable \mathcal{C}-indexed set at X</i> is written $\mathcal{y}X$ (pronounced “yo X”) and defined by</p> $ \begin{aligned} (\mathcal{y}X)(A) &= \mathcal{C}(A, X) \\ (\mathcal{y}X)(s : A' \rightarrow A) : (\mathcal{y}A) &\rightarrow (\mathcal{y}A') \\ &= (f : A \rightarrow X) \mapsto (f \circ s : A' \rightarrow X) \end{aligned} $

This definition, on first glance, is very hard to unpack. So we will spend some time with it. The intuition you should have for $\mathcal{y}X$ is that it is “the external view of the object X in the category \mathcal{C} ”. Let’s start by drawing a balloon diagram to picture one of these representables. Consider the following example category:

$$\begin{array}{ccccc}
 X_1 & \xrightarrow{g_1} & X_2 & \xrightarrow{g_2} & X_3 \\
 & \searrow f_1 & \downarrow f_2 & \swarrow f_3 & \\
 & & Z & &
 \end{array} \tag{8.1}$$

Let’s draw a balloon diagram that visualizes $\mathcal{y}Z$:



To build intuition, you should think of \mathbb{Z} as describing “the perspective every other object in \mathcal{C} has on Z .” Unpacking this more:

- Let’s look at the action of \mathbb{Z} on Z first. Unfolding the definition, we have:

$$(\mathbb{Z})(Z) = \mathcal{C}(Z, Z) = \{\text{id}_Z\}$$

So, somewhat vacuously, we might say: “the perspective Z has on itself is the identity morphism”. And, for X_1 , we have:

$$(\mathbb{Z})(X_1) = \mathcal{C}(X_1, Z) = \{f_1, f_2 \circ g_1, f_3 \circ g_2 \circ g_1\}$$

This is slightly more interesting: we would say, “ X_1 has 3 perspectives on Z : (1) via the f_1 morphism, (2) by following $f_2 \circ g_1$; (3) by following $f_3 \circ g_2 \circ g_1$ ”.

- Now let’s look at the action on morphisms. These will correspond to “changes in perspective”. Let’s examine the action of \mathbb{Z} on the morphism f_1 :


$$\begin{aligned} (\mathbb{Z})(f_1 : X_1 \rightarrow Z) : \mathbb{Z} &\rightarrow \mathbb{Z} \\ &= \text{id}_Z \mapsto \text{id}_Z \circ f_1 = f_1 \end{aligned}$$

One can understand this as: “to move from “ X_1 ’s perspective on Z to Z ’s perspective on Z , follow the f_1 morphism”. Somewhat more interestingly, we can also shift perspectives between X_1 and X_2 by precomposing with the g_1 morphism:

$$\begin{aligned} (\mathbb{Z})(g_1 : X_1 \rightarrow X_2) : (\mathbb{Z})(X_2) &\rightarrow (\mathbb{Z})(X_1) \\ &= \begin{cases} f_2 \mapsto f_2 \circ g_1 \\ f_3 \circ g_2 \mapsto f_3 \circ g_2 \circ g_1 \end{cases} \end{aligned}$$

8.1.1 Representables in the divisibility category

Let's consider what representables look like in a familiar category to get a better sense of their properties. Consider the following subset of the natural numbers ordered by divisibility:



Let's consider the representables of each element in this category. First let's inspect \mathbb{N}_{12} . We can compute its action on objects:


$$\begin{aligned}
 (\mathbb{N}_{12})(12) &= \{\text{id}_{12}\} & (\mathbb{N}_{12})(6) &= \{f\} & (\mathbb{N}_{12})(3) &= \{f \circ e\} \\
 (\mathbb{N}_{12})(2) &= \{f \circ d\} & (\mathbb{N}_{12})(1) &= \{f \circ d \circ c\} & (\mathbb{N}_{12})(5) &= \{\}
 \end{aligned}$$

Note that in this category there is at most 1 morphism between any 2 objects (i.e., we have that the morphisms $f \circ d \circ c = f \circ e \circ b$, so we only include 1 of them in $(\mathbb{N}_{12})(1)$). So, each representable \mathbb{N}_X is either a singleton (meaning $X \preceq 12$), or it is an empty set (meaning $X \not\preceq 12$). The action of \mathbb{N}_{12} on morphisms is not complicated; let's see what it is by unfolding definitions:

$$\begin{aligned}
 (\mathbb{N}_{12})(c) : (\mathbb{N}_{12})(2) &\rightarrow (\mathbb{N}_{12})(1) \\
 &= (f \circ d) \mapsto (f \circ d) \circ c
 \end{aligned}$$

In this case we see that to shift perspective from 1 to 2, we should precompose by the c morphism. From this, we can observe that there is a single morphism $\mathbb{N}_Y \rightarrow \mathbb{N}_X$ if and only if there is a morphism $X \rightarrow Y$ in \mathcal{C} .

Now, let's closely examine the structure of a representable \mathbb{N}_{12} and try to gain a more global picture of what data it contains. We can summarize \mathbb{N}_{12} by the following balloon, where we've replaced each object by its balloon:



We see that the collection of objects with a \star in them are exactly those objects that are at or below $\downarrow 12$ in the order. This collection of objects has a name in order theory: it is called the **principal ideal** of $\downarrow 12$. Formally:

Definition 22: Principal ideal

Let (X, \leq) be a preorder. Then, the **principal ideal** of an element $a \in X$, written $\downarrow a$, is the smallest collection of all elements less than or equal to a in X , i.e. $\downarrow a = \{x \in X \mid x \leq a\}$.

So, we can summarize: the representable $\downarrow 12$ is the principal ideal $\downarrow 12$, i.e., it contains exactly as much information as is contained in the principal ideal of $\downarrow 12$. In general, for any category formed by a preorder, it is the case that the representable of an element is a principal ideal in the preorder. This order-theoretic viewpoint sheds some light on the intuition that the representable of an element is “all the ways of viewing the element in the category”. One can imagine “peering up” at an element X from all the other objects: $\downarrow X$ is then all the vantage points below X from which it can be seen.

8.2 Representables in STLC

So far we’ve been studying representables in thin categories, where there is a very pleasant order-theoretic understanding what they mean in terms of principal ideals. Let’s see an example of representables in a category with more than 1 morphism between objects. Recall the syntax of STLC:

$$\begin{aligned}
 A, B &::= \text{Unit} \mid A \times B \mid A \rightarrow B \\
 M, N &::= \langle \rangle \mid \langle M, N \rangle \mid \text{fst } M \mid \text{snd } M \mid \lambda x. M \mid MN \mid x
 \end{aligned}
 \tag{STLC}$$

Recall the category of simply-typed lambda calculus terms quotiented by equations:

Definition 23: The category STLC

The category of simply-typed lambda calculus terms (STLC) is the category where:

- Objects are types A ;
- Morphisms $A \rightarrow B$ are equivalence classes of terms $[x : A \vdash e : B]$ where equality is given by the equational theory of the β and η laws;
- Composition is given by substitution:

$$[b : B \vdash N : C] \circ [a : A \vdash M : B] = [a : A \vdash N[M/b] : C]$$

- The identity morphism is the class of terms $[a : A \vdash a : A]$.

Let's see the representables in this category. For example, what do the representables $\mathfrak{y}(\mathbf{Unit} \times \mathbf{Unit})$ look like? Considering a concrete example:

$$\begin{aligned} \mathfrak{y}(\mathbf{Unit} \times \mathbf{Unit})(\mathbf{Unit}) &= \{\text{the ways of viewing terms of type } \mathbf{Unit} \times \mathbf{Unit} \text{ from terms of type } \mathbf{Unit}\} \\ &= \{[a : \mathbf{Unit} \vdash M : \mathbf{Unit} \times \mathbf{Unit}] \mid M \text{ an stlc term}\} \\ &= [a : \mathbf{Unit} \vdash \langle a, a \rangle : \mathbf{Unit} \times \mathbf{Unit}] \end{aligned}$$

In the case above, there is only 1 element of the set $\mathfrak{y}(\mathbf{Unit} \times \mathbf{Unit})(\mathbf{Unit})$: this is a quirk of our choice of types. Intuitively, the representables $(\mathfrak{y}B)(A)$ for two types A and B is the collection of all equivalence classes of terms M of type $A \rightarrow B$. Then, naturally, the action on morphisms $(\mathfrak{y}X)(f : A \rightarrow B) : \mathfrak{y}B \rightarrow \mathfrak{y}A$ ought to be the ways of shifting viewpoints from A to B , i.e.:

$$\begin{aligned} (\mathfrak{y}X)([a : A \vdash M : B]) : (\mathfrak{y}X)(B) &\rightarrow (\mathfrak{y}X)(A) \\ &= [b : B \vdash N : X] \mapsto [b : B \vdash N : X] \circ [a : A \vdash O : B] \\ &= [b : B \vdash N : X] \mapsto [a : A \vdash N[O/b] : X] \end{aligned}$$

Carefully check all the types in the above equations.

8.3 Representability

There can be many \mathcal{C} -indexed sets for a category, but the representables are very special due to their relationship with the indexing category. As usual in category theory, we are concerned not with particular objects but rather with equivalence classes of objects up to isomorphism, a notion we make formal here:

Definition 24: Isomorphism of \mathcal{C} -indexed sets

Two \mathcal{C} -indexed sets P, Q are *isomorphic* if there are \mathcal{C} -indexed functions $\alpha : P \Rightarrow Q$ and $\beta : Q \Rightarrow P$ such that $\alpha \circ \beta = \text{id}_Q$ and $\beta \circ \alpha = \text{id}_P$.

With this definition in hand we can make precise the set of \mathcal{C} -indexed sets that are tied to the indexing category:

Definition 25: Representability

A \mathcal{C} -indexed set P is *representable* if there is an isomorphism $\alpha : P \cong \text{y}X : \beta$ for some object X of \mathcal{C} .

Proposition 8.3.1. The FinSet -indexed set $\text{y}X \times \text{y}Y$ is representable.

Proof. for any finite set A , there is an isomorphism of sets

$$\alpha_A : \text{FinSet}(A, X \times Y) \cong \text{FinSet}(A, X) \times \text{FinSet}(A, Y).$$

Concretely, α_A is the function that sends a $f \in \text{FinSet}(A, X \times Y)$ to $(\pi_1 \circ f, \pi_2 \circ f) : \text{FinSet}(A, X) \times \text{FinSet}(A, Y)$. Its inverse is defined by $\beta_A(f, g) = \langle f, g \rangle$. Both α and β can be checked satisfy the naturality condition. Thus $X \times Y$ represents $\text{y}X \times \text{y}Y$. \square

Representability is the bridge between the internal view and the external view of a category: it lets us reflect properties presheaf category on \mathcal{C} back into \mathcal{C} itself. Concretely, in the next chapter, we will use the Yoneda lemma in full to show:

Theorem 8.3.2. Let P, A, B be objects in a category \mathcal{C} . Then, P is a product of A and B if and only if it represents the product in the presheaf category $\text{y}A \times \text{y}B$.

We will not prove this theorem here – it is a direct consequence of the Yoneda lemma, which we are now ready to appreciate – but one can understand this theorem as: an object P is a product of A and B internally if and only if it behaves like a product externally.

9

Yoneda IV: universal constructions, elements, and properties

Now we are ready to appreciate the Yoneda lemma:

Theorem 9.0.1 (Baby Yoneda). Let \mathcal{C} be a category, X be an object in \mathcal{C} , and P be a \mathcal{C} -indexed set. Then, there is a bijective correspondence:¹

$$\{\mathcal{C}\text{-indexed functions } \mathcal{C}X \Rightarrow P\} \cong P(X)$$

On first glance this theorem is profoundly mysterious because it stipulates that there is a very strong relationship between two very different-seeming things. Let's unpack this slowly to really understand what is going on by instantiating the Yoneda lemma to a specific category and \mathcal{C} -indexed set. Let's consider a finite category with 3 objects:

$$X \xrightarrow{f} Y \xrightarrow{g} Z \tag{9.1}$$

Now, let's define a simple \mathcal{C} -indexed set P :

$$P(X) = \{a, b\} \quad P(Y) = \{c, d\} \quad P(Z) = \{e, h\}$$

$$P(f) = \begin{cases} c \mapsto a \\ d \mapsto b \end{cases} \quad P(g) = \begin{cases} e \mapsto c \\ h \mapsto c \end{cases}$$

Continuing to unpack the data, we can define $\mathcal{C}Z$, which is also a \mathcal{C} -indexed set:


$$(\mathcal{C}Z)(Z) = \{\text{id}_Z\} \quad (\mathcal{C}Z)(Y) = \{g\} \quad (\mathcal{C}Z)(X) = \{g \circ f\}$$

$$(\mathcal{C}Z)(f) = g \mapsto g \circ f \quad (\mathcal{C}Z)(g) = \text{id}_Z \mapsto g$$

Next, let's enumerate the natural transformations $\mathcal{C}X \Rightarrow P$. If the Yoneda lemma is true, we don't have too much work to do: we know

¹ This is called the "baby Yoneda lemma" because the full Yoneda lemma further stipulates that this isomorphism satisfies additional naturality conditions. The baby Yoneda lemma is still quite strong, however, and is much easier to prove.

there should be *only two of them*, since $P(X)$ has only two elements in it. Let's enumerate them on a balloon diagram:



Here we've drawn the two possible natural transformations α_1 and α_2 in solid red and dashed blue respectively. Now, let's understand why the Yoneda lemma holds by tracing the possible paths through this balloon diagram with your fingers. Recall the two finger rule: a \mathcal{C} -indexed function $\mathcal{X}Z \Rightarrow P$ satisfies the naturality condition if, starting at any element in any balloon $\mathcal{X}X$, your two fingers move in synchrony along all possible paths. So, let's start with $\mathcal{X}Z = \{\text{id}_Z\}$. Let's play the two finger game together to design the two possible natural transformations:

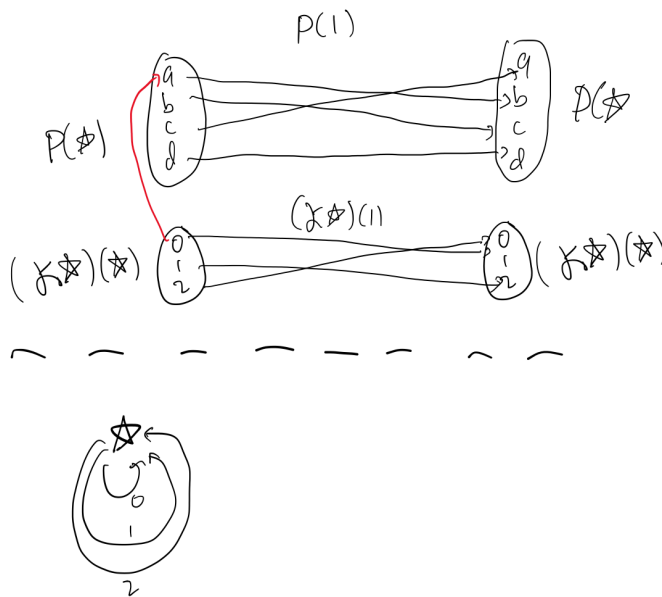
- Your left finger moves to g . Where can your right finger go? It can go either to e or h . So, choose h and so we are designing α_1 . *The rest of the moves are forced*: your right finger must now move to c , then to a . Your left finger is also forced: it must move to g then to $f \circ g$. Now, all other decisions are forced, and we can fill in α_1 so that all moves are in harmony.
- Next, we can choose moving our right finger to e , and so we are designing α_2 . *Again, all moves are forced*: the right finger follows the functions in P , and the left finger follows the functions in $\mathcal{X}Z$. Then, we are forced to choose each α to exactly line up exactly with how these fingers move: there is only one choice for α at each step.

Notice what is happening: because, in this example, there is a single element in $(\mathcal{X}Z)(Z)$, then there must be *exactly 1* natural transformation for each possible element in $P(Z)$. Naturality is a very strong restraint to place on α !

So, it is now hopefully quite clear that, if id_Z is the only element of $\mathcal{X}Z$, the baby Yoneda lemma is true, since the entire action of a

natural transformation is determined by its action on the identity morphism of Z . But, what if there is more than one morphism in $(\downarrow Z)(Z)$?

The crux is that the choice of $(\downarrow Z)(id_Z)$ fully determines the action of $\downarrow Z$ on all the other morphisms $Z \rightarrow Z$. Let's see this visually. Suppose we have as our indexing category the monoid corresponding to the natural numbers mod 3 under addition, the category \mathbb{Z}_3 . This category has a single object \star and 3 morphisms: the identity morphism 0, and morphisms 1 and 2, together with equations stipulating that $1 \circ 1 = 2, 1 \circ 2 = id, 2 \circ 2 = 1, 2 \circ 1 = id$. Let's visualize some of this data:



Here we've drawn an arbitrary \mathcal{C} -indexed set P with its action on the 1 morphism, and we've chosen arbitrarily that the \mathcal{C} -indexed function α_1 maps 0 to a . Note that we've drawn the balloon twice for the same object; this is for visualization purposes only.

Now, the key: *does this choice of α fully determine how α behaves on all other elements of $(\downarrow \star)(\star)$?* Indeed it does! Observe: by naturality, we have that $P(1) \circ \alpha = \alpha \circ (\downarrow \star)(1)$. This gives us a system of constraints that uniquely defines α . Observe:

$$\begin{aligned} ((P(1) \circ \alpha))(0) &= b = (\alpha \circ (\downarrow \star)(1))(0) \\ \Rightarrow b &= \alpha((\downarrow \star)(1)(0)) \\ \Rightarrow b &= \alpha(1). \end{aligned}$$


A similar chain of logic lets us derive $\alpha(2) = c$. Notice what is happening here: P is a functor, so it must "behave like the monoid": i.e., it has the same equations and morphism compositions. Here, the choice of $\alpha(0)$ is *picking an element of P that behaves like the identity*.

Now, how do we see that there are exactly 4 natural transformations $\alpha : \mathbb{Z}_3 \Rightarrow P$? Let's look at P . Since P is a functor, it has copies of the \mathbb{Z}_3 in its codomain. We can see two copies: the $\{a, b, c\}$ subset is an exact copy – the elements a, b , and c correspond to distinct numbers, and α will pick an arbitrary 0 point – and the $\{d\}$ subset is a degenerate copy with only a single number 0. So, there are 3 natural transformations that send 0 to $\{a, b, c\}$, each picking a different 0 point, and 1 natural transformation into $\{d\}$ that sends all elements of $(\mathbb{Z}_3)(\star)$ to d .

9.1 Proving the baby Yoneda lemma

Now we are ready to prove Theorem 9.0.1. Let \mathcal{C} be a category, P be a \mathcal{C} -indexed set, and X be an object in \mathcal{C} . Then, we need to construct a bijection:

Thanks to Bex Golovanov for the scribe notes for the next two subsections. up a bit.



What does this bijective correspondence statement actually mean? We have some $func$ that, given a $p \in PX$, can construct a natural transformation $\alpha : \mathbb{Z}X \Rightarrow P$, and some g that can extract p from a given natural transformation.

1. Constructing natural transformation from $p \in PX$

- Given a $p \in PX$, we defined $func(p)$ via a family of functions $func(p)_A$, where $func(p)$ is a natural transformation. We can think of the $func(p)_A$ as the components of $func(p)$, where for each object A in \mathcal{C} , the component $func(p)_A$ maps from $(\mathbb{Z}X)(A)$ to $P(A)$.

$$\begin{aligned}
 func(p)_A : (\mathbb{Z}X)(A) &\rightarrow PA \\
 f &\mapsto (Pf)(p)
 \end{aligned}$$

where $f : A \rightarrow \mathbb{Z}X$. The element p by definition has type PX , and P is contravariant so Pf maps PX to PA , giving us the desired codomain.

- Naturality of $func(p)$. Let $g : B \rightarrow A$. I don't really want to

justify this but here is the square

$$\begin{array}{ccc}
 (\downarrow X)(A) & \xrightarrow{func(p)_A} & P(A) \\
 (\downarrow X)(g) \downarrow & & \downarrow P(g) \\
 (\downarrow X)(B) & \xrightarrow{func(p)_B} & P(B)
 \end{array}$$

2. Getting p from natural transformation

- How did we extract $p \in PX$ from some given natural transformation α ? We know α_X maps $(\downarrow X)(X) \rightarrow PX$, so applying α_X to id_X gives us an element of PX , we can set $p := \alpha_X(id_X)$.
- Now we need to show $\alpha_A = func(p)_A$. Since α_A acts on the elements of $(\downarrow X)(A)$, it suffices to show $\alpha_A(f) = func(p)_A = (Pf)(p)$ for every $f \in \mathcal{C}(A, X)$.
- We have $(Pf)(p) = (Pf)(\alpha_X(id_X))$. We will use the naturality of α . Specifically, $Pf \circ \alpha_X = \alpha_A \circ (\downarrow X)(f)$.

$$\begin{array}{ccc}
 (\downarrow X)(X) & \xrightarrow{\alpha_X} & P(A) \\
 (\downarrow X)(f) \downarrow & & \downarrow P(f) \\
 (\downarrow X)(A) & \xrightarrow{\alpha_A} & P(B)
 \end{array}$$

$$\begin{aligned}
 (Pf)(\alpha_X(id_X)) &= (Pf \circ \alpha_X)(id_X) = (\alpha_A \circ (\downarrow X)(f))(id_X) \\
 &= \alpha_A((\downarrow X)(f)(id_X)) = \alpha_A(id_X \circ f) \\
 &= \alpha_A(f)
 \end{aligned}$$

So $(Pf)(p) = \alpha_A(f)$, and we are done!

Remark 9.1.1. Defining some natural transformation $\alpha := f(p)$ by defining how $f(p)$ acts on each object A is like a "pointwise" definition of α .

9.2 The category of elements

Definition 26: Category of elements

Let P be a \mathcal{C} -indexed set. The *category of elements* of P , written $\int_{\mathcal{C}} P$, is a category where:

1. objects : pairs $(X, p \in PX)$
2. morphisms : from $(X, p \in PX)$ to $(Y, q \in PY)$ are morphisms $f : X \rightarrow Y$ of \mathcal{C} such that $(Pf)(q) = p$. Alternatively, $q \cdot_P f =$

p . A morphism of this category of elements is a morphism from the original category with the *additional* structure of preserving the points of PX .

Definitions of composition and identity follow immediately from the definitions in \mathcal{C} , with the additional restriction of point preservation. It is relatively simple to show the desired categorical properties hold.

Proposition 9.2.1. Suppose \mathcal{C} -indexed set P is representable by X . There exists a $u \in PX$ called the *P -universal element* which satisfies the following universal property : for any object Γ of \mathcal{C} and any $g \in P\Gamma$, there exists a unique $\hat{g} : \Gamma \rightarrow X$ such that $g = u \cdot_P \hat{g}$. Alternatively, $g = (P\hat{g})(u)$.

Proposition 9.2.2. A \mathcal{C} -indexed set P is representable if $P \cong \downarrow X$ for some object X of \mathcal{C} . Suppose P is representable. Then there exists a *universal element* $u \in PX$ and u is a terminal object in the *category of elements* of P .

Example 9.2.3 (Category of elements of the terminal presheaf). For any category \mathcal{C} , the presheaf category on \mathcal{C} has a terminal object $\mathbf{1}$ defined:

$$\mathbf{1}(X) = \{\star\}$$

$$\mathbf{1}(X \xrightarrow{h} Y) = \star \mapsto \star$$

The category of elements $\int \mathbf{1}$ has objects (X, p) where $p = \star$ and X are objects of \mathcal{C} , and has morphisms $(X, p) \xrightarrow{h} (Y, q)$ where $F(h)(p) = q$; this point preservation is clearly trivially satisfied since $p = q = \star$. So, we can conclude that $\int \mathbf{1} = \mathcal{C}$, i.e., the category of elements of the terminal object in \mathcal{C} is the original category \mathcal{C} .

We have from Theorem 9.2.2 the definition of the universal property that we want these mysterious universal elements to satisfy. We can think about this definition as stating that we can “factor” any $g \in P\Gamma$ (for any Γ) into u and a unique \hat{g} .

What does it mean for $u \in P(X)$ to be terminal in the category of elements of P ? That any object has a unique morphism mapping to (X, u) . Explicitly, for any $(A, p \in PA)$, there is a unique map $f : A \rightarrow X$ such that $(Pf)(u) = p$ (or $p = u \cdot_P f$). This looks pretty similar to the universal property given in Theorem 9.2.2 – we can “factor” any $p \in PA$ into f and u .

How do we get Theorem 9.2.2 from the Yoneda Lemma? These are the key ideas:

- X representing P means we have an isomorphism between $\downarrow X$ and P . This isomorphism is a \mathcal{C} -indexed function, which we

have learned is usually called a natural transformation. So technically we have a bijective natural transformation (a natural isomorphism?)

- The Yoneda Lemma says that for every natural transformation $\natural X \Rightarrow P$, we should have a corresponding $u \in PX$ from which we can construct said transformation. This means the natural isomorphism between $\natural X$ and P has a *corresponding* $u \in PX$ (suggestively labeled).
- Let α denote the natural isomorphism $\natural X \Rightarrow P$. Because α is bijective, and α can be defined component-wise, each α_A for each object A must also be bijective. So consider this bijection $\alpha_A : (\natural X)(A) \Rightarrow PA$. From surjectivity, for any $p \in PA$, there exists some \hat{p} in $(\natural X)(A)$ such that $\alpha_A(\hat{p}) = p$. From injectivity, this \hat{p} is unique.
- How do we connect this \hat{p} to the corresponding $u \in PX$ we identified? (Look at how we constructed the natural transformations α_A in our proof of baby yoneda).
- From Yoneda Lemma we know we have a function $func(p)$ defined as a family of functions $func(p)_A$, where $func(p)_A$ maps p to the component α_A . Specifically, α_A is defined to take elements of $(\natural X)(A)$, e.g. some $f : A \rightarrow X$, and map them to $(Pf)(p) = p \cdot_P f$, elements of PA . Then we know $\alpha_A(\hat{p}) = u \cdot_P \hat{p}$ and $\alpha_A(\hat{p}) = p$, meaning $p = u \cdot_P \hat{p}$!

9.3 Calculating a universal property

Let's bring all these ideas together to find an object of a category that *internalizes* a \mathcal{C} -indexed set. This is an essential application of the Yoneda lemma: sometimes it is easier to describe objects by how they behave externally, and the Yoneda lemma lets us bring those external descriptions back to objects of the category. Let's use this idea to calculate the universal property of an *equalizer*.

Definition 27: Equalizer

An equalizer of $A \begin{matrix} \xrightarrow{f} \\ \xrightarrow{g} \end{matrix} B$ is a representation for the \mathcal{C} -indexed set with action on objects:

$$\text{Eq}(f, g)(X) = \{X \xrightarrow{p} A \mid f \circ p = g \circ p\}$$

and action on morphisms works by precomposition:

$$\begin{aligned} \text{Eq}(X \xrightarrow{h} Y) : \text{Eq}(f, g)(Y) &\rightarrow \text{Eq}(f, g)(X) \\ &= (Y \xrightarrow{p} A \mid f \circ p = g \circ p) \mapsto \\ &\quad (X \xrightarrow{h} Y \xrightarrow{p} A \mid f \circ p \circ h = g \circ p \circ h) \end{aligned}$$

From this external definition, we can calculate the universal property for an internal definition of equalizer by computing a terminal object in the category of elements $\int \text{Eq}(f, g)$. Let's calculate what it is. The objects are pairs (X, p) where $X \xrightarrow{p} A$ and $f \circ p = g \circ p$ (i.e., $p \in \text{Eq}(f, g)(X)$). Morphisms from (X, p) to (Y, q) is a \mathcal{C} -morphism $X \xrightarrow{h} Y$ such that $\text{Eq}(f, g)(h)(q) = p$. Unfolding $\text{Eq}(f, g)(h) = q \circ h$, we have that $q \circ h = p$.

Now we have a clearer picture of what objects in this category of elements looks like. They are exactly this configuration of objects:

$$E \xrightarrow{h} A \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} B$$

Say we have another object in this category of elements:

$$Q \xrightarrow{h} A \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} B$$

Then, morphisms in the category of elements are morphisms $Q \xrightarrow{f} P$:

$$\begin{array}{ccc} R & \xrightarrow{h} & A \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} B \\ f \uparrow & \nearrow & \\ Q & & \end{array}$$

where the above diagram commutes. We call the terminal object in the category of elements $\int \text{Eq}(f, g)$ the equalizer. From this we can finally calculate the internal definition of an equalizer: it is the terminal object in this category of elements.

Proposition 9.3.1. Let $(E, E \xrightarrow{p} A)$ be an equalizer of $A \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} B$. Then, $\mathfrak{z}E \cong \text{Eq}(f, g)$.

Functors and natural transformations

We are ready to broaden our categorical perspective and language to begin comparing two categories with one another. We will see how some of the notions we've been encountering so far in our journey through the Yoneda lemma – in particular, the functoriality and naturality properties of \mathcal{C} -indexed sets and functions – are special cases of more general phenomena in category theory called functors and natural transformations.

10.1 Functors

A functor is a structure-preserving map between categories:

Definition 28: Functor

Given two categories \mathcal{C} and \mathcal{D} , a *functor* F from \mathcal{C} to \mathcal{D} , written $F : \mathcal{C} \rightarrow \mathcal{D}$, consists of

- An “action on objects”, which is a function sending each object X of \mathcal{C} to an object $F(X)$ of \mathcal{D}
- An “action on morphisms”, which is a function sending each morphism $X \xrightarrow{f} Y$ of \mathcal{C} to a morphism $F(X) \xrightarrow{F(f)} F(Y)$ of \mathcal{D}

such that the following *functoriality conditions* are satisfied:


- Identity preservation: for all objects X of \mathcal{C} it holds that $F(\text{id}_X) = \text{id}_{F(X)}$
- Composition preservation: for all composable pairs of morphisms $X \xrightarrow{f} Y \xrightarrow{g} Z$ of \mathcal{C} it holds that $F(g \circ f) = F(g) \circ F(f)$.

Intuitively, if there is a functor $F : \mathcal{C} \rightarrow \mathcal{D}$, then there is an “abstract copy of \mathcal{C} inside \mathcal{D} ”.

10.1.1 Examples of functors

The simplest example of a functor is one between two finite categories. Consider the following example of a functor between two finite categories \mathcal{C} and \mathcal{D} :

A good exercise: how many functors are there from \mathcal{C} to \mathcal{D} ?



The functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is drawn in red; we've only shown its actions on objects here. Written out, we have:

$$F(Y) = Z, \quad F(X) = W, \quad F(Z) = Z$$

This action on objects defines the action on morphisms. This isn't always the case; there may be more than one morphism between two objects in \mathcal{C} and \mathcal{D} , in which case, there may be a choice on which morphisms to map to. But here, it is simple, and we can conclude:

$$F(\text{id}_Y) = \text{id}_Z \quad F(\text{id}_X) = \text{id}_W \quad F(\text{id}_Z) = \text{id}_Z$$

$$F(g) = \text{id}_Z \quad F(f) = k \quad F(h) = k \quad F(g \circ f) = k$$

We need to check that the functoriality properties are satisfied. Clearly, identity preservation is satisfied. The only interesting case to check is the compositional property that $k = F(g \circ f) = F(g) \circ F(f) = \text{id}_Z \circ k = k$.

As usual, it's useful to gain an intuition for what functors look like between categories that represent preorders. They will correspond to monotone functions:

Definition 29: Monotone function
Let (X, \leq_X) and (Y, \leq_Y) be preorders. Then, a function $f : X \rightarrow Y$ is called <i>monotone</i> if, for any $x_1, x_2 \in X$ such that $x_1 \leq_X x_2$, it is the case that $f(x_1) \leq_Y f(x_2)$.

Now, let's see how functors between preorders correspond to monotone functions. Let \mathcal{C} and \mathcal{D} be categories corresponding to preorders and let $F : \mathcal{C} \Rightarrow \mathcal{D}$ be a functor between these two categories. Then, functoriality of F requires that, for any morphism $X \xrightarrow{f} Y$ in \mathcal{C} , there is a morphism $F(X) \xrightarrow{F(f)} F(Y)$. This is exactly the monotonicity requirement: the morphism f , considered order-theoretically, denotes that $X \leq Y$; similarly, the morphism $F(f)$ denotes that $F(X) \leq F(Y)$.

10.2 Natural transformations

Natural transformations are structure-preserving maps between functors:

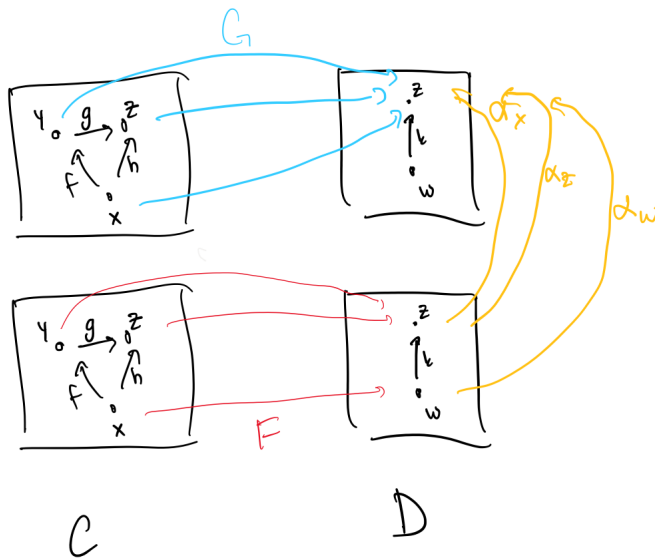
Definition 30: Natural transformation

Given two categories \mathcal{C} and \mathcal{D} , and two functors $F, G : \mathcal{C} \rightarrow \mathcal{D}$, a *natural transformation* α from F to G , written $\alpha : F \Rightarrow G$, consists of a function sending each object X of \mathcal{C} to a *morphism* $F(X) \xrightarrow{\alpha_X} G(X)$ of \mathcal{D} such that the following *naturality condition* is satisfied: for each morphism $X \xrightarrow{f} Y$ of \mathcal{C} , the following square commutes:

$$\begin{array}{ccc} FX & \xrightarrow{\alpha_X} & GX \\ Ff \downarrow & & \downarrow Gf \\ FY & \xrightarrow{\alpha_Y} & GY \end{array}$$

In other words, $\alpha_Y \circ Ff = Gf \circ \alpha_X$ for all $X \xrightarrow{f} Y$.

Let's consider two functors F and G between finite categories and a natural transformation $\alpha : F \Rightarrow G$:



Breaking the data in this figure down, we have that $\alpha_x = \text{id}_z, \alpha_z = \text{id}_z$, and $\alpha_w = k$. Now we can check the naturality square for every morphism:

- Checking for f : we need to show that $\alpha_y \circ F(f) = G(f) \circ \alpha_x$. Substituting, this is $\text{id}_z \circ k = \text{id}_z \circ k$.

10.2.1 More examples of natural transformations

Think back to \mathcal{C} -indexed functions. We drew these as squares involving two Balloon diagrams stacked together. It may help to think of the naturality condition as an abstraction of such a picture, where Balloons are replaced by objects of \mathcal{D} :

$$\begin{array}{ccc} GX & \xrightarrow{Gf} & GY \\ \alpha_X \uparrow & & \uparrow \alpha_Y \\ FX & \xrightarrow{Ff} & FY \end{array}$$

$$X \xrightarrow{f} Y$$

Definition 31: Identity natural transformation

Let $F : \mathcal{C} \rightarrow \mathcal{D}$ be a functor between two categories \mathcal{C}, \mathcal{D} . The identity natural transformation $\text{id}_F : F \Rightarrow F$ is defined by $\text{id}_{F,X} = \text{id}_{F(X)}$. Naturality is verified by the following “abstract balloon diagram”:

$$\begin{array}{ccc} FX & \xrightarrow{Ff} & FY \\ \text{id}_{FX} \uparrow & & \uparrow \text{id}_{FY} \\ FX & \xrightarrow{Ff} & FY \end{array}$$

$$X \xrightarrow{f} Y$$

Definition 32: Composition of natural transformations

Given $\alpha : F \Rightarrow G$ and $\beta : G \Rightarrow H$, the composition $\beta \circ \alpha : F \Rightarrow H$ is defined by $(\beta \circ \alpha)_X = \beta_X \circ \alpha_X$. Naturality is verified by the

following “abstract balloon diagram”:

$$\begin{array}{ccc}
 HX & \xrightarrow{Hf} & HY \\
 \beta_X \uparrow & & \uparrow \beta_Y \\
 GX & \xrightarrow{Gf} & GY \\
 \alpha_X \uparrow & & \uparrow \alpha_Y \\
 FX & \xrightarrow{Ff} & FY
 \end{array}$$

$$X \xrightarrow{f} Y$$

Examples

- Suppose you have two functors $F, G : X \rightarrow Y$ between preorders X, Y , aka two monotone functions. There is at most one natural transformation $\alpha : F \Rightarrow G$, which exists if and only if $F(x) \leq G(x)$ for all $x \in X$.
- Mind-bending example: there is an identity functor $\text{id} : \text{FinSet} \rightarrow \text{FinSet}$. What does a natural transformation $\alpha : \text{id} \Rightarrow \text{id}$ look like? Think polymorphism: intuitively, α must have type “ $\forall X. \text{FinSet}(\text{id}(X), \text{id}(X))$ ”, aka “ $\forall X. X \rightarrow X$ ”. This suggests that α is the identity, i.e., $\alpha_X = \text{id}_X$ for all objects X of FinSet . Indeed, we can verify this follows from α being natural, as follows.
 - Let X be an arbitrary object of FinSet . We will show α_X is the identity function on X .
 - By function extensionality it’s enough to show that for all $x \in X$ it holds that $\alpha_X(x) = x$.
 - Fix $x \in X$ arbitrary. We saw in a homework that x corresponds to a FinSet morphism $\hat{x} : 1 \rightarrow X$, where 1 is the terminal object of FinSet (aka the one-point set). Now, by naturality of α , the following square commutes:

$$\begin{array}{ccc}
 X & \xrightarrow{\alpha_X} & X \\
 x \uparrow & & \uparrow x \\
 1 & \xrightarrow{\alpha_1} & 1
 \end{array}$$

Because 1 is terminal, there can only be one morphism $1 \rightarrow 1$, namely the identity. Hence $\alpha_1 = \text{id}_1$. Thus this square collapses to the following triangle:

$$\begin{array}{ccc}
 X & \xrightarrow{\alpha_X} & X \\
 x \uparrow & \nearrow x & \\
 1 & &
 \end{array}$$

But now this triangle expresses the fact that $\alpha_X(x) = x$, which is exactly what was to be shown.

10.3 Functor categories

Definition 33: Functor category

Let \mathcal{C} and \mathcal{D} be categories. Suppose that

1. The collection of functors from \mathcal{C} to \mathcal{D} forms a set
2. For any two functors $F, G : \mathcal{C} \rightarrow \mathcal{D}$, the collection of natural transformations from F to G forms a set

Then there is a *functor category*, written $[F; G]$, whose

- Objects are functors \mathcal{C} to \mathcal{D}
- Morphisms from F to G are natural transformations $\alpha : F \Rightarrow G$

Identity morphism is the identity natural transformation, and composition is composition of natural transformations.

Examples

- As you can see, there are small set-theoretic nuisances with this definition. We will come back to this
- $\text{FinSet}^{\text{loop}}$ is a functor category! It's the functor category $[\mathcal{C}; \text{FinSet}]$ where \mathcal{C} is the category with one object \star , a special non-identity morphism f , and all composites one can make from this (aka f^n for $n \in \mathbb{N}$). (In other words, \mathcal{C} is the category you get from the monoid $(\mathbb{N}, +, 0)$ of natural numbers under addition.)
- $\text{FinSet}^{\bullet \rightarrow \bullet}$, the category of finite functions, is a functor category! It's $[(\bullet \rightarrow \bullet); \text{FinSet}]$ where $\bullet \rightarrow \bullet$ is the category with two objects and one non-identity arrow between them.
- As you can see, functor categories are kind of like "categories of diagrams": the objects of a functor category look like diagrams whose shape is given by the domain category. Keep this in mind; it'll come back later (limits and colimits)
- You may wonder: are \mathcal{C} -indexed sets a functor category? Hold that thought...

10.4 Full and faithful functors

We noted that a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ ensures an abstract copy of \mathcal{C} is inside \mathcal{D} . But, what if we want *more* of \mathcal{C} to be visible inside \mathcal{C} ? We can coax more structure out by enforcing that the functor does not collapse any structure.

Note that for each functor $F : \mathcal{C} \rightarrow \mathcal{D}$ and each pair of objects X, Y in \mathcal{C} , you get a function $\mathcal{C}(X, Y) \rightarrow \mathcal{D}(FX, FY)$.

- A functor is *full* if this function is surjective: for any morphism $FX \xrightarrow{\hat{f}} FY$ in \mathcal{D} , there exists a morphism $X \xrightarrow{f} Y$ in \mathcal{C} such that $\hat{f} = F(f)$.
- A functor is *faithful* if this function is injective: for any morphism for any morphisms $f, g : X \rightarrow Y$ in \mathcal{C} , if $Ff = Fg$ as morphisms $FX \rightarrow FY$ in \mathcal{D} , then $f = g$.

In particular, if a functor is *full and faithful*, then each function $\mathcal{C}(X, Y) \rightarrow \mathcal{D}(FX, FY)$ is a bijection, showing that $\mathcal{D}(FX, FY) \cong \mathcal{C}(X, Y)$ for all objects X, Y of \mathcal{C} .

10.5 The category of sets

Naive definition, paralleling Definition 2: the category *Set* consists of

- Objects: the set of all sets
- Morphisms: tuples (A, f, B) where A, B are sets and $f : A \rightarrow B$ is a function

Obviously this definition is unworkable: there is no set of all sets. We have to finally confront the set-theoretic minutiae that we have swept under the rug so far. Russell famously showed that the collection of all sets does not form a set; it is “too big”. We will adopt a way out which has become standard when doing highly categorical things: postulate the existence of a “very big set” that contains “enough sets to look like the set of all sets”. Such a set is called a *Grothendieck universe*. We won’t give the formal definition here; it’s not essential to understand the details.¹ For our purposes, all that matters is that every Grothendieck universe \mathcal{U} is a model of ZFC, hence “looks like the set of all sets”. From now on we assume the following axiom.

Axiom 10.5.1. There exists a Grothendieck universe \mathcal{U} .

The set \mathcal{U} splits the world of sets into two halves:

- On one side there are the sets X such that $X \in \mathcal{U}$. These sets are “small enough to fit into \mathcal{U} ”, so called *\mathcal{U} -small*.

¹ If you’re curious, check out the nLab page for Grothendieck universe.

- On the other side there are sets X such that $X \notin \mathcal{U}$. These sets are “too large to fit”, so called \mathcal{U} -large.

The key point is that, while we cannot define a category of *all* sets, we can define a category of \mathcal{U} -small ones.

Definition 34: the category of (\mathcal{U} -small) sets

Set is the category whose

- Objects are elements of \mathcal{U} , aka \mathcal{U} -small sets
- Morphisms from X to Y are functions $f : X \rightarrow Y$

Just as sets can be \mathcal{U} -small or \mathcal{U} -large, so too can categories. There are actually three tiers of “small”-ness for categories. First, a category can fit perfectly inside of the universe \mathcal{U} .

Definition 35

A category \mathcal{C} is \mathcal{U} -small if both of its sets of objects and morphisms are.

More generally, a category could be too large to fit inside of \mathcal{U} , but it could be the case that for any two objects, the collection of *morphisms* between those two objects does fit inside of \mathcal{U} . Intuitively, this is like saying the category is \mathcal{U} -small if you zoom in on the morphisms between two given objects; hence this condition is called being “locally” small.

Definition 36

A category \mathcal{C} is \mathcal{U} -locally-small if, for each pair of objects X, Y , the set $\mathcal{C}(X, Y)$ is \mathcal{U} -small.

Finally, a category is \mathcal{U} -large if it is neither small nor locally small.

10.5.1 Indexed set theory as a category

Now we can show that \mathcal{C} -indexed sets form a category, so long as we take Set to mean the category of \mathcal{U} -small sets, and \mathcal{C} is locally \mathcal{U} -small.

Definition 37: category of \mathcal{C} -indexed sets

Let \mathcal{C} be a locally \mathcal{U} -small category. Then \mathcal{C} -indexed sets and functions between them form a functor category $[\mathcal{C}^{\text{op}}; \text{Set}]$.

Proposition 10.5.2. Let \mathcal{C} be a locally \mathcal{U} -small category. The Yoneda embedding defines a functor $\mathfrak{y} : \mathcal{C} \rightarrow [\mathcal{C}^{\text{op}}; \text{Set}]$, and this functor is full and faithful.

Proof. Let X and Y be objects. To show that \mathfrak{y} is full and faithful, we must show that $\mathcal{C}(X, Y) \cong [\mathcal{C}^{\text{op}}; \text{Set}](\mathfrak{y}X, \mathfrak{y}Y)$.

Let X and Y be objects. Then, by specializing the baby Yoneda lemma, $\{\alpha : \mathfrak{y}X \Rightarrow \mathfrak{y}Y\} \cong \mathfrak{y}(Y)(X)$. We have immediately that $\{\alpha : \mathfrak{y}X \Rightarrow \mathfrak{y}Y\} = [\mathcal{C}^{\text{op}}; \text{Set}](\mathfrak{y}X, \mathfrak{y}Y)$ by the definition of a functor category, and that $\mathfrak{y}(Y)(X) = \mathcal{C}(X, Y)$ by the definition of \mathfrak{y} . \square

11

Monoidal Categories

11.1 Categorifying Monoids

Thanks Jialu Bao for this lecture, and Shubh Agrawal for these scribe notes.

We have been working thus far with typed PLs, specified using typing rules. We used categories to give them semantics and diagrams to depict the categories. We will now discuss PLs with resources, and use *monoidal categories* to give them semantics, and use string diagrams to depict them.

First, let us recall the definition of monoid:

Definition 38: monoid

A triple (M, \cdot, e) is a monoid if it satisfies:

1. **Associativity:** $\forall x, y, z \in M. (x \cdot y) \cdot z = x \cdot (y \cdot z)$
2. **Identity:** $\forall x \in M. e \cdot x = x = x \cdot e$


We have seen one way of categorifying monoids: construct a category with one object and a morphism for each element of the monoid. The monoid operation then corresponds to sequential composition and the monoid axioms correspond to the category axioms.

What if we instead categorify monoids by making the *objects* form a monoid and where the monoid operation corresponds to some notion of parallel composition. Such categories are called *monoidal categories*. Monoidal categories are a category \mathcal{C} together with an operation \otimes , called the tensor product, and a distinguished object I . The axioms are then encoded as a certain relationship between corresponding objects. Formally:

Definition 39: monoidal category

A category \mathcal{C} is a monoidal category if:

1. there exists a bifunctor $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$
2. there exist natural isomorphisms:
 - $\alpha_{A,B,C} : A \otimes (B \otimes C) \Rightarrow (A \otimes B) \otimes C$
 - $\lambda_A : I \otimes A \Rightarrow A$
 - $\rho_A : A \Rightarrow A \otimes I$
3. such that the following diagrams commute:




The conditions that these diagrams commute are often simply referred to as “coherence conditions.” The intuition is that the three natural isomorphisms need to interact with each other well in the sense that if we take any two paths between objects, then they are equivalent. Proving those two diagrams commute is sufficient to prove the *coherence theorem*, which makes the notion described above precise.

Let us define some of the other terms in the definition:

Definition 40: bifunctor
A functor from the product category $\mathcal{C} \times \mathcal{C}$ to \mathcal{C} .

Definition 41: natural isomorphism
natural isomorphism: a natural transformation $\eta : F \Rightarrow G$ is a natural isomorphism if for every object $X \in \mathcal{C}$, the component η_X is an isomorphism. More concretely, we have the following diagram for every $f : X \rightarrow Y$ in \mathcal{C} :



With these in hand, we can now define how to turn a monoid (M, \cdot, e) into a monoidal category:

- The objects are the elements of M
- Morphisms are only the identity morphism on each object
- The tensor product $a \otimes b$ is the monoid operation $a \cdot b$
- The tensor product on morphisms $f \otimes g$ is trivial because both f and g are identity morphisms; thus, we have $\text{id}_A \otimes \text{id}_B = \text{id}_{A \otimes B}$.
- We must show that the \otimes defines a functor. It clearly respects identity by the definition above. showing it respects composition amounts to proving the following: $(f_1 \otimes f_2) \circ (g_1 \otimes g_2) = f_1 \circ g_1 \otimes f_2 \circ g_2$. This holds because all morphisms are the identity
- The objects we need to define natural transformations between are exactly equal by the monoid axioms, so we just define the transformations to be the identity on each component.


11.2 String Diagrams

In string diagrams, we represent objects with wires and morphisms with boxes (including an identity). Sequential compositions correspond to wiring boxes together. Tensor product for objects corresponds to juxtaposing wires next to each other, and tensor products of wires is putting boxes next to each other. The monoid unit corresponds to a “phantom wire” which need not be drawn.

Concept	String Diagram
Object X	$ _x$
Morphism $f: X \rightarrow Y$	$\begin{array}{c} y \\ \boxed{f} \\ x \end{array}$
Tensor $X \otimes Y$	$ _x _y$
Tensor $f \otimes g$	$\begin{array}{c} y \\ \boxed{f} \\ x \end{array} \quad \begin{array}{c} w \\ \boxed{g} \\ z \end{array}$
Unit I	\square
Identity morphism id_x	$\begin{array}{c} x \\ \boxed{id} \\ x \end{array}$ or $ _x$

Many of the laws correspond to extremely intuitive facts about diagrams:

Isomorphism of objects	Equality of diagrams
$I \otimes X \cong$ $X \cong X \otimes I$	
$(X \otimes Y) \otimes Z$ $\cong X \otimes (Y \otimes Z)$	




In any monoidal category, we have a notion of “generating objects and morphisms,” which can’t be decomposed into the tensor product of other objects; the category is then formed by taking all the possible tensor products.

11.3 Example: Real Numbers

Consider the preorder of real numbers viewed as a category. We can define the tensor product on objects to be addition and the monoidal unit to be 0. Given two morphisms $n_1 \rightarrow n_2$, $m_1 \rightarrow m_2$, the tensor product is the morphism $n_1 + m_1 \rightarrow n_2 + m_2$, which we know exists by monotonicity of addition. Similar to the monoid example, the natural isomorphisms are the identity on each component because the corresponding objects are equal.

We can also use string diagrams to prove facts about this preorder. Here, we have diagrams representing proofs of $v + x \leq y$, $w + u \leq$

$x + z, t \leq v + w$. We can then wire the boxes together to get a proof of $t + u \leq y + z$:




11.4 Example: Cartesian Categories

If \mathcal{C} has a product \times and terminal object T , then this forms a monoidal category with $\otimes = \times$ and $I = T$.

The natural isomorphism for associativity was proven in Homework 1: $\alpha = \langle \langle \pi_1, \pi_1 \circ \pi_2 \rangle, \pi_2 \circ \pi_2 \rangle$. We also have $\lambda = \pi_1$ and $\rho = \langle \star, \text{id}_A \rangle$.

In the typed PL perspective, we had rules for operating with products and the terminal object. We could then use these to construct more complicated terms. For example, we might ask given terms $x : A \vdash M : B$ and $\langle x, y \rangle : A \times B \vdash N : C$, how can we construct a term T such that $x : A \vdash T : C$? We can do this with our term language. We can also do it in string diagrams:




Doing this requires a *copy* operation, which allows us to duplicate A .

11.5 Additional Structure

11.5.1 Symmetry






Notice that in both examples given, we have that the tensor product is commutative. The categories where this holds are called *symmetric monoidal categories*, which come equipped with an additional natural transformation $\gamma : A \otimes B \Rightarrow B \otimes A$. We get the following laws:



11.5.2 Copy and Discard

In order to encode all the same operations we expected in our STLC with pairs and unit, we needed a copy operation. The copy operation is a morphism $A \rightarrow A \otimes A$ for each A . In general, monoidal categories won't always have a copy operation. This corresponds to the intuition of monoidal categories as representing resources – if our resource can be copied, then our category would have a copier. If not, then it wouldn't.

Often, when we have a copy operation, we also have a discard operation $A \rightarrow I$. These operations will satisfy the following rules:



Limits and colimits

So far, we've been building a "categorical dictionary" of structure that categories might have that makes them useful models for programming languages. Categorical products let us model products in programming languages, exponential objects let us model higher-order functions. Limits and colimits take these ideas to their extreme: we will see that a category that has limits and colimits has a rich vocabulary of universal constructions of which products and exponentials are a special case. The utility of limits and colimits is in their generality, both in the kinds of constructions they give you and also how they are characterized. We will see that limits are *preserved* or *created* by the presence of certain functors, enabling us to characterize a large amount of categorical structure with relatively little work.

As an example of how limits generalize products, let

$$\text{Point} = \{x : \mathbb{R}, y : \mathbb{R}, z : \mathbb{R}\}$$

This leads to the idea of indexed products.

Definition 42: Indexed product


Let $(X_i)_{i \in I}$ be an I -indexed family of objects of a category \mathcal{C} . An *indexed product* of $(X_i)_{i \in I}$ consists of an object P and a family of morphisms $(\pi_i : P \rightarrow X_i)_{i \in I}$ satisfying the following universal property: for any object Y and any family of morphisms $(f_i : Y \rightarrow X_i)_{i \in I}$, there exists a unique morphism $\langle f_i \rangle_{i \in I} : Y \rightarrow P$ such that $\pi_i \circ \langle f_i \rangle_{i \in I} = f_i$ for all i in I .

While records are typically finite, the definition of indexed product works just as well for infinite products. For example,


Proposition 12.0.1. 1 is the indexed product of $(k)_{k \in \mathbb{N}}$ in the category of natural numbers under divisibility.

More generally, indexed products in preorders are like infima.

Indexed products are in some sense the “largest” way to map into a tuple of objects. For example, the product for Point is the “largest” way of filling in the question marks in this diagram:



Similarly, 1 is the “largest” way to fill in the question marks in this diagram:




All of these diagrams are “discrete”: they consist only of objects, with no morphisms connecting them. More generally, we can ask if it is possible to find the largest way of mapping into an arbitrary diagram consisting of both objects and morphisms. This is called a *limit*.

To be more precise about this, it is useful to give a formal definition of what a diagram is:

Definition 43: Diagram

For two categories \mathcal{I} and \mathcal{C} , a *diagram of shape \mathcal{I}* is a functor $F : \mathcal{I} \rightarrow \mathcal{C}$. The category \mathcal{I} is called the *shape of the diagram*.

It is typical for the shape of a diagram to be a simple category with only a few objects and morphisms. For example, we can consider $\mathcal{I} = \bullet \rightarrow \bullet$; then, a diagram of shape \mathcal{I} in some category \mathcal{C} would pick out a pair of two objects in \mathcal{C} . Here is a simple example of a diagram:



The essential idea of a limit is that it is the largest way of “mapping into a diagram”. The way to formalize this is via the notion of a *cone*:

Definition 44: Cone

Let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram of shape \mathcal{I} in a category \mathcal{C} . A *cone* over D is a pair (P, p) where P is an object of \mathcal{C} and p is a family of morphisms $(p_i : P \rightarrow D(i))_{i \in I}$ such that $D(f) \circ p_i = p_j$ for all morphisms $f : i \rightarrow j$ of \mathcal{I} .


Let's understand this definition for the simple diagrams of shape $\mathcal{I} = \bullet \rightarrow \bullet$ before we consider more more interesting shapes with morphisms in them. In the above example diagram, the tuple (P, p) is a cone where $p_L = f$ and $p_R = g$.

Cones form a category:

Definition 45: Morphism of cones


Let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram of shape \mathcal{I} in a category \mathcal{C} . Let (P, p) and (Q, q) be two cones over D . A *morphism of cones* from (P, p) to (Q, q) is a morphism $f : P \rightarrow Q$ such that $p_i = q_i \circ f$ for all $i \in \mathcal{I}$.

A morphism of cones can be visualized as a kind of 3d prism:




At the bottom of this prism, we have drawn a piece of the diagram D . The “bottom left face” of the prism is the cone (P, p) ; the “bottom right face” is the cone (Q, q) . A morphism of cones from (P, p) to (Q, q) is a morphism $f : P \rightarrow Q$ as shown that makes all the triangles in the prism commute.

Looking at specific choices of diagram D recovers the kinds of figures we were drawing. For example, suppose D is the discrete diagram on two objects X and Y . We get that a cone (P, p) is a diagram of shape




and a morphism from one such cone (P, p) to another (Q, q) is a

diagram of shape:



Now imagine taking the object P and “dragging it up and to the right”, until it’s in the same plane as the triangle involving Q, X, Y :



Hopefully this diagram looks familiar: up to a rotation, it is the kind of diagram we were drawing when examining the universal property of product.

Together, cones over a diagram D and morphisms between them form a category:

Definition 46: category of cones

Let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram in a category \mathcal{C} of shape \mathcal{I} . The category of cones over D , written $\text{Cone}(D)$, is the category whose objects are cones over D and whose morphisms are morphisms of cones. (Because cone morphisms are simply morphisms of \mathcal{C} satisfying a certain property, the definition of composition for cone morphisms and of the identity cone morphism are inherited from composition and identity in \mathcal{C} .)

The category of cones is designed so that it captures all the ways of mapping into a diagram. Recall our motivation that a limit ought to be the biggest way to map into a diagram. This hints at our final definition: the “biggest” object in a category is the terminal object,

Definition 47: Limit

Let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram of shape \mathcal{I} in a category \mathcal{C} . A *limit* of D is a terminal object in $\text{Cone}(D)$.

Now we can see how products are terminal in the categories of cones of shape $\mathcal{I} = \bullet \quad \bullet$. In this category, objects look like this:


12.1 Examples of limits

We have seen two examples of limits already, in the homework.

- Given two morphisms $X \xrightarrow{f} Z \xleftarrow{g} Y$, the *pullback* is the largest way of mapping into f, g , in other words the largest way of filling in the question marks in this diagram:

$$\begin{array}{ccc} ? & \xrightarrow{?} & Y \\ ? \downarrow & & \downarrow g \\ X & \xrightarrow{f} & Z \end{array}$$


Explicitly, a pullback is a 3-tuple (W, p, q) where $p : W \rightarrow X$ and $q : W \rightarrow Y$ and $fp = gq$ satisfying the following universal property: for any other 3-tuple (W', p', q') where $p' : W' \rightarrow X$ and $q' : W' \rightarrow Y$ and $fp' = gq'$, there exists a unique morphism $w : W' \rightarrow W$ such that $pw = p'$ and $qw = q'$.




- Given two morphisms $X \xrightarrow{f} Y \xleftarrow{g} Y$, their *equalizer* is the largest way of mapping into f, g in other words the largest way of filling in the question marks in this diagram:

$$\begin{array}{ccc} & ? & \\ ? \swarrow & & \searrow ? \\ X & \xrightarrow{f} & Y \\ & \xrightarrow{g} & \end{array}$$

Explicitly, the equalizer is a 3-tuple (W, p, q) where $p : W \rightarrow X$ and $q : W \rightarrow Y$ and $fp = q$ and $gp = q$ satisfying the following universal property: for any other 3-tuple (W', p', q') where $p' : W' \rightarrow X$ and $q' : W' \rightarrow Y$ and $fp' = q'$ and $gp' = q'$ there exists a unique $w : W' \rightarrow W$ such that $pw = p'$ and $qw = q'$.




This definition can be simplified a bit: notice that in the 3-tuple (W, p, q) , the morphism q is forced to be fp . So it is equivalent to say that an equalizer is a 2-tuple (W, p) such that $fp = gp$, which satisfies the universal property that for any other (W', p') with $f p' = g p'$, there exists a unique $w : W' \rightarrow W$ such that $w p = p'$ and $w q = q'$.



The equalizer looks like a refinement type:

$$\text{equalizer of } X \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} Y \quad \rightsquigarrow \quad \{x : X \mid f(x) = g(x)\}$$

The pullback looks like a “fiber product”:



There is an equivalent way of defining cones, that requires a bit more machinery.

Definition 48
<p>Given an object X of a category \mathcal{C}, and an arbitrary category \mathcal{I}, the <i>constant functor at X from \mathcal{I} to \mathcal{C}</i>, written $\Delta(X) : \mathcal{I} \rightarrow \mathcal{C}$, is the functor defined by $\Delta(X)(i) = X$ on objects i of \mathcal{I} and by $\Delta(X)(i \xrightarrow{f} j) = \text{id}_X$ on morphisms.</p>

Proposition 12.1.1. A cone over $D : \mathcal{I} \rightarrow \mathcal{C}$ is equivalent to a pair (P, α) where P is an object of \mathcal{C} and α is a natural transformation $\Delta(P) \Rightarrow D$.

Proof. Unwinding definitions, the natural transformation α is a \mathcal{I} -indexed family of morphisms $(\alpha_i : \Delta(P)(i) \rightarrow D(i))_{i \in \mathcal{I}}$ satisfying the

following naturality square for all morphisms $i \xrightarrow{f} j$ of \mathcal{I} :

$$\begin{array}{ccc} \Delta(P)(i) & \xrightarrow{\alpha_i} & D(i) \\ \Delta(P)(f) \downarrow & & \downarrow D(f) \\ \Delta(P)(j) & \xrightarrow{\alpha_j} & D(j) \end{array}$$

But by definition of $\Delta(P)$, it holds that $\Delta(P)(i) = \Delta(P)(j) = P$ and $\Delta(P)(f) = \text{id}_P$. Hence the left arrow in this commutative square is the identity, and the square can be collapsed into the following triangle which has to commute for all $i \xrightarrow{f} j$ in \mathcal{I} :

$$\begin{array}{ccc} & & D(i) \\ & \nearrow \alpha_i & \downarrow D(f) \\ P & & \\ & \searrow \alpha_j & \\ & & D(j) \end{array}$$

Hence the natural transformation $\alpha : \Delta(P) \Rightarrow D$ amounts to a family of morphisms $(\alpha_i : P \rightarrow D(i))_{i \in \mathcal{I}}$ satisfying the property that $D(f) \circ \alpha_i = \alpha_j$ for all $i \xrightarrow{f} j$ in \mathcal{I} . This is precisely the property of being a cone over D . \square

Last time we saw by the Yoneda lemma that a representation for a \mathcal{C} -indexed set P is equivalent to a terminal object in $\int P$, the category of elements of P . This fact provides an interesting equivalent definition of limit:

Proposition 12.1.2. Let \mathcal{I} be a \mathcal{U} -small category. A limit of a diagram $D : \mathcal{I} \rightarrow \mathcal{C}$ is equivalent to a representation for the \mathcal{C} -indexed set P defined by

$$P(X) = \text{the set of natural transformations } \Delta(X) \Rightarrow D.$$

Proof. The \mathcal{C} -indexed set P is designed so that the category of cones over D is equivalent to the category of elements of P . Hence, to have a terminal cone is to have a terminal object in this category of elements, which is equivalent to having a representation for P . \square

12.1.1 Some important facts

Definition 49: Limit preservation

Let \mathcal{C} and \mathcal{D} be categories, let $F : \mathcal{C} \rightarrow \mathcal{D}$, and let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram in \mathcal{C} . Suppose D has a limit (P, p) in \mathcal{C} .

Then, the functor F *preserves* this limit if the diagram $\tilde{D} = F \circ D$ in \mathcal{D} has limit $(F(P), (F(p_i) : P \rightarrow D(i))_{i \in \mathcal{I}})$.

Proposition 12.1.3. The Yoneda embedding preserves limits.

We say a functor *reflects* limits if the converse of Definition 49 holds.

Definition 50: Limit reflection

Let \mathcal{C} and \mathcal{D} be categories, let $F : \mathcal{C} \rightarrow \mathcal{D}$, and let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram in \mathcal{C} . Let (P, p) be a cone of shape D in \mathcal{C} , and suppose the $\tilde{D} = F \circ D$ has limit $(F(P), (F(p_i) : P \rightarrow D(i))_{i \in \mathcal{I}})$ in \mathcal{D} . We say this limit is *reflected* into \mathcal{C} if the cone (P, p) is limiting in \mathcal{C} .

Proposition 12.1.4. Any full and faithful functor reflects limits.

12.2 Every limit is the equalizer of a product

We've seen two important examples of simple diagrams: the discrete diagram where there are no non-identity morphisms, and the equalizer diagram where there are exactly two morphisms between exactly two objects. These are both simple enough on their own, but what is interesting and powerful is that *all limits can be built out of these two simple diagrams*:

Theorem 12.2.1. Every limit is an equalizer of an indexed product.

Consequently, a category with indexed products and equalizers has all limits.

12.3 Colimits

The essence of colimits is essentially to flip the direction of all the arrows of limits. This seems simple at first, but surprisingly, colimits behave quite differently than limits. Let's quickly state the relevant definitions:

Definition 51: Cocone

Let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram of shape \mathcal{I} in a category \mathcal{C} . A *cocone* over D is a pair (P, p) where P is an object of \mathcal{C} and p is a family of morphisms $(p_i : D(i) \rightarrow P)_{i \in \mathcal{I}}$ such that $p_i \circ D(f) = p_j$ for all morphisms $f : i \rightarrow j$ of \mathcal{I} .

Definition 52: Morphism of cocones

Let $D : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram of shape \mathcal{I} in a category \mathcal{C} . Let (P, p) and (Q, q) be two cocones over D . A *morphism of cocones* from (P, p) to (Q, q) is a morphism $f : P \rightarrow Q$ such that $f \circ p_i = q_i$ for all $i \in \mathcal{I}$.

Then, the category of cocones for a diagram \mathcal{D} is written $\text{CoCone}(\mathcal{D})$, analogous to the case for the category of cones $\text{Cone}(\mathcal{D})$. Now we can state the definition of colimit:

Definition 53: Colimit

Let $\mathcal{D} : \mathcal{I} \rightarrow \mathcal{C}$ be a diagram of shape \mathcal{I} in \mathcal{C} . A *colimit* of \mathcal{D} is a cocone (P, p) that is initial in $\text{CoCone}(\mathcal{D})$.

Two examples in finite sets:

- coproduct, end up “tagged unions”
- coequalizer is unfortunately much trickier. It ends up being a “generalized quotient”

Adjunctions

- So far we've mostly focussed on characterizing properties of an individual category. Now let's broaden our focus to discuss in more detail relationships between categories.
- First, let's study the strongest possible relationship between two categories: does it mean for two categories to be equivalent? Following intuition from set isomorphism, we will want to design two functors between the categories that behave like inverses.
- For this definition we will need a way of talking equivalence of functors:

Definition 54: Natural isomorphism of functors

Let \mathcal{C} and \mathcal{D} be categories and $F, G : \mathcal{C} \rightarrow \mathcal{D}$ be functors. Then, F and G are *naturally isomorphic*, written $F \cong G$, if there is a natural transformation $\alpha : F \Rightarrow G$ whose components are isomorphisms.

- Now, to give an equivalence of categories, we must show that there exists a pair of functors F and G between them whose compositions are naturally isomorphic to the identity functors on each category:

Definition 55: Equivalence of categories

Let \mathcal{C} and \mathcal{D} be categories. Then, \mathcal{C} is equivalent to \mathcal{D} , written $\mathcal{C} \cong \mathcal{D}$, if there are two functors $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$ satisfying $F \circ G \cong \text{id}_{\mathcal{D}}$ and $G \circ F \cong \text{id}_{\mathcal{C}}$, where $\text{id}_{\mathcal{C}}$ and $\text{id}_{\mathcal{D}}$ are the identity functors on \mathcal{C} and \mathcal{D} respectively.

- There are some surprising categorical equivalences that are worth thinking about:

Proposition 13.0.1. The category of finite pointed sets FinSet^* is equivalent to the category of finite sets and partial functions $\text{FinSet}_{\text{par}}$.

The intuition behind this equivalence is interesting. First let's identify the functors that form the equivalence. Let $F : \text{FinSet}_{\text{par}} \rightarrow \text{FinSet}^*$. The action on objects is straightforward: $F(X) = (X \uplus \{\star\}, \star)$ i.e., F sends a set X to a pointed set with X extended by some element \star with \star as the point. The action on morphisms is interesting. Let $f : X \rightarrow Y$ be a partial function. Then, we define:

$$\begin{aligned} F(f) : (X \uplus \{\star\}, \star) &\rightarrow (Y \uplus \{\star\}, \star) \\ &= x \mapsto \begin{cases} \star & \text{if } f(x) \text{ undefined} \\ f(x) & \text{otherwise.} \end{cases} \end{aligned}$$

Intuitively, the \star stands in for the subset of X on which f is undefined. The inverse G is also straightforward: its action on objects is to forget remove point, $G((X, \star)) = X \setminus \star$. Then, its action on morphisms map the point to undefined:

$$\begin{aligned} G(g) : X \setminus \star &\rightarrow Y \setminus \star \\ &= x \mapsto \begin{cases} \perp & \text{if } g(x) = \star \\ g(x) & \text{otherwise} \end{cases} \end{aligned}$$

Now, show that there is a natural isomorphism $\alpha : F \circ G \Rightarrow \text{id}_{\text{FinSet}^*}$. This is not too hard and a good exercise. It requires showing that there is a FinSet^* isomorphism $(X, a) \cong ((X \setminus a) \uplus \star, \star)$.

- Categorical equivalences are very powerful when they arise but the condition is often too strong. A natural weakening of this notion will be an *adjunction*, which will be pairs of functors between categories are a form of “weak equivalence”.
- The natural notion of weakening here will be as follows. Let \mathcal{C} and \mathcal{D} be categories and $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$ be functors. Then, instead of requiring that round-trips of F and G be naturally isomorphic identity functors, an adjunction will require instead that there simply exist some pair of natural transformations:


$$\eta : \text{id}_{\mathcal{C}} \Rightarrow G \circ F \quad \epsilon : F \circ G \Rightarrow \text{id}_{\mathcal{D}} \quad (13.1)$$

The natural transformation η is called the **unit** of the adjunction and ϵ is called the **counit**, the functor F is called the **left adjoint**, and the functor G the **right adjoint**.

- There is one more coherence requirement we place on these natural transformations in order for them to form an adjunction. But, before we do that, let's examine this notion of notion of "weak equivalence" in a simple preorder category to understand what it means before we give the full definition of adjoint functors for general categories.

13.1 Adjunctions for preorder categories: Galois connections

- As usual it's very useful to study new category theory definitions in pre-order categories to gain order-theoretic intuition about them before generalizing to the general setting where there can be more than one morphism between any two objects.
- Let $(X, \leq_X), (Y, \leq_Y)$ be preorders. Recall that a *monotone map* is a function $f : X \rightarrow Y$ that respects the preorder, i.e., if $x_1 \leq_X x_2$ then $f(x_1) \leq_Y f(x_2)$. Recall that functors between preorder categories are exactly monotone maps.
- Let's interpret the definition of adjoint in this setting. First, let's give two example functors between two preorder categories:



Here we've drawn $F : C \rightarrow D$ in blue and $G : D \rightarrow C$ in red.

- Let's check if these functors satisfy the adjoint requirement as we've seen it so far. First, let's try to find a unit by finding a natural transformation $\eta : \text{id}_C \Rightarrow G \circ F$. Then, η is a C -indexed family of morphisms in C :

$$\eta_1 = 1 \rightarrow 2 \quad \eta_2 = \text{id}_2 \quad \eta_3 = \text{id}_3$$

We can also build the counit, which is a D -indexed family of morphisms in D :

$$\epsilon_A = \text{id}_A \quad \epsilon_B = \text{id}_B \quad \epsilon_C = B \rightarrow C$$

Intuitively, these two natural transformations are characterizing what happens after round trips of following F and G . The unit η says that, after following a round trip of F and then G , we must end up "in a higher place" than we started; dually, the counit ϵ is saying that after a round trip of G then F we must end up "in a lower place" than we started.

- Now we can reverse engineer an order-theoretic understanding for the unit and the counit of the adjunctions F and G . This relationship has a name in order theory:

Definition 56: Galois connection

Let (X, \leq_X) and (Y, \leq_Y) be preorders and let $\alpha : X \rightarrow Y$ and $\gamma : Y \rightarrow X$ be monotone functions. Then, α and γ form a *Galois connection* if, for all $x \in X$ and $y \in Y$:

$$x \leq \gamma(\alpha(x)) \quad \text{and} \quad \alpha(\gamma(y)) \leq y$$

- Pause for a moment to absorb how close a Galois connection is to an inverse. If γ and α above were inverses, then indeed this would be an inverse.
- An equivalent definition to the above is that, for all $x \in X, y \in Y$, it holds that:

$$\alpha(x) \leq y \text{ if and only if } x \leq \gamma(y) \quad (13.2)$$

Exercise: Why is this equivalence true?

- A programming languages example: Galois connections are important in defining abstract interpreters. The map α is called the *abstraction function* and γ is called the *concretization function*, and α is left adjoint to γ .

13.1.1 Adjoints preserve structure

- An important property of adjoints is that left adjoints preserve colimits and right adjoints preserve limits. This property is of course true for Galois connections, so let's see what it means in that setting first:

Theorem 13.1.1. Let (X, \leq) and (Y, \leq) be complete partial orders (i.e., they both have all meets and joins) and let $\alpha : X \rightarrow Y$ and $\gamma : Y \rightarrow X$ be monotone functions where α is left adjoint to γ . Then, for any $x_1, x_2 \in X$ we have that:

$$\alpha(x_1 \sqcup x_2) = \alpha(x_1) \sqcup \alpha(x_2) \quad (13.3)$$

and, for any $y_1, y_2 \in Y$ we have that:

$$\gamma(y_1 \sqcap y_2) = \gamma(y_1) \sqcap \gamma(y_2) \quad (13.4)$$

Proof. Let's prove that the left adjoint preserves meets; the other one is a good exercise. By monotonicity, $\alpha(x_1) \sqcup \alpha(x_2) \leq \alpha(x_1 \sqcup x_2)$.

So, we need to show $\alpha(x_1 \sqcup x_2) \leq \alpha(x_1) \sqcup \alpha(x_2)$. We will use Eq. (13.2) to show this. Let $m = x_1 \sqcup x_2$ and $m^* = \alpha(x_1) \sqcup \alpha(x_2)$. Then,

$$\alpha(m) \leq m^* \iff m \leq \gamma(m^*)$$

Proceeding on the right hand side, we use the fact that γ is monotone to build an intermediate inequality:


$$\gamma(\alpha(x_1)) \sqcup \gamma(\alpha(x_2)) \leq \gamma(\alpha(x_1) \sqcup \alpha(x_2))$$

So, if we can show $x_1 \sqcup x_2 \leq \gamma(\alpha(x_1)) \sqcup \gamma(\alpha(x_2))$, we are done. This follows directly from adjoint closure: $x \leq \gamma(\alpha(x))$ for all x . □


13.2 Adjunctions for preorders, continued

Notes from John's lecture

- Recall the Galois connection example:



- We defined a pair of functions $\alpha : \mathcal{C} \rightarrow \mathcal{A}$ and $\gamma : \mathcal{D} \rightarrow \mathcal{C}$, where α is left-adjoint to γ . We defined an example γ as follows (shown in red):




- To get to the general definition of preorder, we will study the properties of α and γ . We will want our general definition of adjoint to preserve *all* of this structure
- What is the property we want α to have, given this definition of γ ? In some sense, it is *the most precise abstraction* would can make.
- For example, we can define a really bad α that is very imprecise. For instance, we could let $\alpha(_) = ?$. But if we did this, then α would *not be an adjoint*.
- There is a collection of interesting properties we want α to have so that it is in some sense “most precise”:
 1. $\alpha(c)$ is the smallest $a \in \mathcal{A}$ such that $c \leq \gamma(a)$.
 2. $\gamma(a)$ is the largest $c \in \mathcal{C}$ such that $\alpha(c) \leq a$.
 3. For all $a \in \mathcal{A}$ and $c \in \mathcal{C}$, $\alpha(c) \leq a$ if and only if $c \leq \gamma(a)$.
- These fun facts are more than just fun facts: they are in fact *equivalent definitions* of adjoints (at least for preorders).

Theorem 13.2.1. Let $\mathcal{C} \begin{matrix} \xrightarrow{\alpha} \\ \xleftarrow{\gamma} \end{matrix} \mathcal{A}$ be a pair of functors. Then the following are equivalent:

1. $\alpha \dashv \gamma$
2. For any $c \in \mathcal{C}$, $\alpha(c)$ is the smallest $a \in \mathcal{A}$ such that $c \leq \gamma(a)$
3. For any $a \in \mathcal{A}$, $\gamma(a)$ is the largest $c \in \mathcal{C}$ such that $\alpha(c) \leq a$.
4. For all $a \in \mathcal{A}$ and $c \in \mathcal{C}$, $\alpha(c) \leq a$ if and only if $c \leq \gamma(a)$.

- This is going to be a big proof. First we'll prove that (1) implies (2) and (3). Then we'll prove that (2) and (3) both imply (4). Finally, we'll prove that (4) implies (1), giving the following strongly connected graph:



- (1) implies (3): Suppose $\alpha \dashv \gamma$. Then by definition $\alpha(\gamma(a)) \leq a$. So it only remains to show that $\gamma(a)$ is the largest c satisfying $\alpha(c) \leq a$. Suppose $c \in \mathcal{C}$ such that $\alpha(c) \leq a$. Need to show that $c \leq \gamma(a)$. By γ monotone,

$$\gamma(\alpha(c)) \leq \gamma(a)$$

By definition of adjoint, $c \leq \gamma(\alpha(c))$, which completes the proof.

- (1) implies (2): analogous to the previous case.
- Proof that (2) implies (4).

Suppose that for any $c \in \mathcal{C}$, $\alpha(c)$ is the smallest $a \in \mathcal{A}$ such that $c \leq \gamma(a)$. Want to show that for all $a \in \mathcal{A}$, $c \in \mathcal{C}$, $\alpha(c) \leq a \iff c \leq \gamma(a)$.

(\Rightarrow): Have that $\alpha(c) \leq a$, need to show that $c \leq \gamma(a)$. By monotonicity, we have that $\gamma(\alpha(c)) \leq \gamma(a)$. By assumption, $c \leq \gamma(\alpha(c))$.¹

(\Leftarrow): Have $c \leq \gamma(a)$. Want $\alpha(c) \leq a$. Since $\alpha(c)$ is the smallest a' such that $c \leq \gamma(a')$, and a is one such a' , it must be that $\alpha(c) \leq a$.

- Proof that (4) implies (1). If for all $a \in \mathcal{A}$, $c \in \mathcal{C}$ it holds that $\alpha(c) \leq a$ iff $c \leq \gamma(a)$, then for all $a \in \mathcal{A}$, $c \in \mathcal{C}$ it holds that $\alpha(\gamma(a)) \leq a$ and $c \leq \gamma(\alpha(c))$.

First, show $\alpha(\gamma(a)) \leq a$. By assumption this is true if and only if $\gamma(a) \leq \gamma(a)$, which vacuously holds. The other side is identical.

- Define some helpful notation: the sets $c \downarrow \gamma = \{a \mid c \leq \gamma(a)\}$ and $\alpha \downarrow a = \{c \mid \alpha(c) \leq a\}$. Then the previous theorem can be re-expressed as

$$\begin{aligned}
 & \alpha \dashv \gamma \\
 \iff & (\forall c. \alpha(c) = \min(c \downarrow \gamma)) \\
 \iff & (\forall a. \gamma(c) = \max(\alpha \downarrow a)) \\
 \iff & (\forall a c. \alpha(c) \leq a \iff c \leq \gamma(a)).
 \end{aligned}$$

¹ Note: cannot appeal to adjoint roundtrip property here. This is the assumption in equivalent condition (2).

13.3 Generalizing beyond preorders

- We want a general definition of adjoints that generalizes the above four-way equivalence.
- To get there, we will first become more formal about this downarrow.
- Let \mathcal{C} and \mathcal{A} be categories. Then, for an object $c \in \mathcal{C}$, the category $c \downarrow G$ is a category whose objects are morphisms $c \rightarrow G(a)$ for each object $a \in \mathcal{A}$. Morphisms are commuting triangles:

$$\begin{array}{ccc}
 c & \xrightarrow{f} & G(a) \\
 & \searrow^{f'} & \downarrow G(g) \\
 & & G(a')
 \end{array}$$

where $c \xrightarrow{f} G(a)$ and $c \xrightarrow{f'} G(a')$ are two objects.

- Now that we have a category we have a notion of “smallest element” that generalizes categorically. This means: the categorically generalization of property (2) will be that for all objects $c \in \mathcal{C}$, the category $c \downarrow G$ has an initial object.

This is a *definition of adjoint*.

- Generalizing (3), we need to build a category $F \downarrow a$. This category has objects $F(c) \xrightarrow{f} a$ for objects c and morphisms f . Then, a morphism in $F \downarrow a$ is again a triangle:

$$\begin{array}{ccc}
 F(c) & \xrightarrow{f} & a \\
 \uparrow & \nearrow^{f'} & \\
 F(c') & &
 \end{array}$$

- Let $\mathcal{C} \begin{array}{c} \xrightarrow{F} \\ \xleftarrow{G} \end{array} \mathcal{A}$

Finally generalizing (1) is a bit tricky. Let $\eta : \text{id} \Rightarrow G \circ F$ and $\varepsilon : F \circ G \Rightarrow \text{id}$. Then, we have *coherence conditions*²

- The final generalization is (4). Categorically, this says for all $c \in \mathcal{C}$ and $a \in \mathcal{A}$, $\mathcal{A}(F(c), a) \cong \mathcal{C}(c, G(a))$ that is natural in a and c .

What does “natural in a and c ” mean? Well, we need two functors with compatible domains. Intuition: take all the things that the definitions says are natural and “poke a hole in them”. Here are

This is an instance of a *comma category*, and explains the notation we chose for the sets $c \downarrow \gamma$ and $\alpha \downarrow a$ in the previous section.

² The idea of a coherence condition is essentially “equations you add to your definitions when you generalize from preorders to categories”. Here, we’ve introduced extra ways to get between \mathcal{C} and \mathcal{A} via the units. This could add lots of extra morphisms to our category if we aren’t careful. So, the coherence conditions restrict this proliferation of morphisms. It’s useful to think about what happens without these coherence conditions; the easiest is to see that initiality is violated in the category $c \downarrow G$.

the holes: $\mathcal{A}(F(-_1), -_2) : \mathcal{C}^{\text{op}} \times \mathcal{A} \rightarrow \text{Set}$ and $\mathcal{C}(-_1, G(-_2)) : \mathcal{C}^{\text{op}} \times \mathcal{A} \rightarrow \text{Set}$. Then, the definition says there must be a natural isomorphism between these two functors.

- Packaging all these definitions up:

Definition 57: Adjunction

Let $\mathcal{C} \begin{matrix} \xrightarrow{F} \\ \xleftarrow{G} \end{matrix} \mathcal{D}$ be a pair of functors. Then, the functor F is *left adjoint* to G , written $F \dashv G$, if any of the following equivalent definitions hold:

- *Unit & co-unit*: There is a pair of natural transformations $\eta : \text{id}_{\mathcal{C}} \Rightarrow G \circ F$ and $\epsilon : F \circ G \Rightarrow \text{id}_{\mathcal{D}}$ that satisfy the *triangle identities*:

$$\begin{array}{ccc}
 F & \xrightarrow{F\eta} & FGF \\
 \searrow 1_F & & \downarrow \epsilon F \\
 & & F
 \end{array}
 \qquad
 \begin{array}{ccc}
 G & \xrightarrow{G\eta} & GFG \\
 \searrow 1_G & & \downarrow \epsilon G \\
 & & G
 \end{array}$$

- *Natural isomorphism of hom-sets*: For any objects $c \in \mathcal{C}$ and $d \in \mathcal{D}$ there is an isomorphism:

$$\mathcal{D}(Fc, d) \cong \mathcal{C}(c, Gd)$$

that is *natural in c and d*, i.e., there is a natural isomorphism between the two functors:

$$\begin{array}{ccc}
 & & \mathcal{D}(F-, -) \\
 \mathcal{C}^{\text{op}} \times \mathcal{D} & \xrightarrow{\quad} & \text{Set} \\
 & & \mathcal{C}(-, G-)
 \end{array}$$

- *Initiality condition*: For every object $c \in \mathcal{C}$, the category $c \downarrow G$ has an initial object.
- *Terminality condition*: For every object $d \in \mathcal{D}$, the category $F \downarrow d$ has a terminal object.

13.4 Examples

- A category \mathcal{C} has terminal objects if and only if there is a functor $F : \mathcal{C} \rightarrow 1$ (1 is the 1-object category) that has a right adjoint.
-

14

Elementary topoi

- A *topos* is a category whose internal language is like set theory. If you like, you can think of it as a “#lang for math”: it gives a reinterpretation internal to some category of all the mathematical symbols like $\wedge, \vee, \exists, \forall, \Rightarrow$, set-builder notation, etc.
- Many of the categories we’ve been discussing so far – presheaf categories, the category of G -sets, the category of finite transition systems – are examples of toposes.
- Today we’ll discuss elementary toposes, which are toposes whose internal language is intuitionistic higher order logic. This means that the law of excluded middle and the axiom of choice may not hold in these topoi (i.e., logical statements like $\neg\neg\varphi$ may not entail φ)
- To get there, we will study the critical set-theoretic relationship of a *subset*, which is a special relationship between two objects that says one object a *part* of another. Once we have the categorification of subset in hand, we will use it to define other logical connectives.

This chapter is heavily based on [21] chapters 32 and 33. This chapter is work-in-progress.

14.1 Parts of objects

- Recall the definition of a monomorphism, which is the categorification of an inclusion map (or injective function): from the homework:

Definition 58: Monomorphism

Let \mathcal{C} be a category. A morphism $X \xrightarrow{f} Y$ is a *monomorphism* if for any morphisms $Z \begin{matrix} \xrightarrow{g_1} \\ \xrightarrow{g_2} \end{matrix} X$, if $f \circ g_1 = f \circ g_2$ then $g_1 = g_2$

g2. This property is called *left cancellation* of f . Monomorphisms are typically drawn with a hook arrow like $X \hookrightarrow Y$.

- In the homework we saw that, in the category of finite sets, a function is injective if and only if it is a monomorphism. Injective maps identify subsets, so monomorphisms will identify subobjects.
- Monomorphisms give us a natural notion of a categorification of subset, which will be called a *part of an object*:

Definition 59: Part of an object

Let \mathcal{C} be a category and X be an object in \mathcal{C} . Then, Y is a *part of* X if there is a monomorphism $Y \hookrightarrow X$.

14.2 Subobject classifier

- The subobject classifier is the categorification of a characteristic function.

Definition 60: Characteristic function

Let S be a set and $A \subseteq S$. Then, the *characteristic function* of A , written $\chi_A : A \rightarrow \{\top, \perp\}$, is the function:

$$\chi_A = x \mapsto \begin{cases} \top & \text{if } x \in A \\ \perp & \text{otherwise.} \end{cases}$$

- Intuitively, characteristic functions ought to be in bijection with the subsets they are defined over. We can see this as follows. Suppose we have the following diagram in the category of finite sets:

$$\begin{array}{ccc} & \{\star\} & \\ & \downarrow \text{true} & \\ S & \xrightarrow{\chi_A} & \{\top, \perp\} \end{array}$$

Define $\text{true} : \star \mapsto \top$. Then, there is a set B that is the pullback along the morphism $! : B \rightarrow \star$ and χ_A , and this set B must be isomorphic to A .

- This setup lends itself naturally to a categorical definition:

Definition 61: Subobject classifier

Let \mathcal{C} be a category with terminal objects and Ω be an object in \mathcal{C} . Then, a morphism $true : \mathbf{1} \rightarrow \Omega$ is called a *subobject classifier* of \mathcal{C} if for every monomorphism $j : U \hookrightarrow X$ there is a unique morphism $\chi_j : X \rightarrow \Omega$ called the *classifying morphism for the subobject represented by j* such that the following diagram:

$$\begin{array}{ccc} U & \xrightarrow{!} & \mathbf{1} \\ \downarrow j & & \downarrow true \\ X & \xrightarrow{\chi_j} & \Omega \end{array}$$

is a pullback along χ_j and $!$. In this situation, the object Ω is called the *object of truth values*.

- In the category of finite sets, the two-element set $\{\top, \perp\}$ is the object of truth values and the map $true : \{\star\} \rightarrow \top$ is the subobject classifier.
- The subobject classifier can be much more interesting in different categories! Let's see an example of one.

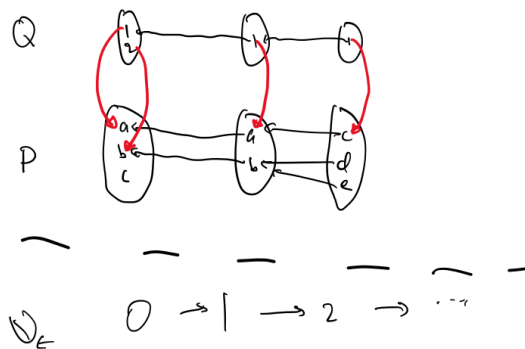
Definition 62: Topos of trees

Consider the following category called \mathbb{N}_{\rightarrow} :

$$0 \longrightarrow 1 \longrightarrow 2 \longrightarrow 3 \longrightarrow \dots$$


The *topos of trees* is the category of presheaves on \mathbb{N}_{\leftarrow} written $\text{Psh}(\mathbb{N}_{\rightarrow})$.

- We haven't defined what a topos is yet, so just treat the above as the definition for a special kind of presheaf category. But, let's see what a subobject classifier is in this category. First, let's look at what monomorphisms look like. Here is an example of a monomorphism:



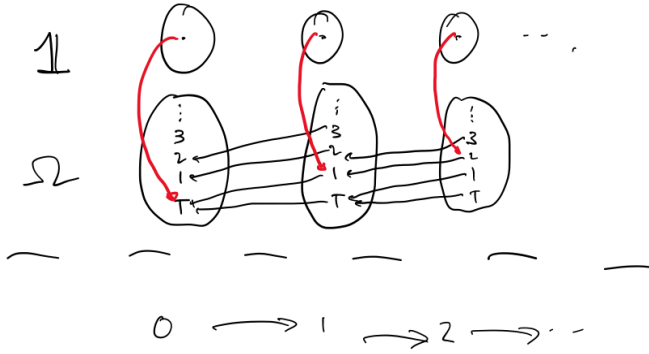
The red arrow is a natural transformation $Q \Rightarrow P$ where each component of the natural transformation is an injective function. In a situation analogous to the one in finite sets, a natural transformation $\alpha : P \Rightarrow Q$ is a monomorphism in a presheaf category if and only if each component is an injective map.

- Objects of this category can be thought of as “sets of facts that are true at particular times”. They will correspond to *Kripke models*, and a monomorphism will correspond to a Kripke sub-world.
- Now we can think about the object of truth values. It is this:



Formally, Ω is a presheaf that maps each object to the set $\{\top, 1, 2, \dots\}$, and each morphism to the predecessor function that maps 1 to \top and \top to \top . Intuitively, think of each element of each set as representing “the number of steps until something is true”. With this intuition, a non-zero value represents that something is false at a particular time step.

- Given this object of truth values, we can design the subobject classifier as taking each time step and mapping it to the number of steps of “gas” that timestep has left:



- Let's use this object of truth values to find the classifying morphism for α above. We need to find a morphism $P \xrightarrow{\chi_\alpha} \Omega$
- The mere existence of subobject classifiers already forces your category to behave in set-like ways. An example: Every category with a subobject classifier is *balanced*, meaning that if there is an epimorphism and monomorphism between two objects, they are isomorphic. This is just like in sets where the existence of an injection and surjection implies an isomorphism.

14.3 Generalized elements

Important picture: the subobject classifier gives us an interpretation for set-builder.

Think of $\varphi(x)$ as an open term with free variable of type X .

$$\begin{array}{ccc} \{x \in X \mid \varphi(x)\} & \longrightarrow & 1 \\ \downarrow & & \downarrow \text{true} \\ X & \xrightarrow{\varphi(x)} & \Omega \end{array}$$

We can say *an element of X satisfies φ* if we can find a morphism x_0 such that the following commutes:

$$\begin{array}{ccccc} & & \{x \in X \mid \varphi(x)\} & \longrightarrow & 1 \\ & \nearrow x_0 & \downarrow & & \downarrow \text{true} \\ 1 & \longrightarrow & X & \xrightarrow{\varphi(x)} & \Omega \end{array}$$

This is often too strong, e.g. Kripke semantics. Enter generalized elements:

We say $U \models \varphi$ (read " U forces φ ").

14.4 Kripke-Joyal Semantics

- The Kripke-Joyal semantics gives us a full picture of working internal to a topos

- These will look very familiar if you are familiar with separation logic papers

14.5 *Elementary topos*

- The Kripke-Joyal semantics tells us what features we need our category to support, which will be an elementary topos:

Definition 63: Elementary topos

A category \mathcal{C} is an elementary topos if and only if it has products, coproducts, exponentials, an initial object, a terminal object, an object of truth values, and satisfies that for every object X the slice category \mathcal{C}/X has products.

- Every presheaf category is an elementary topos (!!)

Sheaf semantics

Last time we mentioned that sheaves serve as a “hash-lang for math”. More precisely, this means that sheaves provide a model of *higher-order intuitionistic logic*. In this logic it is possible to express a large amount of math; as a rough approximation, think of what can be written down as using Prop in Rocq (and without Inductives), or what can be expressed in Isabelle/HOL. While the standard model of this logic interprets propositions as classical truth values (either true or false), each category of sheaves provides a non-standard interpretation which often turns out to be convenient for modeling situations that arise in PL. This chapter provides a pared down presentation of the key idea behind how sheaves accomplish this. Rather than covering full higher-order intuitionistic logic, we will cover the intuitionistic propositional logic:

$$\varphi, \psi ::= \top \mid \perp \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \Rightarrow \psi \mid a$$

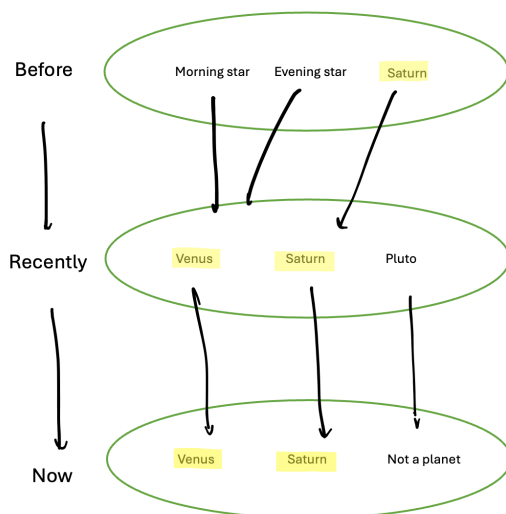
This logic consists of the usual intuitionistic logical connectives, plus some set of atomic propositions a .

Eventually we will see how to interpret this logic using the so-called *sheaf semantics*. We will build up to it by first talking about Kripke semantics, then presheaf semantics.

15.1 Kripke semantics

Kripke semantics are concerned with describing beliefs over time. For example, in the past, we believed that in the night sky there were 3 objects: a morning star and an evening star. Then, recently, we learned that in fact the morning star and evening star were both venus, and discovered new planets like Saturn and Pluto.

This chapter is heavily based on these notes from Alex Simpson: <https://alexksimpson.github.io/Talks/TutorialOnSheafSemantics.pdf>



We can describe this situation by interpreting the connectives of intuitionistic propositional logic with a generalized set of truth values. The key idea of Kripke semantics is to introduce a set W of *possible worlds*, whose elements represent states of knowledge. These states of knowledge are partially ordered according to what an *accessibility relation* \leq ; the relation $w \leq w'$ (pronounced “world w' is accessible from w ”) represents the situation where a state of knowledge w can evolve into w' . For the planetary example above, one could set W to the three-element set $\{\text{before, recently, now}\}$ ordered by $\text{before} \leq \text{recently} \leq \text{now}$.

Each proposition φ is interpreted as a subset of W giving the set of possible worlds in which φ holds. Crucial to this interpretation is a *monotonicity* property that models the fact that knowledge is stable over time: if φ holds in world w , then it also holds in all worlds accessible from w . The Kripke semantics of intuitionistic propositional logic is as follows.

$$\begin{aligned} \llbracket - \rrbracket &: \text{Formula} \rightarrow \wp(W) \\ \llbracket \top \rrbracket &= W \\ \llbracket \perp \rrbracket &= \emptyset \\ \llbracket \varphi \wedge \psi \rrbracket &= \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket \\ \llbracket \varphi \vee \psi \rrbracket &= \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket \\ \llbracket \varphi \Rightarrow \psi \rrbracket &= \{w \mid \text{for all } w' \geq w, \text{ if } w' \in \llbracket \varphi \rrbracket \text{ then } w' \in \llbracket \psi \rrbracket\} \end{aligned}$$

Of these definitions, the interpretation of \Rightarrow is the most interesting. A naive interpretation that says $\varphi \Rightarrow \psi$ holds in world w if and only if $w \in \llbracket \varphi \rrbracket$ implies $w \in \llbracket \psi \rrbracket$ is not valid here, because it fails to satisfy monotonicity.

These semantics are designed so that you have sound interpre-

tations of the proof rules of intuitionistic propositional logic: if $\varphi_1, \dots, \varphi_n \vdash \psi$ then $[[\varphi_1]] \cap \dots \cap [[\varphi_n]] \subseteq [[\psi]]$, where \vdash is intuitionistic syntactic entailment. In this sense the Kripke semantics provides a “domain-specific reinterpretation” of the usual rules of propositional logic, where each formula φ denotes a monotone subset $[[\varphi]]$ of the poset W .

As an example of this kind of reinterpretation, consider the poset \mathbb{N} ordered by *reverse* inclusion, so that a natural number n is accessible from m if n is *smaller* than m . This gives a semantics of intuitionistic propositional logic that is used in *step-indexing*¹.

¹ Derek Dreyer, Amal Ahmed, and Lars Birkedal. Logical step-indexed logical relations. In *2009 24th Annual IEEE Symposium on Logic In Computer Science*, pages 71–80. IEEE, 2009

15.2 Presheaf semantics

Presheaf semantics generalizes the poset (W, \leq) from Kripke semantics to a category \mathcal{C} . Elements $w \in W$ are generalized to objects c of \mathcal{C} and the accessibility relation $w' \geq w$ is generalized to a morphism $f : c' \rightarrow c$. There are two things to note about this generalization. First, there can now be multiple such morphisms f between any two given objects; intuitively, this is like a generalization of traditional Kripke semantics where there can be multiple different “ways” of going from one world to another. Second, there is a bit of reversal of direction going on: the accessibility relation $w' \geq w$ becomes the existence of a morphism $c' \rightarrow c$. This reversal of direction has to do with the fact that presheaves on a category \mathcal{C} are contravariant; i.e., functors $\mathcal{C}^{\text{op}} \rightarrow \text{Set}$.

Rather than interpreting each formula as a monotone subset of W , the presheaf semantics interprets each formula as what is called a *sieve* on \mathcal{C} .

Definition 64

A *sieve* on a category \mathcal{C} is a collection of morphisms S in \mathcal{C} that is closed under precomposition in the sense that if $c' \xrightarrow{f} c$ is in S then so is $c'' \xrightarrow{g} c' \xrightarrow{f} c$ for any morphism $c'' \xrightarrow{g} c'$.

Intuitively, a sieve S containing a morphism $c' \xrightarrow{f} c$ represents a proposition φ where, if one is currently in world c , then φ will become true if one transitions to world c' via f . The requirement that S be closed under precomposition provides a generalization of monotonicity.

The *presheaf semantics* is the corresponding generalization of

Kripke semantics where propositions are interpreted as sieves:

$$\begin{aligned} \llbracket - \rrbracket &: \text{Formula} \rightarrow \{\text{sieves on } \mathcal{C}\} \\ \llbracket \top \rrbracket &= \text{the set of all morphisms in } \mathcal{C} \\ \llbracket \perp \rrbracket &= \emptyset \\ \llbracket \varphi \wedge \psi \rrbracket &= \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket \\ \llbracket \varphi \vee \psi \rrbracket &= \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket \\ \llbracket \varphi \Rightarrow \psi \rrbracket &= \{c' \xrightarrow{f} c \mid \text{for all } c'' \xrightarrow{g} c', \text{ if } f \circ g \in \llbracket \varphi \rrbracket \text{ then } f \circ g \in \llbracket \psi \rrbracket\} \end{aligned}$$

Just like in the Kripke semantics, the presheaf semantics for Kripke logic will also validate intuitionistic logic entailments.

The following proposition verifies that the presheaf semantics indeed generalizes the posetal case.

Proposition 15.2.1. For any poset (W, \leq) , monotone subsets of W are the same as sieves on W^{op} , where W is considered as a thin category, and the presheaf semantics agrees with the Kripke semantics.

In particular, the Kripke semantics for step-indexed logic corresponds to the presheaf semantics for the usual poset of natural numbers (\mathbb{N}, \leq) made into a thin category. (Note that, unlike in the last section, this poset has the natural numbers ordered in the usual way. The appearance of op in the presheaf semantics automatically performs the reversal of the ordering relation that we did manually in the last section.)

A more interesting example is the presheaf semantics for name generation. We will take as our indexing category the category of finite injective maps $\text{FinInj}^{\text{op}}$. Objects are finite sets Γ and homs are injection functions $\Gamma \rightarrow \Gamma'$. Intuitively, this represents how names can change over time: an injective function $\Gamma \rightarrow \Gamma'$ expands the set of available names from Γ to Γ' , while possibly renaming the variables in Γ along the way.

In the resulting presheaf semantics, the interpretation $\llbracket \varphi \rrbracket$ of a proposition φ is the set of renamings $r : \Gamma \rightarrow \Gamma'$ such that φ holds when Γ' is “available”, i.e., Γ is rewritten into some new context Γ' that validates φ .

Notice that this is *proof-relevant* in the sense that the kind of substitution performed matters: there can be many injective maps $\Gamma \rightarrow \Gamma'$. As usual the interesting case here is implication, which can be read off directly from the Kripke semantics above.

15.3 Sheaves

At this point we’ve hit most of the common PL applications. Presheaf semantics get you very far: it is not that common to need to use

sheaves.

Whereas presheaf and Kripke semantics are well-suited to modeling *when* a proposition becomes true, sheaves are well-suited to modeling *where* a proposition becomes true. Sheaves interpret objects c of the indexing category \mathcal{C} not as states of knowledge, but rather as *spaces*; morphisms $f : c' \rightarrow c$ represent *subspace inclusions* that map a smaller space c' into a larger space c . The sheaf semantics then models the spaces where a given proposition is true. The monotonicity requirements of Kripke and presheaf semantics continue to be relevant under this spatial reading: if a proposition holds in a space c , then one would intuitively still expect it to hold in a subspace c' picked out by some inclusion $f : c' \rightarrow c$. But the spatial setting provided by sheaf semantics adds an additional twist not present in the more temporal reading provided by presheaf and Kripke semantics: it is possible for a family of “subspaces” $\{f_i : c'_i \rightarrow c\}_{i \in I}$ to jointly *cover* a space c .

The spatial interpretation offered by sheaf semantics affects the reading of disjunction. The Kripke/presheaf reading of disjunction says that $\varphi \vee \psi$ holds on a space w if either φ holds on all of w or ψ holds on all of w . But there is a third possibility: perhaps φ holds on some piece w_1 of w , and ψ holds on some piece w_2 of w , and together w_1, w_2 cover w .

A simple example is given by interpreting “space” as “subset of \mathbb{R} ” and “inclusion” as subset inclusion. These spaces along with their inclusion relations forms a thin category $\wp(\mathbb{R})$. Families of morphisms with a common codomain U in this setting are families $(U_i)_{i \in I}$ of subsets of \mathbb{R} where $U_i \subseteq U$ for all $i \in I$. It is then natural to say that such a family *covers* U if they union to U : $\bigcup_i U_i = U$. The sheaf semantics for this notion of cover interprets each proposition φ as a collection $\llbracket \varphi \rrbracket$ of subsets of \mathbb{R} ; a subset U is in $\llbracket \varphi \rrbracket$ if “ φ is true everywhere in U ”. For example, one can define an atomic proposition φ = “is less than one” whose interpretation $\llbracket \varphi \rrbracket$ is the collection of all those subsets U of \mathbb{R} contained in the $(-\infty, 1)$, and an atomic proposition ψ = “is greater than zero” whose interpretation $\llbracket \psi \rrbracket$ is the collection of all those subsets contained in $(0, \infty)$. Now consider the disjunction $\varphi \vee \psi$. Intuitively, this proposition should hold *everywhere* in \mathbb{R} , as every real number is either less than one or greater than zero: we should expect $\llbracket \varphi \vee \psi \rrbracket = \llbracket \top \rrbracket =$ the set of all subsets of \mathbb{R} . But the presheaf semantics says that $\llbracket \varphi \vee \psi \rrbracket = \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket$, which consists only of those subsets of \mathbb{R} are either contained in $(-\infty, 1)$ or $(0, \infty)$. The sheaf semantics corrects this by taking into account the fact that $(0, \infty)$ and $(-\infty, 1)$ form a *cover* of \mathbb{R} : a disjunction $\varphi \vee \psi$ holds on a subset U of \mathbb{R} under the sheaf semantics if there is a covering $U = U_1 \cup U_2$ of U by two subspaces U_1, U_2 such that φ holds on

U_1 and ψ holds on U_2 .

More generally, the sheaf semantics is parameterized by a relation T called a *coverage* that says when a given family of morphisms $\{f_i : c_i \rightarrow c\}_{i \in I}$ with common codomain c counts as a cover of c .

Definition 65

A *coverage* on a category \mathcal{C} is a relation T between families of morphisms with common codomain c and objects c of \mathcal{C} . If T relates a family $\{f_i : c_i \rightarrow c\}_{i \in I}$ to an object c we say that $\{f_i : c_i \rightarrow c\}_{i \in I}$ is a *T-cover* of c . The relation T must satisfy the following properties:

- Reflexivity: the singleton set $\{\text{id}_c : c \rightarrow c\}$ is in T for every object c
- Transitivity: if $\{f_i : c_i \rightarrow c\}_{i \in I}$ is in T and for every $i \in I$ there is a family $\{f_{ij} : c_{ij} \rightarrow c_i\}_{j \in J_i}$ in T , then the family $\{f_i f_{ij} : c_{ij} \rightarrow c\}_{i \in I, j \in J_i}$ is in T
- Stability: if $\{f_i : c_i \rightarrow c\}_{i \in I}$ is in T and $g : c' \rightarrow c$ is an arbitrary morphism into c , then there exists a family $\{f'_j : c'_j \rightarrow c'\}_{j \in J}$ such that every gf'_j factors through some f_i ; that is, for every $j \in J$ there exists an $i \in I$ such that $gf_j = f_i g'$ for some g' .

These coverage conditions are best visualized as closure conditions on families of morphisms drawn as trees. Reflexivity says that the trivial tree with just an identity morphism counts as a cover; transitivity says that a family of covering trees can be grafted onto another to form a new covering tree; stability says that a covering tree can be “pulled back” along a morphism $g : c' \rightarrow c$.

With the definition of coverage in hand, the sheaf semantics of intuitionistic propositional logic is as follows.

$$\llbracket - \rrbracket : \text{Formula} \rightarrow \{\text{sieves on } \mathcal{C}\}$$

$$\llbracket \top \rrbracket = \text{the set of all morphisms in } \mathcal{C}$$

$$\llbracket \perp \rrbracket = \{f : c' \rightarrow c \mid \emptyset \text{ is a } T\text{-cover of } c\}$$

$$\llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket$$

$$\llbracket \varphi \vee \psi \rrbracket = \{f : c' \rightarrow c \mid \text{there is a } T\text{-cover } \{f_i : c_i \rightarrow c'\}_{i \in I} \text{ of } c' \text{ such that } ff_i \in \llbracket \varphi \rrbracket \text{ or } ff_i \in \llbracket \psi \rrbracket \text{ for all } i\}$$

$$\llbracket \varphi \Rightarrow \psi \rrbracket = \{c' \xrightarrow{f} c \mid \text{for all } c'' \xrightarrow{g} c', \text{ if } f \circ g \in \llbracket \varphi \rrbracket \text{ then } f \circ g \in \llbracket \psi \rrbracket\}$$

Like the presheaf semantics, this semantics interprets each proposition as sieve. But the sheaf semantics satisfies an additional invariant. Each sieve produced by the sheaf semantics is what is known as *T-closed*; we omit the formal definition here, but intuitively this says

that if a proposition holds for every component of a covering family, then it holds for the space covered.

Examples of coverages:

- The notion of cover for subsets of \mathbb{R} forms a coverage on $(\wp(\mathbb{R}), \subseteq)$ considered as a thin category.
- One can place a coverage on the category *Formula* of propositional formulae that says that a finite set $\varphi_1, \dots, \varphi_n$ covers a formula φ if $\varphi_1 \vee \dots \vee \varphi_n \dashv\vdash \varphi$. Intuitively, this notion of coverage corresponds to the notion of cover one would get by interpreting formulas as sets of models.
- A more exotic notion of coverage, for the category $\text{FinInj}^{\text{op}}$, is given by declaring any singleton set to be a cover (that is, $\{r : \Gamma \rightarrow \Gamma'\}$ covers Γ for any Γ, Γ', r). This coverage corresponds to a category of sheaves known as the *Schanuel topos*, and produces a sheaf semantics that can be used to talk about name generation as modeled by *nominal sets*².

² Andrew M Pitts. *Nominal sets: Names and symmetry in computer science*. Cambridge University Press, 2013

15.3.1 Where are the categories?

One might have noticed that categories of presheaves or sheaves have not played a role in any of the above. This is because we have restricted the discussion to propositional logic. Moving from propositional to first-order logic requires interpreting formulas-in-context $\Gamma \vdash \varphi$, which in turn requires interpreting contexts Γ and upgrading the semantics of formulas to include substitutions $\gamma \in \llbracket \Gamma \rrbracket$. It will turn out that this upgrade requires interpreting each $\llbracket \Gamma \rrbracket$ as a presheaf (in the presheaf semantics) or sheaf (in the sheaf semantics).

We briefly sketch how this interpretation works. Recall the planetary picture from before. Turning this picture on its side, we obtain a balloon diagram lying over the category W^{op} , where W is the poset before \leq recently \leq now. This balloon diagram defines a presheaf *Sky* on W^{op} , whose elements model our conception of objects in the night sky over time. This presheaf can serve as the interpretation of a basic sort *Sky* of a first-order intuitionistic logic. Logical predicates $x : \text{Sky} \vdash \varphi(x)$ then denotes subpresheaves: for instance, the points in the diagram highlighted in yellow depict a subpresheaf corresponding to the predicate φ that out those objects that are considered planets through to the present day.

16

Monads

Monads are widely used for structuring effectful computation in functional programs. For example, we can have a small functional programming language with a monadic type-former \mathbb{T} :

$$\begin{aligned} A, B ::= & \text{Unit} \mid A \times B \mid A \rightarrow B \mid \mathbb{T} A \\ M, N ::= & \langle \rangle \mid \langle M, N \rangle \mid \text{fst } M \mid \text{snd } M \mid \lambda x. M \mid MN \mid x \\ & \mid \text{pure } M \mid x \leftarrow M; N \end{aligned} \quad (16.1)$$

The syntax `pure M` lifts a pure computation into a monadic one, and the syntax `x ← M; N` is a *monadic bind* that enables computation inside the monad. The way that these syntactic forms are used is best understood via their typing judgments:

$$\frac{\Gamma \vdash M : A}{\Gamma \vdash \text{pure } M : \mathbb{T} A} \quad \frac{\Gamma \vdash M : \mathbb{T} A \quad \Gamma, x : A \vdash N : \mathbb{T} B}{\Gamma \vdash x \leftarrow M; N : \mathbb{T} B}$$

If you’ve programmed in Haskell, you’re quite familiar with examples of monads. A simple example is the `Maybe` monad, which lets you program with values that can be either `Nothing` or `Just` a value. This makes it a convenient pattern for implementing and sequencing partial functions, as the following Haskell code illustrates:

```
1 divide :: Double -> Double -> Maybe Double
2 divide a 0 = Nothing
3 divide a b = Just $ a / b
4
5 my_computation :: Double -> Maybe Double
6 my_computation x = do
7   div1 <- divide 10 x;
8   div2 <- divide 20 x;
9   divide div1 div2
```

Let’s consider now the denotational semantics of monads. Let’s start with the type former $\mathbb{T} M$. We typically think of the interpretation of types as “a space of values that behave like a type”, where

A monad is just a monoid in the category of endoburritos.

we use universal properties to characterize what it means to behave like a type. A monad is a new and exciting creature when viewed from this lens: it *takes in a type* A (so it is higher-order), and it produces a *new type* $\llbracket T A \rrbracket$ that has A living inside of it. This feels very functorial. We can couple this intuition with the requirement that denotational semantics be defined compositionally, and we can arrive at the conclusion that $\llbracket T \rrbracket$ is a functor:

$$\llbracket T A \rrbracket = \llbracket T \rrbracket (\llbracket A \rrbracket) \quad (16.2)$$

Concretely, suppose we are interpreting our small higher-order language in some Cartesian-closed category \mathcal{C} . Then, $\llbracket T \rrbracket$ is an *endo-functor* $\llbracket T \rrbracket : \mathcal{C} \rightarrow \mathcal{C}$, a particular kind of functor from a category into itself. One way of understanding what T is is that it maps elements to *generalized elements* that are imbued with the additional structure of the monad.

As usual we don't know what kind of properties the functor $\llbracket T A \rrbracket$ must have yet; to see those, we'll need to reverse-engineer them from how monads are used. Let's start by examining `pure`. By our usual translation of typing judgments into morphisms, we can conclude that $\llbracket \Gamma \vdash \text{pure } M : T A \rrbracket$ must have type $\llbracket \Gamma \rrbracket \rightarrow \llbracket T \rrbracket (\llbracket A \rrbracket)$. We need to build this morphism. We know how to get close to it: $\llbracket M : A \rrbracket$ has type $\llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$. But, we need an extra morphism to get from A to $T A$: this is called the *unit* of the monad at $\llbracket A \rrbracket$, and will be denoted $\eta_{\llbracket A \rrbracket}$:

$$\llbracket \Gamma \vdash \text{pure } M : T A \rrbracket = \llbracket \Gamma \rrbracket \xrightarrow{\llbracket M \rrbracket} \llbracket A \rrbracket \xrightarrow{\eta_{\llbracket A \rrbracket}} \llbracket T \rrbracket (\llbracket A \rrbracket)$$

In Haskell, η is called `pure` or `return`. Notably, in Haskell, `pure` is a polymorphic function of type `a -> m a` for some monad `m` and type parameter `a`. This is a polymorphic function, so it makes sense to think of η as a *natural transformation* $\eta : \text{id}_{\mathcal{C}} \Rightarrow T$.

```
1 class Functor m => Monad m where
2   return :: alpha -> m alpha
3   join   :: m (m alpha) -> m alpha
```

Interpreting `bind`:

$$\begin{aligned} \left[\frac{\Gamma \vdash M : T A \quad \Gamma, x : A \vdash N : T B}{\Gamma \vdash x \leftarrow M; N : T B} \right] &: \llbracket \Gamma \rrbracket \rightarrow \llbracket T \rrbracket (\llbracket B \rrbracket) \\ &= \llbracket \Gamma \rrbracket \xrightarrow{\langle \text{id}, \llbracket M \rrbracket \rangle} \llbracket \Gamma \rrbracket \times \llbracket T \rrbracket (\llbracket A \rrbracket) \\ &\quad \xrightarrow{\text{strength}} \llbracket T \rrbracket (\llbracket \Gamma \rrbracket \times \llbracket A \rrbracket) \\ &\quad \xrightarrow{\llbracket T \rrbracket (\llbracket N \rrbracket)} \llbracket T \rrbracket (\llbracket T \rrbracket (\llbracket B \rrbracket)) \\ &\quad \xrightarrow{\mu} \llbracket T \rrbracket (\llbracket B \rrbracket) \end{aligned}$$

[1] has a very nice discussion of this motivation of monads as “generalized elements of a category”.

16.1 What's the problem?

"A monad is just a monoid in the category of endofunctors, what's the problem?"

We've seen how a monad consists of (1) an endofunctor $T : \mathcal{C} \rightarrow \mathcal{C}$; (2) a *unit* natural transformation $\eta : \text{id}_{\mathcal{C}} \Rightarrow T$ (3) a *multiplication (or join)* natural transformation $\mu : T^2 \Rightarrow T$. As usual, these special morphisms must satisfy certain laws in order to behave coherently. In particular, these two natural transformations need to satisfy natural *monoid laws*. Recall the definition of a monoid:

Definition 66: Monoid

A *monoid* is a triple (X, \bullet, e) where X is a set, \bullet is a *multiplication operator* $\bullet : X \times X \rightarrow X$, and $e \in X$ is a distinguished *unit element* satisfying: (1) *associativity*: for any $a, b, c \in X$ it is the case that $a \bullet (b \bullet c) = (a \bullet b) \bullet c$ and (2) *unitality*: for any $a \in X$ it is the case that $a \bullet e = e \bullet a = a$.

A monad will satisfy categorical analogs of the monoid laws: the unit natural transformation η must behave like a monoidal unit, and the multiplication natural transformation μ must behave like a monoidal product. Let's examine associativity first. This requirement can be captured by the following commutative square:

$$\begin{array}{ccc} T^3 & \xrightarrow{T\mu} & T^2 \\ \downarrow \mu T & & \downarrow \mu \\ T^2 & \xrightarrow{\mu} & T \end{array} \quad (16.3)$$

This square has some unfamiliar notation in it: the natural transformation $T\mu$ is given by composing the natural transformation μ with the functor T . Concretely, the natural transformation $T\mu$ has components $(T\mu)_A : T^3(A) \rightarrow T^2(A)$ given by composing T with each component of μ , i.e. $(T\mu)_A = T \circ \mu_A$. The definition for μT is analogous.

The next requirement on the monad natural transformations is that η behaves like a monoidal unit. This is captured by the following commuting diagram:

$$\begin{array}{ccc} T & \xrightarrow{\eta T} & T^2 & \xleftarrow{T\eta} & T \\ & \searrow \text{id}_T & \downarrow \mu & \swarrow \text{id}_T & \\ & & T & & \end{array} \quad (16.4)$$

There are two faces of this diagram each capturing one of the unitality requirements. The left square asserts that $\mu \circ \eta T = \text{id}_T$, and the

right square asserts that $\mu \circ T\eta = \text{id}_T$; these equations correspond naturally to the monoid unitality law. Together, these equations define a monad:

Definition 67: Monad

Let \mathcal{C} be a category. A *monad* on \mathcal{C} is a triple (T, η, μ) where $T : \mathcal{C} \rightarrow \mathcal{C}$ is an endofunctor, $\eta : \text{id}_{\mathcal{C}} \Rightarrow T$ is the *unit* natural transformation, and $\mu : T^2 \Rightarrow T$ is a *multiplication (or join)* natural transformation, satisfying the associativity law in Eq. (16.3) and unit law in Eq. (16.4).

16.2 Examples of monads

16.2.1 Maybe monad

The *maybe monad* is often the first example of a monad one encounters in Haskell. Intuitively, the type `Maybe A` represents values that behave like A extended with the option of being “nothing”.

- Let $F : \text{FinSet} \rightarrow \text{FinSet}$ be a functor:
 - Action on objects: $A \mapsto A \uplus \{\star\}$
 - Action on morphisms $f : A \rightarrow B$:

$$F(f) = a \mapsto \begin{cases} \text{inl } f(x) & \text{if } a = \text{inl } x \\ \text{inr } \star & \text{if } a = \text{inr } \star \end{cases}$$

- The unit $\eta : \text{id}_{\text{FinSet}} \Rightarrow F$ has components that straightforwardly inject A into $F(A)$:

$$\begin{aligned} \eta_A : \text{id}_{\text{FinSet}}(A) &\rightarrow F(A) \\ &= x \mapsto \text{inl } x \end{aligned}$$

- The multiplication $\mu : F^2 \rightarrow F$ merges the two nested distinguished elements \star into a single one and flattens the nesting.

16.2.2 Powerset monad

The *powerset monad* is defined by

- The *powerset functor* $\wp : \text{FinSet} \rightarrow \text{FinSet}$ that sends each finite set A to the set of all subsets $\{A_i \mid A_i \subseteq A\}$ is a monad.
- The unit $\eta : \text{id}_{\text{FinSet}} \Rightarrow \wp$ has components that map sets to the singleton set containing them:

$$\begin{aligned} \eta_A : A &\rightarrow \wp(A) \\ &= \{A\} \end{aligned}$$

- The multiplication $\mu : \wp^2 \Rightarrow \wp$ has components that “unfold the double powerset”:

$$\begin{aligned} \mu_A : \wp(\wp(A)) &\Rightarrow \wp(A) \\ &= \mathcal{S} \mapsto \bigcup_{A \in \mathcal{S}} A \end{aligned}$$

16.3 Monads from adjunctions

A surprising fact: all adjunctions form a monad, and all monads come from an adjunction pair!

Let’s see first how every adjunction forms a monad:

Proposition 16.3.1. Consider functors $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{D} \rightarrow \mathcal{C}$ where $F \dashv G$. Then, there is a unit to the adjunction $\eta : \text{id}_{\mathcal{C}} \Rightarrow G \circ F$ and co-unit $\nu : F \circ G \Rightarrow \text{id}_{\mathcal{D}}$. Then, there is a monad (T, η, μ) :

- Endofunctor T given by the roundtrip $T = G \circ F$
- Unit given by the unit of the adjunction
- Multiplication $\mu : G \circ F \circ G \circ F \Rightarrow G \circ F$ given by $\mu = G \circ \nu \circ F$.

16.3.1 Maybe monad from a free/forgetful adjunction

- Consider the categories FinSet and FinSet^* .
- There is a functor $\text{forget} : \text{FinSet}^* \rightarrow \text{FinSet}$ called the *forgetful functor* that “forgets about the distinguished element”.
- There is a functor $\text{free} : \text{FinSet} \rightarrow \text{FinSet}^*$ that adds on a distinguished element
- There is an adjunction $\text{free} \dashv \text{forgetful}$.
- The round trip $\text{free} \circ \text{forgetful}$ is exactly the monad F we built earlier for the maybe monad.

16.3.2 List monad

- List monad: Free/forgetful adjunction between sets and monoids. Unit is singleton list construction; join is concatenation.

Example: the free \dashv forgetful adjunction between pointed sets and ordinary sets is the Maybe monad.

Modeling effects with syntax

Monads have found extensive application in computer science and language ideas for dealing with effects. The previous chapter showed how to interpret a monadic type, but it was not exactly clear how this captures the notion of an *effectful* computation. Here our aim is to make this clear by elucidating the connection between a monad and its *Klesli category*, which can be thought of as a syntactic representation of the effects being modeled. For any particular monad T on a category \mathcal{C} , the monad will come about as an adjunction between \mathcal{C} and its Klesli category.

We can start by characterizing effects using tiny languages that capture only the effectful fragments of computation. For example, we can represent the computations taking place under the finite probability monad `Dist` using the following tiny grammar:

$$e_{\text{Dist}} ::= \text{pure}(v) \mid \text{if flip}(p) e_1 e_2$$

The syntax is extremely pared down: it does not have any features at all except for the ability to refer to values v , or branch on a random Boolean that is true with probability p . The set of all terms that can be formed out of this grammar is parameterized by the collection of values v : for each set of values \mathcal{V} we denote the set of all terms from `eDist` using values \mathcal{V} as $\text{Exp}_{\text{Dist}}(\mathcal{V})$.

As usual, given a language it is useful to characterize its equivalences. This is a design question: you, as the designer of the effect, must decide which effectful programs ought to behave the same.

Some equations are quite natural from this perspective:

$$\frac{}{\text{if flip}(p) e e \equiv e} \text{ (DROP)}$$

$$\frac{}{\text{if flip}(p) e_1 e_2 \equiv \text{if flip}(1-p) e_2 e_1} \text{ (SWAP)}$$

$$\frac{}{\text{if flip}(1) e_1 e_2 \equiv e_1} \text{ (THEN)}$$

$$\frac{}{\text{if flip}(0) e_1 e_2 \equiv e_2} \text{ (ELSE)}$$

The above equations capture natural reasoning principles. But, is this *all* the necessary equations that characterize? Here is an example of two tiny effectful programs that ought to be equivalent but aren't yet:

$$\text{if flip}(1/2) (\text{if flip}(1/3) (\text{pure } \top) (\text{pure } \perp)) (\text{pure } \perp) \stackrel{?}{\equiv} \text{if flip}(1/6) (\text{pure } \top) (\text{pure } \perp)$$

This motivates one more equation that captures all the equivalences for this effect:

$$\frac{p_1 + p_2 + p_3 = 1 \quad p_{12} = p_1 + p_2 \quad p_{23} = p_2 + p_3}{\text{if flip}(p_1 + p_2) (\text{if flip}(p_1/p_{12}) e_1 e_2) e_3 \equiv \text{if flip}(p_1) e_1 (\text{if flip}(p_2/p_{23}) e_2 e_3)} \text{ (ASSOC)}$$

Now we have a tiny language that captures *only* the effect of allocating and branching on random quantities. Using this tiny language, we can bolt on sampling as a monad in the category of sets Set . Recall that the set $\text{Exp}_{\text{Dist}}(\mathcal{V})$ of probabilistic computations with values in \mathcal{V} is parameterized by a set \mathcal{V} . This can be used to define an endofunctor $D : \text{Set} \rightarrow \text{Set}$:

$$\begin{aligned} D &: \text{Set} \rightarrow \text{Set} \\ D(\mathcal{V}) &= \text{Exp}_{\text{Dist}}(\mathcal{V}) / \equiv \\ D(f : \mathcal{V} \rightarrow \mathcal{V}') &= e \mapsto e[f(v)/v] \end{aligned}$$

The action of this functor on morphisms sends a function f to the function that uses f to substitute each value v in a given computation e with $f(v)$.

This endofunctor D forms a monad. The unit η is defined by the coercion from values to computations:

$$\begin{aligned} \eta_{\mathcal{V}} &: \mathcal{V} \rightarrow D(\mathcal{V}) \\ \eta_{\mathcal{V}}(v) &= \text{pure}(v) \end{aligned}$$

The multiplication is defined by *substitution*: given a computation e that produces computations, the multiplication $\mu(e)$ is “flattening” of

the syntax tree of e :

$$\begin{aligned}\mu_{\mathcal{V}} &: D(D(\mathcal{V})) \rightarrow D(\mathcal{V}) \\ \mu_{\mathcal{V}}(\text{pure}(v)) &= v \\ \mu_{\mathcal{V}}(\text{if flip}(p) e_1 e_2) &= \text{if flip}(p) \mu_{\mathcal{V}}(e_1) \mu_{\mathcal{V}}(e_2)\end{aligned}$$

Note in the first case that the “value” v is an element of $D(\mathcal{V})$. Thinking of computations as abstract syntax trees, this multiplication operation flattens a tree of trees into an ordinary tree. One can check that the monad laws are satisfied; they correspond to properties one might expect to hold of tree flattening.

The finite probability distribution monad $\mathcal{D} : \text{Set} \rightarrow \text{Set}$ sends sets X to the set of finitely-supported functions $\mu : X \rightarrow [0, 1]$ such that $\sum_{x \in X} \mu(x) = 1$.

As more convenient notation, we can describe a probability distribution on X as a *formal series*: a distribution $\mu : X \rightarrow [0, 1]$ can be written as an expression $c_1[x_1] + \cdots + c_n[x_n]$ where $c_i \in [0, 1]$ are real-valued numbers satisfying $\sum_i c_i = 1$, and each $[x_i]$ is a *formal parameter* (i.e., a variable). There is a natural equivalence relation on formal series given by combining the coefficients of formal parameters:

$$c_1[x_1] + \cdots + c_i[x_i] + c'_i[x_i] + \cdots + c_n[x_n] \equiv c_1[x_1] + \cdots + (c_i + c'_i)[x_i] + \cdots + c_n[x_n]$$

The additional equations of formal series are the natural equivalences given by the commutativity and associativity of addition and multiplication.

Now we can give the unit and multiplication natural transformations for \mathcal{D} :

- The unit $\eta : \text{id}_{\text{Set}} \Rightarrow \mathcal{D}$ has components $\eta_X : X \rightarrow \mathcal{D}(X)$ that place all the probability mass on a single point: $\eta_X(x) = 1[x]$.
- Multiplication $\mu : \mathcal{D}^2 \Rightarrow \mathcal{D}$ merges the probability distributions in the following way:

$$\begin{aligned}\mu_X &: \mathcal{D}\mathcal{D}(X) \rightarrow \mathcal{D}(X) \\ &= (c_1[c_{11}[x_1] + \cdots + c_{n1}x[n]] + \cdots + c_n[c_{1n}[x_1] + \cdots + c_{nn}x[n]]) \\ &\mapsto (c_1(c_{11}[x_1] + \cdots + c_{n1}x[n]) + \cdots + c_n(c_{1n}[x_1] + \cdots + c_{nn}x[n]))\end{aligned}$$

Proposition 17.0.1. The monads D and \mathcal{D} are isomorphic; that is, there is a natural isomorphism $D \cong \mathcal{D}$ that commutes with the unit and multiplication operations of D and \mathcal{D} respectively.

Outline:

- From what just happened to “Kleisli category is syntactic description of a monad”.
 - Taking stock of what happened with D and \mathcal{D} : starting from a syntactic description of probability, we recovered finitely-supported PMFs.
 - Plotkin+Power’s idea, inspired by Moggi’s proposal to use monads: this is a recipe for modelling effects more generally: start from a syntactic description, add equations one expects should hold, and try to see what mathematical objects come out (c.f. algebraic effects)
 - Examples:
 - * Failure: $e ::= v \mid \text{fail}$. This recovers the Maybe monad.
 - * Nondeterminism: $e ::= v \mid \text{amb}(e_1, e_2)$. This recovers the finite nonempty powerset monad.
 - * Logging: $e ::= v \mid \text{log}(s); e$ (where s is drawn from some set of strings). This recovers the writer monad.
 - * Global state (one Boolean): $e ::= v \mid x \leftarrow \text{get}_b; e(x) \mid \text{put}_b; e$. This recovers the state monad, if you work hard to find the right equations ¹.
- This approach of starting from syntax is essentially universal algebra. One nice consequence is that we can use results from universal algebra to study effects. We look at one example now: the idea of a Kleisli category of a monad.
 - Definition of Kleisli category. Morphisms $\mathcal{V} \rightarrow T\mathcal{W}$ are like substitutions, showing how to transform each value \mathcal{V} into a term $T\mathcal{W}$.
 - Example: for Dist, Kleisli morphisms are transition matrices. In the continuous case, Kleisli morphisms are Markov kernels.
 - Example: for Maybe, Kleisli morphisms are partial functions.
- Given a monad, find its Kleisli category and corresponding adjunction.
- The Kleisli monad models call-by-value computation. Universal algebra shows how to model call-by-name: Eilenberg-Moore, from the perspective of “models of tiny language”.

¹ Gordon Plotkin and John Power. Notions of computation determine monads. In *International Conference on Foundations of Software Science and Computation Structures*, pages 342–356. Springer, 2002

Realizability

Realizability originated from Kleene [18] as a semantics for Heyting arithmetic. In this article, Kleene motivates the concept by calling logical statements “incomplete communication” and “partial judgment” from the perspective of constructivists and intuitionists. For instance, a proposition “ A implies B ”, when interpreted intuitionistically, is an incomplete communication of another statement which gives a general procedure that: when receiving the information that completes A , it finds the information that completes B . In other words, Kleene is seeking the computational content of logical statements. To achieve this, Kleene presents a way of encoding the items of information with partial recursive functions¹. Then, we may say that certain recursive function *realizes* a statement, as a kind of intuitionistic truth notion of the statement.

In the context of PL research, the more relevant way of viewing realizability is motivated by a generalized question: “given a mathematical structure, what should its implementation look like?” [4] Such interpretation of realizability is of great interest in the area of compiler correctness [5, 29], since it naturally establishes a relation between implementations and specifications. Also, instead of partial recursive functions in Gödel numbers, we will use a generalization of the λ -calculus that allows applications to be partially defined for our realizers, namely the *partial combinatory algebra*, or PCA. But before defining PCAs, let’s first introduce an even more generalized structure underlying them, namely the *partial applicative structure*, or PAS.

Definition 68: Partial applicative structure (PAS)

A *partial applicative structure* (A, \cdot) is a set A equipped with a *partial* binary operation:

$$\cdot : A \times A \rightarrow A$$

This chapter was written by Ming-tong Lin. It is heavily based on [4]. The reader is only assumed to have acquaintance with previous chapters.

¹ In Gödel numbers, which is what you will get if you search for “realizability” on Wikipedia.

The discussion on PAS, PCA, and combinatory completeness is largely taken from [15].

For $f, a \in A$, if $(f, a) \in \text{dom}(\cdot)$, we write $f \cdot a \downarrow$, otherwise $f \cdot a \uparrow$. Sometimes people may omit \cdot and just write fa for $f \cdot a$. Application associates to the left, just like that in λ -calculus, so abc stands for $(ab)c$.

The definition of PAS reveals a key feature of combinatorial structures. Consider $f, a \in A$ and the application $f \cdot a$, the element f acts as a function, while the element a acts as an argument. In other words, the program (functions) and the data (arguments) live in the same realm and are represented in the same way.

Using our intuition from λ -calculus, a PAS defines an operation analogous to λ -applications, but what about λ -abstractions? To talk about them, we first need to incorporate variables.

Definition 69: Terms (PAS)

Let $V = \{x_0, x_1, \dots\}$ be a countable set of fresh variables, the set of *terms* over A , written $T(A)$, is the least set satisfying:

- if $x \in V$, then $x \in T(A)$;
- if $a \in A$, then $a \in T(A)$;
- if $t_1, t_2 \in T(A)$, then $t_1 \cdot t_2 \in T(A)$.

We will use the familiar notations from the λ -calculus for free variables and substitutions.

Notation 18.0.1. $FV(t)$ denotes the set of free variables in the term $t \in T(A)$. A term t is *closed* if $FV(t) = \emptyset$.

Notation 18.0.2. $t[a/x]$ stands for “substituting a for x in term t ”, exactly as in the λ -calculus.

Unfortunately, we are still not yet able to get something like λ -abstraction, but what exactly is going wrong?

Example 18.0.3. Let $A := \{1, 2\}$ and $\cdot : A \times A \rightarrow A$ be natural number multiplication. Of course, we are able to write square of elements in A in our PAS (A, \cdot) :

$$\begin{aligned} 1 \cdot 1 &\downarrow \\ 2 \cdot 2 &\uparrow \end{aligned}$$

Now, suppose we want to abstract over the squaring operation, immediately we try to write a term $t \in T(A)$ such that

$$t = x \cdot x$$

Next, we need an element $f \in A$ in our PAS that behaves like t , so that we can actually apply it for the square of an arbitrary a :

$$f \cdot a = t[a/x] = a \cdot a$$

However, by inspection, none of the elements in A suffices.

The above example reveals the discrepancy that, when given a term $t \in T(A)$, we might fail to find an element that behaves like it in A . In order to fill the gap, we shall first formalize what “behaves like” generally means for partial functions.

Unfortunately, the notation for Kleene equality collides with that for categorical equivalence. We will reserve \simeq for Kleene equality exclusively throughout this chapter.

Definition 70: Kleene equality

Kleene equality establishes the quality between partial functions. Given an argument, either both functions are undefined, or both are defined and their values on that arguments agree. We write $f \simeq g$ for Kleene equality.

Then, we are able to state what it means to have an element $a \in A$ that behaves like a specific term t .

Definition 71: Representability

Let $t \in T(A)$ be a term with $FV(t) \subseteq \{x_0, \dots, x_n\}$ and $f \in A$ an element of A , we say that f represents t if for all $a_0, \dots, a_n \in A$:

- $f \cdot a_0 \cdot \dots \cdot a_{n-1} \downarrow$
- $t[a_0/x_0, \dots, a_n/x_n] \simeq f \cdot a_0 \cdot \dots \cdot a_n$

In words, f is a total function in the first n arguments, and as an $n+1$ -ary partial function, f is equal to t .

When a term t is representable (i.e., represented by one or more elements in A), we have the internalization of t in A . In other words, if all of the terms $t \in T(A)$ are representable, we can always find at least one element that encodes our desired abstraction (e.g., the squaring operation in Theorem 18.0.3).

Definition 72: Combinatory completeness

A PAS (A, \cdot) is *combinatorially complete* if every term $t \in T(A)$ can be represented by some $a \in A$.

Remark 18.0.4. From the perspective of recursion theory, the definition of PAS abstracts the concept of the *enumeration theorem* [19]²: let A be \mathbb{N} (the Gödel numbers) and define \cdot to be the universal function Φ , we have $f \cdot a := \Phi(f, a)$.

² Note that a PAS is not necessarily *universal*, therefore we are not referring to the *utm theorem*.

Conversely, combinatory completeness corresponds to the *smn theorem*. Thus, an equivalent definition of combinatory completeness is: for a term $t \in T(A)$ and a variable x , there is a term t' such that $x \notin FV(t')$ and $t' \downarrow$, satisfying $t' \cdot a \simeq t[a/x]$ for all $a \in A$.

Definition 73: Partial combinatory algebra (PCA)

A PAS (A, \cdot) is called a *partial combinatory algebra* if it is combinatory complete.

Proposition 18.0.5. A PAS (A, \cdot) is combinatory complete if and only if there exists $K, S \in A$ such that for all $a, b, c \in A$:

- $K \cdot a \cdot b = a$, and
- $S \cdot a \cdot b \cdot c \simeq (a \cdot b) \cdot (b \cdot c)$ and $S \cdot a \cdot b \downarrow$.

Fact 18.0.6. The first Kleene algebra \mathcal{K}_1 , combinatory logic **CL**, and untyped λ -calculus Λ are all PCAs.

Availing us with the language of PCAs, we can now return to the discussion of realizability.

Definition 74: Assembly

An *assembly* X over a PCA (A, \cdot) is a pair $(|X|, \Vdash_X)$ where $|X|$ is the underlying set and \Vdash_X is a relation between A and $|X|$, such that for every $x \in |X|$, there exists an $r \in A$ and $r \Vdash_X x$. We read $r \Vdash_X x$ as “ r realizes x ”.

The assembly requires every $x \in |X|$ to be realized by some $r \in A$, but there is no constraint about their correspondence. Therefore, an $x \in |X|$ may be realized by one or more realizers in A , and an $r \in A$ may also realize one or more elements in $|X|$.

Definition 75: Tracker

For two assemblies $(|X|, \Vdash_X)$ and $(|Y|, \Vdash_Y)$ over the same PCA (A, \cdot) , and a function $f : |X| \rightarrow |Y|$, a *tracker* for f is an element $t \in A$ such that for every $a \Vdash_X x$, we have $t \cdot a \downarrow$ and $t \cdot a \Vdash_Y f(x)$. We say that f is *realized* (or *tracked*) by t .

Definition 76: Assembly map

An *assembly map* $X \rightarrow Y$ between assemblies X and Y over a PCA

The usual way of defining PCAs is to say that a PCA is a PAS with $K, S \in A$. This simplification is due to Curry [10], known as the *bracket abstraction* algorithm. For more related details in the context of PCAs, see [25].

John always has a comment that “when you ask Amal about realizability, she will likely tell you it’s just *binary logical relations*.” Indeed, an equivalent formulation is to view the realizability relation as the subset: $\Vdash_X \subseteq A_\tau \times |X|$. Note that we just sneaked in types for our PCA (we intend to not discuss typed-PCAs [23, 24] due to the lack of space). Then, we can write the realizability relation as a binary \mathcal{V} -relation:

$$(r, x) \in \mathcal{V}[\tau]$$

\mathbb{A} is the pair (f, t) where $f : |X| \rightarrow |Y|$ and $t \in A$ tracks f .

Definition 77: Assembly category (Asm)

Assemblies over a PCA $\mathbb{A} = (A, \cdot)$ form a category $\text{Asm}(\mathbb{A})$ where:

- objects are assemblies $X = (|X|, \Vdash_X)$ over \mathbb{A} ;
- morphisms are assembly maps;
- for two morphisms (f, t) and (g, t') where $f : |X| \rightarrow |Y|$ and $g : |Y| \rightarrow |Z|$, composition is defined as

$$(g, t') \circ (f, t) = (g \circ f, B \cdot t' \cdot t)$$

where $B \in A := S \cdot (K \cdot S) \cdot K$;

- the identity morphism is given by the pair (id, I) , where id is the identity function and $I \in A := S \cdot K \cdot K$.

The combinator B used in the above definition is not in the standard SKI basis of combinatory logic, but is a handy combinator for composition:

$$B \cdot g \cdot f \cdot x \simeq g \cdot (f \cdot x)$$

Example 18.0.7 (Unit assembly $\mathbb{1}$). The *terminal* assembly $\mathbb{1}$ over \mathbb{A} is $(\{\star\}, \Vdash_{\mathbb{1}})$, where for all $r \in A$, we have the trivial realizability relation $r \Vdash_{\mathbb{1}} \star$. The unit assembly $\mathbb{1}$ is the terminal object in $\text{Asm}(\mathbb{A})$.

Previously, we said that a realizer $a \in A$ may realize more than one elements in $|X|$, which is not ideal for some cases. Therefore, we introduce a restricted version of an assembly in which no two elements share a common realizer.

Definition 78: Modest set

A *modest* assembly X , also called a *modest set*, is an assembly in which elements do not share realizers. That is, for any two $x, y \in |X|$, if an $r \in A$ realizes both x and y , we have $x = y$.

Fact 18.0.8. Modest sets over a PCA \mathbb{A} form a full subcategory $\text{Mod}(\mathbb{A})$ of $\text{Asm}(\mathbb{A})$.

Proposition 18.0.9. Both $\text{Asm}(\mathbb{A})$ and $\text{Mod}(\mathbb{A})$ are cartesian-closed. The proofs is left as an exercise for the readers ([4, Chapter 3.6] has some of the proofs).

Definition 79: Partial equivalence relation (PER)

Given a modest set X , we may induce a *partial equivalence relation* \approx_X on A which relates r and r' when they realize the same element $x \in |X|$:

$$r \approx_X r' \iff \exists x \in |X|. r \Vdash_X x \wedge r' \Vdash_X x$$

Definition 80: Total realizer

A realizer $r \in A$ is *total* if $r \approx_X r$. The set of total realizers is denoted by $\|X\| = \{r \in A \mid r \approx_X r\}$. Each $r \in \|X\|$ determines the equivalence class $[r]_{\approx_X} = \{a \in A \mid r \approx_X a\}$.

The name for total realizers is confusing—it is not the converse of “partial” in the computational sense. Replacing the other realizer r' with another r in the above definition reduces to:

$$r \approx_X r \iff \exists x \in |X|. r \Vdash_X x$$

which amounts to the fact that r realizes something in $|X|$, and therefore is not some syntactic “junk” in A . A total realizer r can definitely realize a partial function such as integer division. However, a realizer for $n/0$ is not total, since $n/0$ is absent in $|X|$.

Example 18.0.10. Let us consider the modest assembly $B = (\{\top, \perp\}, \Vdash_B)$ over some PCA \mathbb{A} , and the induced PER \approx_B . The set of total realizers $\|B\|$ contains those who realize something in $|B|$, namely for all $r \in |B|$, either $r \Vdash_B \top$ or $r \Vdash_B \perp$. Then, we have the following two equivalence classes:

$$\begin{aligned} [r_\top]_{\approx_B} &= \{r \in A \mid r_\top \approx_B r\} && \text{where } r_\top \Vdash_B \top \\ [r_\perp]_{\approx_B} &= \{r \in A \mid r_\perp \approx_B r\} && \text{where } r_\perp \Vdash_B \perp \end{aligned}$$

Note that since B is modest, $[r_\top]_{\approx_B}$ and $[r_\perp]_{\approx_B}$ are disjoint.

Definition 81: Extensional realizer

Given two PERs \approx_X and \approx_Y , an *extensional realizer* between them is a $p \in A$ such that, for all $r, r' \in A$, if $r \approx_X r'$ then $p r \downarrow, p r' \downarrow$, and $p r \approx_Y p r'$. Extensional realizers p and p' are *equivalent* when $r \approx_X r'$ implies $p r \approx_Y p' r'$.

Remark 18.0.11. An extensional realizer therefore “sends related inputs to related outputs”, where the relation is induced by the modest set.

Example 18.0.12. Let us consider an extensional realizer between B and itself. We have already discussed the set of total realizers and equivalence classes of B above.

An example extensional realizer p sends the input to the same equivalence class for the output: if the input is in $[r_{\top}]_{\approx_B}$, then the output is also in $[r_{\perp}]_{\approx_B}$, and vice versa. If our \mathbb{A} is λ -calculus, we have $p \ r \ \downarrow$ and $p \ r' \ \downarrow$ for free since application of two λ terms is always defined. One qualified p is the identity $\lambda x. x$, and the extensionality condition is trivially true.

Notation 18.0.13. In stead of writing \approx_X for the PER induced by the modest set X , it is common to denote it by X directly. This convention is slightly ambiguous, but is somewhat reasonable since the correspondence is one-to-one.

Remark 18.0.14. Note that in the discussion of extensional realizers, we did not mention what in $|X|$ is realized by p . This is fundamentally different from what we did when defining trackers, where we said $f : |X| \rightarrow |Y|$ is realized by t . In fact, the underlying set $|X|$ only appeared at the very beginning for defining the PER itself throughout our discussion of PERs.

This marks a fundamental perspective shift. In the assembly view, we relate “implementations” (the realizers) to “specifications” (the mathematical objects in $|X|$). In the PER view, we relate implementations to implementations, by quotienting them through a PER.

Definition 82: PER category $\text{Per}(\mathbb{A})$

PERs and equivalence classes of extensional realizers form a category $\text{Per}(\mathbb{A})$ where:

- objects are PERs X on \mathbb{A} ;
- morphisms $X \rightarrow Y$ are equivalence classes of extensional realizers $[p] : X \rightarrow Y$;
- for two morphisms $[p] : X \rightarrow Y$ and $[q] : Y \rightarrow Z$, composition is defined as

$$[q] \circ [p] = [q \circ p]$$

where $[q \circ p]$ is defined as $[B \cdot q \cdot p]$.

- the identity morphism is I .

Proposition 18.0.15. The categories $\text{Mod}(\mathbb{A})$ and $\text{Per}(\mathbb{A})$ are equivalent.

For details of the equivalence proof, consult [4, Chapter 3.3.3]. There is also an Agda mechanization of this proof in [7].

Definition 83: Uniform assembly

Let X be a set and \mathbb{A} a PCA, the *uniform assembly* (or *constant assembly*)

$$\nabla_{\mathbb{A}} X = (X, \Vdash_{\nabla_{\mathbb{A}} X})$$

is the assembly in which the elements are “uniformly realized”. That is, for all $r \in A$ and $x \in X$:

$$r \Vdash_{\nabla_{\mathbb{A}} X} x$$

Notation 18.0.16. When not ambiguous, we write ∇X for $\nabla_{\mathbb{A}} X$.

Fact 18.0.17. Uniform assemblies over a PCA \mathbb{A} form a full subcategory of $\text{Asm}(\mathbb{A})$.

Proposition 18.0.18. The unit assembly $\mathbb{1}$ is uniform.

In a uniform assembly ∇X , every $x \in X$ is realized by every $r \in A$, meaning that the realizability condition is trivial and the realizers provide *no information* about the elements in X . Think about the trackers: normally, a tracker t of f needs to carefully maintain the condition that when applied to an argument that realizes a valid input x of f , it should realize the result $f(x)$. But within uniform assemblies, the condition collapses to only require $t \cdot a \downarrow$, since it is guaranteed to realize $f(x)$. In this sense, ∇X is devoid of any computational content.

Proposition 18.0.19. Given an assembly X and a uniform assembly ∇Y over a PCA \mathbb{A} , every function $|X| \rightarrow Y$ is an assembly map $X \rightarrow \nabla Y$.

Proof. For a function $f : |X| \rightarrow Y$, we pair it with a trivial tracker $\mathbb{1}$, where for all $a \in A$, $\mathbb{1} \cdot a$ is always defined. \square

Definition 84: Uniform assembly functor (∇)

Let $\nabla : \text{Set} \rightarrow \text{Asm}(\mathbb{A})$ be the functor defined as follows:

- on objects, it sends a set X to the uniform assembly ∇X ;
- on Set morphism $f : X \rightarrow Y$, $\nabla f = f$.

Proposition 18.0.20. The functor ∇ is fully faithful.

Proposition 18.0.21. An assembly X is uniform if and only if it is isomorphic to some ∇Y . Therefore, the functor $\nabla_{\mathbb{A}}$ is essentially surjective onto the full subcategory of uniform assemblies over a PCA \mathbb{A} .

Corollary 18.0.22. The category of uniform assemblies over a PCA \mathbb{A} is equivalent to Set .

Definition 85: Forgetful functor (Γ)

Let $\Gamma : \text{Asm}(\mathbb{A}) \rightarrow \text{Set}$ be the forgetful functor defined as follows:

- on objects, $\Gamma X = |X|$;
- on assembly map $(f, t) : X \rightarrow Y$, $\Gamma(f, t) = f$.

Proposition 18.0.23. The the forgetful functor Γ is left adjoint to the uniform assembly functor ∇ .

Proof. Recall the natural transformations $\eta : \text{id}_{\text{Asm}(\mathbb{A})} \Rightarrow \nabla\Gamma$ and $\varepsilon : \Gamma\nabla \Rightarrow \text{id}_{\text{Set}}$ are the unit and counit of the adjunction.

- For any $X \in \text{Asm}(\mathbb{A})$, the component of the unit $\eta_X : X \rightarrow \nabla\Gamma X$ is defined by the assembly map $(\text{id}_{|X|}, \mathbf{I})$, where \mathbf{I} is a trivial tracker.
- For any $S \in \text{Set}$, the component of the counit $\varepsilon_S : \Gamma\nabla S \rightarrow S$ is given by the identity function id_S .
- Show $\varepsilon\Gamma \circ \Gamma\eta = \text{id}_\Gamma$: for an assembly $X \in \text{Asm}(\mathbb{A})$,

$$(\varepsilon\Gamma \circ \Gamma\eta)_X = \varepsilon_{\Gamma X} \circ \Gamma(\eta_X) = \text{id}_{|X|} \circ \text{id}_{|X|} = \text{id}_{|X|}$$

- Show $\nabla\varepsilon \circ \eta\nabla = \text{id}_\nabla$: for a set $S \in \text{Set}$,

$$(\nabla\varepsilon \circ \eta\nabla)_S = \nabla(\varepsilon_S) \circ \eta_{\nabla S} = (\text{id}_S, \mathbf{I}) \circ (\text{id}_S, \mathbf{I}) = (\text{id}_S, \mathbf{I})$$

□

Categorical Noninterference

Noninterference is a property in formal security to establish the safety of private information with respect to public information. Essentially, a program/language/protocol satisfies noninterference if publicly observable operations do not depend on private information; behavior using private data has no distinguishable effect on public data.

This property is very strong, and almost always too strong. One might be motivated to research different models for reasoning about noninterference, with the hope of obtaining a weakened (yet formal) version. As was mentioned at the beginning of this semester, the answer to understanding a problem is rarely to categorify it. But why not give it a try?

In this chapter, we explore Sterling and Harper's approach to noninterference via a sheaf semantics in their paper titled *Sheaf Semantics of Termination-Insensitive Noninterference*¹.

The category theory in this chapter is primarily sheaf-theoretic; the content is almost entirely self-contained. Sources and supplementary reading can be found in these course notes, §I and §II of *Sheaves in Geometry and Logic* by MacLane and Moerdijk, *Category Theory in Context* by Emily Riehl, and of course ncatlab.org.

19.1 Information Flow Basics

This section introduces some terminology and basic ideas around formal reasoning for information flow. We start with the notion of a *security lattice*. The simplest form is the set $\mathcal{L} = \{L, H\}$ with $L \leq H$, where L represents low-security (or public) data, and H represents high-security (or secret) data.

One way to think about this lattice with respect to programming is by "labelling" variables with some $\ell \in \mathcal{L}$, e.g. x_L represents that x is a low-security variable. This labelling can be implemented with *security*

This chapter was written by Bex Golanov.

¹ We do not cover anything pertaining to termination, sensitivity or lack thereof. Our exploration is also localized to pages 6 and 7 of the paper, as well as one observation on page 12 which we prove at the end of this chapter.

types. For example, if we want to talk about x_L of type `Int` with this framework, we would say x has type Int_L , where Int_L represents public ints.

When working with information flow, we ideally want use of public (observable) data to not reveal anything about secret data. The strongest statement of this is requiring that public data never depends on secret data, usually referred to as noninterference.

Definition 86: ℓ -equivalence

Let $L : \text{Var} \rightarrow \mathcal{L}$ be a labelling from variables x to elements ℓ of the security lattice \mathcal{L} .

A state $\sigma : \text{Var} \rightarrow \text{Val}$ maps variables to values. Two states σ, σ' are ℓ -equivalent, denoted $\sigma \sim_\ell \sigma'$, when for all $\ell' \leq \ell$, if $L(x) = \ell'$, then $\sigma(x) = \sigma'(x)$. This requirement is equivalent to $L(x) \leq \ell$ for all x , which we can interpret as x being "public with respect to ℓ ".

Definition 87: Noninterference

A program α satisfies *noninterference* if, for any states σ_1, σ_2 such that $\sigma_1 \sim_L \sigma_2$, and $\sigma_1 \xrightarrow{\alpha} \sigma'_1, \sigma_2 \xrightarrow{\alpha} \sigma'_2$, we have $\sigma'_1 \sim_L \sigma'_2$. If initial states agree on all public values, the respective post-states should agree on all public values.

Let **Rel** be the category whose objects are sets and whose morphisms $A \sim_R B$ are relations $\sim_R \subseteq A \times B$. We can think of ℓ -equivalence as a morphism in **Rel**, where $\sigma \sim_\ell \sigma'$ means $(\sigma, \sigma') \in \sim_\ell \subseteq \Sigma \times \Sigma'$ for Σ, Σ' sets of states. Since functions $f : A \rightarrow B$ can be represented as graphs $f \subseteq A \times B$, where $(a, b) \in f$ if $f(a) = b$, we can consider functions to be relations as well, so $\alpha : \Sigma \rightarrow \Sigma'$ is also a morphism in **Rel**.

Then noninterference can be understood as the requirement that the following diagram commutes:

$$\begin{array}{ccc}
 \sigma_1 & \xleftrightarrow{\sim_\ell} & \sigma_2 \\
 \alpha \downarrow & & \downarrow \alpha \\
 \sigma'_1 & \xleftrightarrow{\sim_\ell} & \sigma'_2
 \end{array}$$

This property is also referred to as σ_1 and σ_2 being *indistinguishable* with respect to public values.

This indistinguishability property is equivalent to requiring any map from some type τ_H to type τ_L to be constant. One characterization of this definition is the following:

Proposition 19.1.1. For any closed function $\cdot \vdash f : \tau_H \rightarrow \tau_L$, there exists a closed $\cdot \vdash v : \tau_L$ such that $f \simeq \lambda _ . v$, meaning f is observationally (extensionally) equivalent to a constant v function.

19.2 Categorical Background Potpourri

In this section, we explain the medley of categorical concepts needed to understand Sterling and Harper's construction.

19.2.1 Lattices and Logic

Definition 88: Lattices

A lattice \mathcal{L} is a partially ordered set which has all *meets* and *joins*.

- Two elements $l, k \in \mathcal{L}$ have a *meet* if there exists an $m \in \mathcal{L}$ such that m is the greatest lower bound of l and k . *Joins* are symmetrically least upper bounds.
- Having all meets and joins means for any pair of elements in the lattice, you can find a greatest lower bound as well as a least upper bound. This property can also be phrased as having all binary products and all binary coproducts.

A lattice with $\mathbf{0}$ and $\mathbf{1}$ (unique initial and terminal objects) has all finite limits and all finite colimits.

Intuitionistic propositional logic can be thought of as a *Heyting algebra* H , which is a lattice with $\mathbf{0}$ and $\mathbf{1}$, and has for each pair x, y an exponential object y^x , usually written as $x \Rightarrow y$. Disjunction and conjunction are joins and meets, respectively, with the initial object $\mathbf{0}$ and terminal object $\mathbf{1}$ as their respective identities. That is,

$$x \vee \mathbf{0} = x, \quad x \wedge \mathbf{1} = x$$

A *distributive lattice* is a lattice in which

$$x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$$

holds for all $x, y, z \in \mathcal{L}$. This identity also implies the dual, $x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$.

The exponential $x \Rightarrow y$ corresponds to implication. Writing out our usual diagram in the above language,

$$\begin{array}{ccc} z \wedge x & & \\ \downarrow & \searrow & \\ y^x \wedge x & \longrightarrow & y \end{array}$$

we can characterize y^x or $x \Rightarrow y$ by the following:

$$z \wedge x \leq y \iff z \leq x \Rightarrow y$$

We can interpret $x \Rightarrow y$ as a least upper bound for z such that $z \wedge x \leq y$. Negation in a Heyting algebra can be defined in the "implies absurdity" manner:

$$\neg x := x \Rightarrow \mathbf{0}$$

With our definition of \Rightarrow , this implies

$$y \leq \neg x \iff y \wedge x = \mathbf{0},$$

where we have $=$ instead of \leq because $\mathbf{0}$ is initial.

Definition 89: Topology

A *topology* on a set X can be defined as a collection \mathcal{T} of subsets of X , called **open sets** and satisfying the following axioms:

1. The empty set and X itself belong to \mathcal{T} .
2. Any arbitrary union of members of \mathcal{T} belongs to \mathcal{T} .
3. The intersection of any finite number of members of \mathcal{T} belong to \mathcal{T} .


For a topological space X , the set of all open sets of X , notated $\mathcal{O}(X)$, forms a Heyting algebra. This implies we can model the intuitionistic propositional calculus with any topological space X , by looking at the associated set $\mathcal{O}(X)$.

An alternative (cursed) answer to the question "what is a topology?": a topology is *just* something you define sheaves over, and sheaves are *just* something you define a cohomology on, so the question you should really be asking is, what is a cohomology?

Definition 90: Special Subsets of Lattices

- An *upper set* U on a poset (\mathcal{P}, \leq) is upwards closed: if $x \in U$ and $x \leq z$, then $z \in U$.
- Symmetrically, a *lower set* U on \mathcal{P} is downwards closed: if $x \in U$ and $z \leq x$, then $z \in U$.
- A *filter* \mathcal{F} on \mathcal{P} is an upper set with the additional condition of every finite subset of \mathcal{F} having a lower bound in \mathcal{F} (filters are closed under finite meets)

Example 19.2.1. Recalling that the natural numbers \mathbb{N} form a poset ordered under divisibility, here is a Hasse diagram of the factors of 60, where the blue numbers are the elements of the upper set generated by 6 (and in this case the upper set is also a filter), and the orange are elements of the lower set generated by 10.



For an example of an upper set which is *not* a filter, consider the lattice of divisors of 12 ordered by division. The upper set defined by $U = \{4, 6, 12\}$ is not a filter, since the meet (in this case gcd) of 4 and 6 is 2, which is not contained in U . By contrast, $\{6, 12\}$ satisfies filter requirements

Definition 91: Alexandroff Topology

The *Alexandroff Topology* is a natural topology on the underlying set of a pre-ordered set, where the open sets are the upper sets.

Let's check that this definition satisfies the axioms for open sets.

Proposition 19.2.2. Upper sets of a poset \mathcal{P} satisfy the following axioms:

1. The empty set and \mathcal{P} itself are upper sets.
2. Upper sets are closed under arbitrary unions.
3. Upper sets are closed under finite intersections.

Proof. .

(Foreshadowing): If we order filters \mathcal{F} on a poset \mathcal{P} by inclusion, we get a poset of the set of filters. This would imply there exists an Alexandroff topology to be defined here, with upper sets of the filters.

1. \emptyset as an upper set is vacuously true; there are no elements to check. For \mathcal{P} , suppose $x \in \mathcal{P}$ and $x \leq y$. Well $y \in \mathcal{P}$ because \mathcal{P} is the entire set.
2. Let $\{U_i\}_{i \in I}$ be a collection of upper sets. We wish to show $\bigcup_{i \in I} U_i$ is an upper set. Suppose $x \in \bigcup_{i \in I} U_i$ and $x \leq y$. Then $x \in U_j$ for some $j \in I$. U_j is an upper set and $x \leq y$, so $y \in U_j$. Then $y \in \bigcup_{i \in I} U_i$ so arbitrary union of upper sets is also an upper set.
3. Let U, V be upper sets. We want to show $U \cap V$ is an upper set. Suppose $x \in U \cap V$ and $x \leq y$. Then $x \in U$ and $x \in V$. Since U and V are both upper sets and $x \leq y$, we have $y \in U$ and $y \in V$. Then $y \in U \cap V$! So upper sets are also closed under finite intersection.

□

19.2.2 Sheaves

In the sheaf semantics chapter we defined a *coverage*, which was defined to be a relation T between families of morphisms with common codomain c and objects c of \mathcal{C} . If a family $\{f_i : c_i \rightarrow c\}$ and object c are in T , then $\{f_i : c_i \rightarrow c\}$ is a T -cover of c . Conditions for T can be found in the aforementioned chapter. A general intuition for these families of morphisms is to think of them as a way to partition an object c by images of f_i .

Definition 92: What is a sheaf, actually?

Let \mathcal{C} be a category with a coverage T . For any object $A \in \mathcal{C}$, we have a T-cover $\{a_i : A_i \rightarrow A\}_{i \in I}$ indexed by some I . A presheaf F on \mathcal{C} is a sheaf if it satisfies two additional conditions for *any* T-cover:

1. *locality*: For any $f, g \in F(A)$, if $f|_{A_i} = g|_{A_i}$ for all $i \in I$, then $f = g$. The notation $f|_{A_i}$ is equivalent to applying $F(p_i)(f)$ where p_i is the morphism $A_i \rightarrow A$.

If two elements agree locally (on each part of the cover), they agree everywhere. We can also think of this as being able to partition any $f \in F(A)$ into the "sum of its parts", where the parts correspond to an arbitrary cover of A .

2. *gluing*: Let $f \in F(A)$. For any $i, j \in I$, we have $F(\pi_i)(f|_{A_i}) = F(\pi_j)(f|_{A_j})$ in $F(A_i \times_A A_j)$, where π_i is the projection $A_i \times_A A_j \rightarrow A_i$, π_j is the projection for A_j and $A_i \times_A A_j$ is the pullback of A along $f|_{A_i}$ and $f|_{A_j}$.

If \mathcal{C} is a preorder with meets, this simplifies to for any $x \in A_i \cap A_j$, we have $f|_{A_i} x = f|_{A_j} x$.

- Consider the case where $F(A)$ maps to sets of functions. In this case, an equivalent definition to the above which may have more of a "gluing" connotation says that, for any $f_i : A_i \rightarrow A$ and $f_j : A_j \rightarrow A$ such that $f_i(a) = f_j(a)$ for $a \in A_i \cap A_j$, there exists a unique f such that $f|_i = f_i$ and $f|_j = f_j$. That is, when we have two maps agreeing at their intersection, we can "glue" them together to get another map.

Proposition 19.2.3. Let \mathcal{C} be a category with a coverage and $F : \mathcal{C}^{op} \rightarrow \mathbf{Set}$. For any $A \in \mathcal{C}$, for any cover $\{a_i : A_i \rightarrow A\}$ of A , F is a sheaf iff F satisfies the equalizer diagram

$$F(A) \xrightarrow{e} \prod_{i \in I} F(A_i) \begin{array}{c} \xrightarrow{p} \\ \xrightarrow{q} \end{array} \prod_{i, j \in I} F(A_i \cap A_j),$$

where $p(f_i) = f_i|_{A_i \cap A_j}$ and $q(f_i) = f_j|_{A_i \cap A_j}$.

Brief sketch on why the above is true:

Lemma 19.2.4. e (or any equalizing map) is a monomorphism. Conversely, the morphism associated to the coequalizer is an epimorphism. These are not too hard to prove.

Can see Solution 19.4.4

- Since F maps to \mathbf{Set} , e is also injective. We can think of e as the map that partitions some $f \in F(A)$, i.e. $e(f) = \{f|_{A_i}\}_{i \in I}$. Then

agreeing locally, $f|_{A_i} = g|_{A_i}$, is equivalent to saying $e(f) = e(g)$. Injectivity of e gives us $f = g$.

- The equalizer-ness of F means $F(A)$ only consists of those f where $p \circ e(f) = q \circ e(f)$. We are guaranteed $(f|_{A_i})|_{A_i \cap A_j} = (f|_{A_j})|_{A_i \cap A_j}$, which is equivalent to saying $f|_{A_i} x = f|_{A_j} x$ for all $x \in A_i \cap A_j$.

Definition 93: Topological Vocabulary

- Let U be an open subset of some S with topology \mathcal{T} . The family of open sets $C = \{U_i\}_{i \in I}$ indexed by some I is an *open cover* of U iff $\bigcup_{i \in I} U_i = U$.
- A family of open subsets $\{B_i\}_{i \in I}$ of S form a *basis* for \mathcal{T} iff for any $U \subset S$, there exists some $J \subset I$ such that $\{B_i\}_{i \in J}$ is an open cover of U . That is, any element of \mathcal{T} (open set) can be represented as a union of some subset of the basis elements.

For a space S , the topological definition of a covering corresponds to a coverage on $\mathcal{O}(S)$, where the families of morphisms are inclusion maps. Recall $\mathcal{O}(S)$ denotes the set of open sets of S , which form a poset (in fact, a heyting algebra) under inclusion.

Taking C as an open cover of U as above, we can write C as $\{U_i \hookrightarrow U\}_{i \in I}$.

Definition 94: Subterminal Sheaf

Fix a category \mathcal{C}

- Given a functor F , the functor G is a *subfunctor* of F if $G(C) \subseteq F(C) \forall C \in \mathcal{C}$. If F is a sheaf on \mathcal{C} , then G is a *subsheaf* if and only if for every open U of \mathcal{C} and every element $f \in FU$, and every open covering $U = \bigcup U_i$, one has $f \in GU$ if and only if $f|_{U_i} \in GU_i$ for all i .
- The terminal sheaf is the terminal object $\mathbf{1}$ of $Sh(\mathcal{C})$. A *subterminal sheaf* F is a subsheaf of $\mathbf{1}$. F is a subfunctor, meaning $FC \subseteq \mathbf{1}C = \mathbf{1}$ for all $C \in \mathcal{C}$. This means FC can only be either $\mathbf{1}$ or \emptyset .

Proposition 19.2.5. A *subobject* of a sheaf F in the category $Sh(\mathcal{C})$ is isomorphic to a subsheaf of F .

Proof. See §II.3 in *Sheaves in Geometry and Logic* by MacLane and Moerdijk. □

Theorem 19.2.6. For any space X , there is an isomorphism

$$\mathcal{O}(X) \cong \text{Sub}_{\text{Sh}(X)}(\mathbf{1})$$

of partially ordered sets.

Proof. \Rightarrow

Given any open set W of X , define a functor S_W on open sets U by $S_W(U) = 1$ if $U \subseteq W$ and \emptyset otherwise. Recall $\mathbf{1}U = 1$ for all open U . Let $\{V_i\}$ be an arbitrary covering of U . If $S_W(V_i) = 1$ for all i , then $\bigcup V_i \subseteq W$, meaning $U \subseteq W$, so $S_W(U) = 1$. In the other direction, if $S_W(U) = 1$, all $V_i \subseteq U$ so $S_W(V_i) = 1$ for all i . So, we can see S_W is a subsheaf of $\mathbf{1}$.

\Leftarrow

Let S be a subsheaf of $\mathbf{1}$. Each $S(U)$ must then be either 1 or \emptyset . If $S(U) = 1$ and $V \subseteq U$, and S requires a morphism $S(U) \rightarrow S(V)$ meaning $1 \rightarrow S(V)$, it follows that $S(V)$ is nonempty and thus $S(V) = 1$. By the equalizer condition, if $\{V_i\}$ is an open cover of U and $SU_i = 1$ for all i , then $SU = 1$.

Define $W = \bigcup\{U \in \mathcal{O}(X) \mid SU = 1\}$. Then by equalizer condition we have $SU = 1$ if and only if $U \subseteq W$. Then we have $S = S_W$, so we have the bijection $W \iff S_W$. □

19.2.3 Special Sheaf Functors

The crux of the Sterling and Harper’s information flow model is based on the following monad definition:

"For any subspace $Q \subset P$, a sheaf $A \in \text{Sh}(P)$ can be restricted to Q , as in $A|_Q \in \text{Sh}(Q)$, and then extended again to P . This composite defines an *idempotent* monad on $\text{Sh}(P)$ that we can interpret as removing any data from P that cannot be seen from Q ."

This construction is a specialization of a much more general property of sheaves, which we need to do some work to define.

A monad T is idempotent when $T^2 = T$ (i.e. $\mu = \text{iso}$).

Definition 95

Let X, Y be topological spaces. A function $f : X \rightarrow Y$ is *continuous* if, for every open $V \subseteq Y$, the preimage $f^{-1}(V)$ is open in X .

Exercise 1

All monotone maps in the Alexandroff topology are continuous.
See Solution 19.4.1

Definition 96: Direct Image Sheaf Functor

Given a continuous f , each sheaf F on X yields a sheaf f_*F on Y defined for V open in Y , by $(f_*F)V = F(f^{-1}V)$. Essentially, f_*F is the composite functor

$$\mathcal{O}(Y)^{op} \xrightarrow{f^{-1}} \mathcal{O}(X)^{op} \xrightarrow{F} \mathbf{Set}$$

This sheaf is called the *direct image* of F under f . The map $f_* : Sh(X) \rightarrow Sh(Y)$ is a functor of sheaf categories.

A continuous $f : X \rightarrow Y$ also induces a functor $f^* : Sh(Y) \rightarrow Sh(X)$, where each sheaf G on Y yields a sheaf f^*G on X . We typically can't define f^* (the Inverse Image Functor) as simply as we were able to define f_* because we don't know if $f(U)$ is open (continuity only gives us information about preimages).

First we have the general and scary definition (though not the most general because we assume sheaves over topological spaces):

Definition 97: Inverse Image Sheaf Functor

Let $f : X \rightarrow Y$ be a continuous map of topological spaces. For any sheaf G on Y , we define the *inverse image* sheaf f^*G on X to be the *sheafification* of the presheaf $U \mapsto \text{colim}_{V \supseteq f(U)} G(U)$, where U is any open set in X , and the limit is taken over all open sets V of Y containing $f(U)$.

Luckily for us, the only continuous maps we care about in this chapter are inclusion maps. Here is some sketchy / informal reasoning as to why $i^*F(U) = F(U)$ when $X \subseteq Y$ is open for open $U \subseteq X$.

- When $X \subseteq Y$ is open any open $U \subseteq X$ is open in Y as well.
- $F(U)$ would be the colimit of $F(V)$ where $V \supseteq i(U) = U$, because U is clearly the greatest lower bound (limit) on sets containing U , so dually $F(U)$ would be the least upper bound (colimit) on $F(V)$ with $V \supseteq U$, since F is contravariant.
- Then $i^*F(U) = F(U)$, and of course F is a sheaf so $i^*F(U)$ automatically satisfies sheaf conditions for open U .

If X is not open, we have a more complicated situation, where open $U \subseteq X$ is not guaranteed to be open in Y .

If we are working in a topology \mathcal{T} where for every (not necessarily open) subset V of a space Y has a smallest open $U \subseteq Y$ such that $V \subseteq U$, then $i^*F(V) = F(U)$. Let's notate this smallest open U for

In the measure theory/qbs lecture, we briefly discussed the pushforward of a measure along a measurable map: $\mu_B(b) = \mu_A(f^{-1}(b))$ where f^{-1} is the preimage. We can similarly think of μ_B as the composite functor $\Sigma_B \xrightarrow{f^{-1}} \Sigma_A \xrightarrow{\mu_A} [0, 1]$, where we might want to denote μ_B by $f_*\mu_A$, and say f_* is a functor of probability measures.

V by $o(V)$. To have i^*F satisfy sheaf conditions, we need covers to "behave well", that is, for $\bigcup_j V_j = V$, we need $\bigcup_j o(V_j) = o(V)$. In general this is not guaranteed.

However, in this chapter we will only be looking at the Alexandroff topology. In this topology the open sets are upper sets, so for any set V , we have $o(V) = \uparrow V$, where $\uparrow V$ represents the upwards closure of V , basically turning V into an upper set. The upwards closure \uparrow respects/distributes across unions, that is $\bigcup_j \uparrow V_j = \uparrow \bigcup_j V_j$. So we do not need to worry about sheafifying anything!

Theorem 19.2.7. If $f : X \rightarrow Y$ is a continuous map, then the functor f^* , sending each sheaf G on Y to its inverse image (preimage) on X , is left adjoint to the direct image functor f_* :

$$Sh(X) \begin{array}{c} \xrightarrow{f_*} \\ \xleftarrow{f^*} \end{array} Sh(Y) , \quad f^* \dashv f_*$$

Proof. See §II.9 in *Sheaves in Geometry and Logic* by MacLane and Moerdijk. □

Theorem 19.2.7 is included to note that $f_* \circ f^*$ is a monad on $Sh(Y)$. Particularly, when $X \subset Y$ and f is the inclusion map i , we have a monad given by $i_* \circ i^*$. It is precisely this composition which defines Sterling et al.'s idempotent monad.

The preimage of the inclusion map $i : Q \hookrightarrow P$ is defined $i^{-1}(V) = V \cap Q$, where V is an open subset of P . With this, let's compute this monad for any inclusion. Fix a sheaf F and V open in P .

$$\begin{aligned} i_*(i^*F)(V) &= (i^*F)(i^{-1}V) \\ &= (i^*F)(V \cap Q) \end{aligned}$$

Recall that in our special case, the functor i^* operates by $i^*F(U) = F(\uparrow U)$ for open U in subspace Q , where $\uparrow U$ is the upward closure of U in P . When Q is open in P , $i^*F(U) = F(U)$.

For any sheaf F , $F(\emptyset) = 1$. We can see this from the equalizer definition, where F in this case is the equalizer of the empty diagram, which is precisely the terminal object.

If $V \cap Q \neq \emptyset$, then $i_*(i^*F)(V) = F(\uparrow(V \cap Q))$. If $V \cap Q = \emptyset$, then $i_*(i^*F)(V) = 1$.


19.2.4 Pushouts

Let X be a topological space and Y an open set in X . One of the monads that Sterling et al. define is the composition of $i_* \circ i^*$ defined by $i : X \setminus Y \hookrightarrow X$. The authors mention that this monad $i_* \circ i^*$ can be identified with a pushout construction

$$Y \sqcup_{Y \times (-)} (-) : Sh(X) \rightarrow Sh(X),$$

where Y here is treated as the subterminal sheaf where $Y(U) = 1$ if $U \subseteq Y$, and \emptyset otherwise. This construction is key for later non-interference proofs, so we provide some intuition for how pushouts work/behave.

Pushouts are dual to pullbacks, but it is not very intuitive to get a sense of what that means exactly. We can try to work it out by staring at the diagram in **Set**:



Where P is the smallest object with morphisms i_C, i_B , such that $i_C \circ g = i_B \circ f$, i.e. $i_C(g(a)) = i_B(f(a))$ for all $a \in A$. Then we can see that this restriction is only for the images of f and g ; we don't have any equality condition on $b \notin im(f), c \notin im(g)$. These maps i_B, i_C are canonically *quotient maps* modulo the relation by f and g , meaning they map $b \in B, c \in C$ to equivalence classes $[b], [c] \in P$, with this extra equivalence condition.

So, we can interpret the commutativity of this square to mean that for b, c where there exists an a such that $f(a) = b, g(a) = c$, we require $i_B(b) = [b] = [c] = i_C(c)$.

In **Set**, P is a "quotient of the disjoint union of A and B ", which can also be written $A \sqcup B / \sim$, where \sim is the equivalence relation generated by the restriction detailed above, i.e. some \sim satisfying $b \sim c$ when $\exists a \in A$ such that $f(a) = b, g(a) = c$. The pushout essentially collapses the images of f and g to equivalence classes in P . We provide a couple of examples in **Set**.

Example 19.2.8 (pushout of product). Let B, C be nonempty sets.

$$\begin{array}{ccc}
 B \times C & \xrightarrow{\pi_1} & B \\
 \pi_2 \downarrow & & \downarrow i_B \\
 C & \xrightarrow{i_C} & B \sqcup_{B \times C} C
 \end{array}
 \cong
 \begin{array}{ccc}
 B \times C & \xrightarrow{\pi_1} & B \\
 \pi_2 \downarrow & & \downarrow i_B \\
 C & \xrightarrow{i_C} & \{*\}
 \end{array}$$

By definition of product, for every $b \in B$ and for every $c \in C$, there exists an $a \in A$, namely (b, c) , such that $\pi_1((b, c)) = b$ and $\pi_2((b, c)) = c$. Then every element of B is in the same equivalence class as every element of C , and by transitivity of equality, every element of B is in the same equivalence class, as well as every element

of C . Then P only has one equivalence class, so P is a singleton set and thus isomorphic to the terminal object.

Example 19.2.9 (pushout of the initial object, 0). Recall that if a category has a terminal object and pullbacks, then it has products. This is the dual for "initial and pushout imply coproducts".

$$\begin{array}{ccc}
 0 & \xrightarrow{!} & B \\
 \downarrow ! & & \downarrow i_B \\
 C & \xrightarrow{i_C} & B \sqcup_0 C
 \end{array}
 \cong
 \begin{array}{ccc}
 0 & \xrightarrow{!} & B \\
 \downarrow ! & & \downarrow i_B \\
 C & \xrightarrow{i_C} & B \sqcup C
 \end{array}$$

The empty set is the initial object in **Set**. There are no b, c such that $f(a) = b, g(a) = c$ because there are no $a \in A \cong \emptyset$. Then each b and c are distinct, resulting in the disjoint union of B and C .

Going back to our product example, suppose B is empty and C is nonempty. The product $B \times C$ is isomorphic to 0 , so we would get $B \sqcup C = \emptyset \sqcup C = C$. If B and C are empty, then naturally the pushout would be empty as well.

Let's look at how pushouts behave with presheaves. Let $F, G \in [C^{op}; Set]$.

Example 19.2.10 (pushout of $F \times G$). This example is interesting because it does not automatically collapse to $\{*\}$. Functors are defined by how they act on objects, so for some $A \in C$ we have

$$\begin{array}{ccc}
 (F \times G)(A) & \xrightarrow{\pi_{1A}} & F(A) \\
 \pi_{2A} \downarrow & & \downarrow i_B \\
 G(A) & \xrightarrow{i_C} & F(A) \sqcup_{(F \times G)(A)} G(A)
 \end{array}$$

Where the projection maps are now natural transformations, also defined component-wise. Recall $(F \times G)(A) \cong F(A) \times G(A)$. We end up with 3 (really 4) cases for the pushout:

1. $F(A)$ and $G(A)$ are nonempty. Then we have the typical product - pushout, so in this case we do get 1 or $\{*\}$.
2. $F(A)$ is nonempty and $G(A)$ is empty. Then $F(A) \times G(A) = \emptyset$. Pushout from the initial object gives us $F(A) \sqcup G(A) = F(A) \sqcup \emptyset = F(A)$. We have a symmetric case for $F(A)$ empty and $G(A)$ nonempty.
3. $F(A)$ and $G(A)$ are empty. This yields $\emptyset \sqcup \emptyset = \emptyset$

Sheaves are a little more complicated. If we compute a pushout of a product of sheaves S, P as above, we can define the result point-wise as a presheaf, but we are not guaranteed that this presheaf satisfies the sheaf conditions.

The general solution to this is to *sheafify* the computed presheaf pushout, which is currently out of the scope of this chapter.

19.3 Sheaf Semantics for Noninterference (putting it all together)

We are now ready to talk about some of the central ideas of the paper. This section will culminate in a proof of a form of noninterference where sheaves are interpreted as types.

Fix a poset \mathcal{P} of security levels closed under finite meets.

Definition 98

- An *abstract behavior* x is a filter on the poset \mathcal{P} . We can interpret this as saying x denotes the security levels where a behavior is permitted.
 - For example, a filter generated by M would include H and \top , capturing that medium-security information/behavior can be used at higher-security levels.
- A *security policy* U is a lower set in \mathcal{P} . The notation $U \vdash x$ is pronounced "U permits x ", and means $U \cap x \neq \emptyset$.
 - Consider the poset $\{L \leq M \leq H\}$. Supposing we have a security policy generated by M : this includes L and M , and can be understood as denoting the security levels *above* which some information/behavior is permitted. Alternatively, the set where this behavior is not allowed.

Definition 99

The authors define P to be the topological space where points are abstract behaviors, and the open sets are of the form $\{x \mid U \vdash x\}$.

Exercise 2

Prove these sets are open with the openness axioms.

See Solution 19.4.2

Each security level $\ell \in \mathcal{P}$ represents a security policy $\langle \ell \rangle$ of the form $\{a \in \mathcal{P} \mid a \leq \ell\}$. The corresponding open subspace of P , $(\{x \mid \langle \ell \rangle \vdash x\})$ is denoted $P_{\langle \ell \rangle}$. The complement, $P \setminus P_{\langle \ell \rangle}$ is denoted $P_{\bullet \langle \ell \rangle}$.

Proposition 19.3.1. The predicate $\langle \ell \rangle \cap x \neq \emptyset$ is equivalent to $\ell \in x$.

To see the above, suppose we have some nonempty intersection for $\langle \ell \rangle$ and x . Then there exists some $k \in \langle \ell \rangle$ such that $k \in x$. Since $k \in \langle \ell \rangle$, this means $k \leq \ell$. Since x is a filter and hence upwards closed, since $\ell \geq k$ and $k \in x$, we have $\ell \in x$. The other direction is immediate since $\ell \in \langle \ell \rangle$ so $\ell \in x$ clearly implies a nonempty intersection.

Proposition 19.3.2. Let $\langle \ell_1 \rangle$ and $\langle \ell_2 \rangle$ be security policies such that $\ell_1 < \ell_2$. Then $P_{\langle \ell_1 \rangle} \subset P_{\langle \ell_2 \rangle}$.

Proof. Let x be an abstract behavior in $P_{\langle \ell_1 \rangle}$, i.e. a filter on \mathcal{P} such that $\langle \ell_1 \rangle \cap x \neq \emptyset$. Then there is some y in this intersection, meaning $y \in x$ and $y \leq \ell_1$. Since $\ell_1 < \ell_2$ and posets have transitivity, $y < \ell_2$, which means $y \in \langle \ell_2 \rangle$. Then $y \in x \cap \langle \ell_2 \rangle$, so $x \in P_{\langle \ell_2 \rangle}$. \square

19.3.1 Sheaf Interpretation

In this subsection, we parse Sterling et al.'s statement: "*our intention is to interpret each type of a dependency core calculus as a sheaf on the space P of abstract behaviors. to see why this interpretation is plausible as a basis for secure information flow, we note that a sheaf on P is the same thing as a presheaf on the poset \mathcal{P}* ".

- In this context, the only relevance of dependency core calculus is that it is a typed programming language where types can be annotated with security levels from some poset.
- This statement assumes a presheaf interpretation over \mathcal{P} is understood to be a plausible basis for modeling information flow, so let's justify this quickly:
 - Consider the presheaf category over \mathcal{P} . This has a natural representation for information flow : for each security level ℓ , we have a set $F(\ell)$ representing what data is visible under security level ℓ . Since presheaves are contravariant functors, for each $k \leq l$ in \mathcal{P} , we get a morphism $F(\ell) \rightarrow F(k)$, demonstrating the restriction of information to that at a lower security level.
- Now we have this note : a sheaf on P is "the same thing" as a presheaf on the poset \mathcal{P} . This sounds like a bijection.

Theorem 19.3.3. There is a bijection between the presheaf category $Psh(\mathcal{P})$ and the sheaf category $Sh(P)$.

Recall the topological space at hand, P , has open sets of the form $\{x \mid U \vdash x\}$, where U is a lower set, x is a filter, and $U \vdash x$ means

$U \cap x \neq \emptyset$. Note that the open sets are defined by the lower sets of \mathcal{P} , so we will denote them by O_U , where U is the lower set of choice. Now we have a series of lemmas.

Lemma 19.3.4. The family $\{P_{\langle \ell \rangle}\}_{\ell \in \mathcal{P}}$ is a basis for the topology on P .

Proof. We know each $P_{\langle \ell \rangle} = \{x \mid \langle \ell \rangle \vdash x\}$ are open by definition.

Let U be an upper set of \mathcal{P} and O_U the corresponding open set of P , so $O_U = \{x \mid x \cap U \neq \emptyset\} = \{x \mid \exists \ell \in U : \ell \in x\}$. We can transform this into

$$\bigcup_{\ell \in U} \{x \mid \ell \in x\} = \bigcup_{\ell \in U} \{x \mid \langle \ell \rangle \vdash x\} = O_U,$$

where the second expression comes from Theorem 19.3.1. So, we have a family of nonempty open sets of the form $P_{\langle \ell \rangle}$ indexed by some subset of \mathcal{P} (as some of the sets in the above union will be naturally empty) such that their union is equal to O_U in P . \square

Lemma 19.3.5. Let $\ell, k \in \mathcal{P}$. We claim $P_{\langle \ell \rangle} \cap P_{\langle k \rangle} = P_{\langle \ell \wedge k \rangle}$

Proof. Note we have $\ell \wedge k$ for any $\ell, k \in \mathcal{P}$ because \mathcal{P} is closed under finite meets. Suppose we have some $x \in P_{\langle \ell \wedge k \rangle}$, meaning $\ell \wedge k \in x$. Since x is upwards closed and $\ell \wedge k$ is the meet of ℓ and k , it follows immediately that $\ell \in x$ and $k \in x$, meaning $x \in P_{\langle \ell \rangle} \cap P_{\langle k \rangle}$.

For the other direction, we begin with $\ell \in x$ and $k \in x$. Since x is a filter, any finite subset of x has a lower bound in x . Considering $\{\ell, k\}$ as our subset, suppose we have some lower bound $q \in x$. By definition of meets, $q \leq \ell \wedge k$, so $\ell \wedge k \in x$, meaning $x \in P_{\langle \ell \wedge k \rangle}$. This reasoning also tells us that filters on a poset are closed under finite meets. \square

Lemma 19.3.6. Any sheaf on P is determined by values on the basis $\{P_{\langle \ell \rangle}\}_{\ell \in \mathcal{P}}$.

Proof. Fix some sheaf F on P and let O_U be open in P . There exists some subset $\mathcal{Q} \subset \mathcal{P}$ such that $O_U = \bigcup_{\ell \in \mathcal{Q}} P_{\langle \ell \rangle}$, so $\{P_{\langle \ell \rangle}\}_{\ell \in \mathcal{Q}}$ is a covering of O_U . Since F is a sheaf, it must satisfy

$$F(O_U) \xrightarrow{e} \prod_{\ell \in \mathcal{Q}} F(P_{\langle \ell \rangle}) \xrightarrow[\prod_{\ell, k \in \mathcal{Q}}]{p} \prod_{\ell, k \in \mathcal{Q}} F(P_{\langle \ell \rangle} \cap P_{\langle k \rangle}),$$

which is equal to

$$F(O_U) \xrightarrow{e} \prod_{\ell \in \mathcal{Q}} F(P_{\langle \ell \rangle}) \xrightarrow[\prod_{\ell, k \in \mathcal{Q}}]{p} \prod_{\ell, k \in \mathcal{Q}} F(P_{\langle \ell \wedge k \rangle}).$$

Then to compute behavior of F on O_U such that the sheaf conditions are satisfied, we only need to know how F acts on the basis, since we have rewritten the rest of the diagram in terms of the basis elements. \square

There is an alternate (perhaps more elegant) theorem in MacLane and Moerdijk demonstrating the same property as Theorem 19.3.6: Let \mathcal{B} be a basis for a topology on a space X . The restriction functor $Sh(X) \rightarrow Sh(\mathcal{B})$ is an equivalence of categories. They leave the proof as an exercise for the reader.

Okay, now we are ready to prove Theorem 19.3.3!

Proof. .

1. $Psh(\mathcal{P}) \rightarrow Sh(P)$: Let E be a presheaf on \mathcal{P} . We will use E to define a sheaf F on P . We have the natural behavior for the basis elements: $F(P_{\langle \ell \rangle}) := E(\ell)$.

As we showed above, all the other open sets follow immediately when we set $F(O_U)$ as the equalizer of the diagram with the covering of choice being the basis elements. Note that this extends to arbitrary coverings, as each open set can be written itself as a union of basis elements, so no matter what we can always reduce a covering to its basis elements.

F maintains all the presheaf requirements as well: suppose $P_{\langle \ell \rangle} \subset P_{\langle k \rangle}$. We've shown this implies $\ell \leq k$. Then we have a restriction morphism $E(k) \rightarrow E(\ell)$, and by definition this yields a restriction morphism $F(P_{\langle k \rangle}) \rightarrow F(P_{\langle \ell \rangle})$. Defining $F(O_U)$ to be the equalizer of the given diagram gives us the desired sheaf conditions, as sketched in Theorem 19.2.3.

2. $Sh(P) \rightarrow Psh(\mathcal{P})$: Set $E(\ell)$ to $F(P_{\langle \ell \rangle})$ for all $\ell \in \mathcal{P}$.

We could also be done after this definition, since we have a unique definition of F on the basis, and the fancier theorem says $F(\mathcal{B})$ is equivalent to $F(P)$, where \mathcal{B} is a basis for P .

□

19.3.2 Transparency and Sealing Monads

These transparency and sealing monads are precisely instances of $i_* \circ i^*$, as defined in Section 19.2.3.

Definition 100

1. The *transparency monad* $\langle \ell \rangle \Rightarrow A$ replaces A with whatever part of it can be viewed under policy $\langle \ell \rangle$. If we define $i : P_{\langle \ell \rangle} \hookrightarrow P$, the composition $i_* \circ i^*$ defines the unit $\eta_{\Rightarrow} : Sh(P) \rightarrow Sh(P)$ where $A \mapsto (\langle \ell \rangle \Rightarrow A)$. When the unit is an isomorphism at A , we say that A is $\langle \ell \rangle$ -transparent.
2. The *sealing monad* $\langle \ell \rangle \bullet A$ removes from A whatever part of it can be viewed under policy $\langle \ell \rangle$. If we define $i : P \setminus P_{\langle \ell \rangle} \hookrightarrow P$, the composition $i_* \circ i^*$ defines the unit $\eta_{\bullet} : Sh(P) \rightarrow Sh(P)$ where $A \mapsto (\langle \ell \rangle \bullet A)$.
 - The sealing monad can be constructed as the pushout $P_{\langle \ell \rangle} \sqcup_{P_{\langle \ell \rangle} \times A} A$. When the unit is an isomorphism at A , we say A is $\langle \ell \rangle$ -sealed.

Proposition 19.3.7. Sterling et al. claim the transparency monad "is" the function space $A^{P_{\langle \ell \rangle}}$, i.e. $A^{P_{\langle \ell \rangle}} \cong \langle \ell \rangle \Rightarrow A$. We will sketch a proof for the general case of open Y in P .

Proof. (sketch)

Let $j : Y \hookrightarrow P$ and A be a sheaf on P . We want $j_* j^* A \cong A^Y$. One definition of the exponential object on sheaves over $\mathcal{O}(X)$ is

$$A^Y(U) \cong \text{Hom}(Y|_U, A|_U)$$

where $Y|_U, A|_U$ are the sheaves restricted to subsets of U . Since $Y|_U(V) = 1$ if $V \subseteq U \cap Y$ and \emptyset otherwise, any morphisms (natural transformations) $Y|_U \rightarrow A|_U$ are characterized by $V \subseteq U \cap Y$.

Let α be a natural transformation and α_V the component associated to $V \subseteq U \cap Y$, i.e. $\alpha_V : Y|_U(V) \rightarrow A|_U(V) = A(V)$. This simplifies to $\alpha_V : 1 \rightarrow A(V)$. There is a natural isomorphisms between $\text{Hom}(1, A)$ and A , for any object A , so α_V is naturally isomorphic to $A(V)$. The subsets V form a cover of $U \cap Y$, so we should be able to glue the nontrivial components of α (i.e., $\alpha|_{U \cap Y}$) to get $A(U \cap Y)$. Then we have

$$\begin{aligned} \text{Hom}(Y|_U, A|_U) &\cong \text{Hom}(Y|_{U \cap Y}, A|_{U \cap Y}) \\ &\cong \{A(V) \mid V \subseteq U \cap Y\} \\ &\cong A(U \cap Y) \end{aligned}$$

You may notice we are treating the open space $P_{\langle \ell \rangle}$ (and general Y) as a sheaf. Recall Theorem 19.2.6, which proves Sterling et al.'s statement that an open set of P is "the same" as a subterminal sheaf.

In Section 19.2.3, we demonstrated $j_*j^*A(U) = A(U \cap Y)$ for $j : Y \hookrightarrow P$, so we have $j_*j^*A \cong A^Y$. \square

Given space P and arbitrary $\langle \ell \rangle \in P$, the subspace defining the transparency monad $\langle \ell \rangle \Rightarrow$ is $P_{\langle \ell \rangle}$, and the subspace defining the sealing monad $\langle \ell \rangle \bullet$ is $P_{\bullet \langle \ell \rangle}$.

Closed sets are defined to be complements of open sets. It is rare that the complement of an open set will also be open (called clopen sets). Note that $P_{\bullet \langle \ell \rangle}$ is not guaranteed to be open – this is where the caveats of Definition 97 with respect to the Alexandroff topology come in. That is, for every open $U \in P_{\bullet \langle \ell \rangle}$, we have $i^*F(U) = F(\uparrow U)$ where $\uparrow U$ is the upward closure of U in P .

Example 19.3.8. Let's fix $\mathcal{P} := \{L \subset M \subset H \subset \top\}$.

There are only four filters on \mathcal{P} , so we can write all of them out:

$$\begin{aligned} p_0 &= \{L, M, H, \top\} \\ p_1 &= \{M, H, \top\} \\ p_2 &= \{H, \top\} \\ p_3 &= \{\top\} \end{aligned}$$

From here, we can also enumerate the principal open sets on P :

$$\begin{aligned} P_{\langle L \rangle} &= \{p_0\} \\ P_{\langle M \rangle} &= \{p_0, p_1\} \\ P_{\langle H \rangle} &= \{p_0, p_1, p_2\} \\ P_{\langle \top \rangle} &= \{p_0, p_1, p_2, p_3\} \end{aligned}$$

To get a better idea of how these monads work, we will work through explicitly what $\langle M \rangle \Rightarrow A$ and $\langle M \rangle \bullet A$ look like.

1. $\langle M \rangle \Rightarrow A$ is the monad defined by composition of the inverse and direct images of $i : P_{\langle M \rangle} \hookrightarrow P$

$$\begin{aligned} (i_* \circ i^*)A(P_{\langle L \rangle}) &= A(P_{\langle L \rangle} \cap P_{\langle M \rangle}) = A(P_{\langle L \rangle}) \\ (i_* \circ i^*)A(P_{\langle M \rangle}) &= A(P_{\langle M \rangle} \cap P_{\langle M \rangle}) = A(P_{\langle M \rangle}) \\ (i_* \circ i^*)A(P_{\langle H \rangle}) &= A(P_{\langle H \rangle} \cap P_{\langle M \rangle}) = A(P_{\langle M \rangle}) \\ (i_* \circ i^*)A(P_{\langle \top \rangle}) &= A(P_{\langle \top \rangle} \cap P_{\langle M \rangle}) = A(P_{\langle M \rangle}) \end{aligned}$$

2. $\langle M \rangle \bullet A$ is the monad defined by composition of the inverse and direct images of $i : P_{\bullet \langle M \rangle} \hookrightarrow P$

$$\begin{aligned} (i_* \circ i^*)A(P_{\langle L \rangle}) &= A(\uparrow (P_{\langle L \rangle} \cap P_{\bullet \langle \ell \rangle})) = A(\emptyset) = 1 \\ (i_* \circ i^*)A(P_{\langle M \rangle}) &= A(\uparrow (P_{\langle M \rangle} \cap P_{\bullet \langle \ell \rangle})) = A(\emptyset) = 1 \\ (i_* \circ i^*)A(P_{\langle H \rangle}) &= A(\uparrow (P_{\langle H \rangle} \cap P_{\bullet \langle \ell \rangle})) = A(\uparrow (\{p_2\})) = A(P_{\langle H \rangle}) \\ (i_* \circ i^*)A(P_{\langle \top \rangle}) &= A(\uparrow (P_{\langle \top \rangle} \cap P_{\bullet \langle \ell \rangle})) = A(\uparrow (\{p_2, p_3\})) = A(P_{\langle \top \rangle}) \end{aligned}$$

This aligns with intuition that sealing at level $\langle \ell \rangle$ preserves all data "higher" than ℓ , but collapses/makes all data at or "lower" than ℓ indistinguishable.

Exercise 3

The authors mention that the sealing monad also has a pushout construction. Check that evaluating $P_{\langle M \rangle} \sqcup_{P_{\langle M \rangle} \times A} A$ at the principal sets $(P_{\langle \ell \rangle})$ for $\ell \in \mathcal{P}$ agrees with the composition of the direct and inverse images of the inclusion map.

(Hint : use Theorem 19.2.10; don't worry about sheafification)

See Solution 19.4.3

Proposition 19.3.9. The $\langle \ell \rangle$ -transparent part of an $\langle \ell \rangle$ -sealed sheaf is trivial, i.e. $(\langle \ell \rangle \Rightarrow (\langle \ell \rangle \bullet A)) \cong \{*\}$.

Proof. Let's expand the definition of $\langle \ell \rangle \Rightarrow$ in the above expression:

$$i_*(i^*(\langle \ell \rangle \bullet A))(V) = (i^*(\langle \ell \rangle \bullet A))(i^{-1}V) = (i^*(\langle \ell \rangle \bullet A))(V \cap P_{\langle \ell \rangle}) = (\langle \ell \rangle \bullet A)(V \cap P_{\langle \ell \rangle})$$

We know $V \cap P_{\langle \ell \rangle} \subset P_{\langle \ell \rangle}$, so from our computation above we can see $(\langle \ell \rangle \bullet A)(V \cap P_{\langle \ell \rangle}) = \{*\}$ for any $V \in P$. \square

19.3.3 Noninterference

We've made it!

The noninterference property will be represented by the constant definition. That is, if we have a morphism from higher security to lower security, this morphism satisfies noninterference if it is constant.

Recall that we aim to view these sheaves as types: we can think of $A(U)$ where $U \subset P$ as the data of type A which has some associated security restrictions/permissions. If we say a morphism $A \rightarrow B$ is constant, this means that the behavior of f does not reveal any information about data of type A .

Theorem 19.3.10. Any map $\langle \ell \rangle \bullet A \rightarrow \text{bool}$ is constant.

Proof. The boolean sheaf bool assigns to each open the set of boolean values $\{\text{true}, \text{false}\}$.

Let's see what happens when we apply $\langle \ell \rangle \Rightarrow \text{bool}$. This evaluates to $\text{bool}(P_{\langle \ell \rangle} V) = \{\text{true}, \text{false}\}$ for any open $V \in P$. By function extensionality, $\langle \ell \rangle \Rightarrow \text{bool} = \text{bool}$ for all ℓ , so bool is $\langle \ell \rangle$ -transparent for all ℓ , meaning $\langle \ell \rangle \Rightarrow \text{bool} \cong \text{bool}$.

The unit $\eta_{\Rightarrow} : \text{Id} \rightarrow (\langle \ell \rangle \Rightarrow _)$ is a natural transformation, so the following diagram commutes:

$$\begin{array}{ccc}
 \langle \ell \rangle \bullet A & \xrightarrow{f} & \mathbf{bool} \\
 \eta_{\Rightarrow} \downarrow & & \downarrow \eta_{\Rightarrow} \\
 \langle \ell \rangle \Rightarrow (\langle \ell \rangle \bullet A) & \xrightarrow{g} & \langle \ell \rangle \Rightarrow \mathbf{bool}
 \end{array}
 \cong
 \begin{array}{ccc}
 \langle \ell \rangle \bullet A & \xrightarrow{f} & \mathbf{bool} \\
 \eta_{\Rightarrow} \downarrow & & \downarrow \eta_{\Rightarrow} \\
 \{*\} & \xrightarrow{g} & \mathbf{bool}
 \end{array}$$

The diagram on the left comes from Theorem 19.3.9 and the fact that \mathbf{bool} is $\langle \ell \rangle$ -transparent for all ℓ . This simplifies to

$$\begin{array}{ccc}
 \langle \ell \rangle \bullet A & \xrightarrow{f} & \mathbf{bool} \\
 & \searrow ! & \uparrow g \\
 & & \{*\}
 \end{array}$$

Since f factors through the terminal object, f must be constant. \square

Since \mathbf{bool} is a constant sheaf, it is essentially public. So, the above theorem is stating the mapping to \mathbf{bool} from any type of $\langle \ell \rangle$ sealed data will not reveal anything about this "more secure" data.

Theorem 19.3.11. Morphisms $\langle \ell \rangle \bullet A \rightarrow B$ are in bijective correspondence with morphisms $A \rightarrow B$ that restricts to a constant function under $\langle \ell \rangle$.

Proof. We take "restricts under $\langle \ell \rangle$ " to mean "under application of the $\langle \ell \rangle$ transparency monad". Recall η_{\bullet} refers to the unit for sealing.

1. Begin with $f : \langle \ell \rangle \bullet A \rightarrow B$.

We can use the pushout definition of the sealing monad to get the following diagram:

$$\begin{array}{ccccc}
 P_{\langle \ell \rangle} \times A & \longrightarrow & A & & \\
 \downarrow & & \downarrow \eta_{\bullet} & \searrow f' & \\
 P_{\langle \ell \rangle} & \longrightarrow & \langle \ell \rangle \bullet A & \xrightarrow{f} & B
 \end{array}$$

where f' is defined to be $f \circ \eta_{\bullet}$. Based on our understanding of restriction under $\langle \ell \rangle$, we want to see how $\langle \ell \rangle \Rightarrow f'$, mapping $\langle \ell \rangle \Rightarrow A \rightarrow \langle \ell \rangle \Rightarrow B$ behaves, so really we only care about

$$\begin{array}{ccc}
 & \xrightarrow{f'} & \\
 A & \xrightarrow{\eta_{\bullet}} \langle \ell \rangle \bullet A & \xrightarrow{f} B
 \end{array}$$

We now use the functoriality of $\langle \ell \rangle \Rightarrow$ to have that the following diagram commutes,

$$\begin{array}{c}
 \langle \ell \rangle \Rightarrow A \xrightarrow{\langle \ell \rangle \Rightarrow \eta_\bullet} \langle \ell \rangle \Rightarrow (\langle \ell \rangle \bullet A) \xrightarrow{\langle \ell \rangle \Rightarrow f} \langle \ell \rangle \Rightarrow B \\
 \langle \ell \rangle \Rightarrow A \xrightarrow{\langle \ell \rangle \Rightarrow f'} \langle \ell \rangle \Rightarrow B
 \end{array}$$

which simplifies to

$$\begin{array}{c}
 \langle \ell \rangle \Rightarrow A \xrightarrow{!} \{*\} \xrightarrow{\langle \ell \rangle \Rightarrow f} \langle \ell \rangle \Rightarrow B \\
 \langle \ell \rangle \Rightarrow A \xrightarrow{\langle \ell \rangle \Rightarrow f'} \langle \ell \rangle \Rightarrow B
 \end{array}$$

So, since $\langle \ell \rangle \Rightarrow f'$ factors through the terminal object, it must be constant, meaning the f' defined by $f \circ \eta_\bullet$ restricts to a constant function under $\langle \ell \rangle$.

2. Begin with an $f' : A \rightarrow B$ such that $\langle \ell \rangle \Rightarrow f'$ is constant.

We again want to observe the pushout diagram for $\langle \ell \rangle \bullet A$:

$$\begin{array}{ccc}
 P_{\langle \ell \rangle} \times A & \longrightarrow & A \\
 \downarrow & & \downarrow \eta_\bullet \\
 P_{\langle \ell \rangle} & \longrightarrow & \langle \ell \rangle \bullet A
 \end{array}
 \begin{array}{c}
 \searrow f' \\
 B
 \end{array}$$

Note we do not have a defined arrow $\langle \ell \rangle \bullet A \rightarrow B$. By the universal property of the pushout, if we have morphisms $A \rightarrow B$ and $P_{\langle \ell \rangle} \rightarrow B$, then there exists a unique morphism $\langle \ell \rangle \bullet A \rightarrow B$ such that everything is nice etcetera. To obtain a morphism $P_{\langle \ell \rangle} \rightarrow B$, we use the constant under $\langle \ell \rangle$ assumption for f' .

$\langle \ell \rangle \Rightarrow f'$ constant means for any $x \in \langle \ell \rangle \Rightarrow A$, we have $(\langle \ell \rangle \Rightarrow f')(x) = c$ for some constant $c \in \langle \ell \rangle \Rightarrow B$. Theorem 19.3.7 states that the transparency monad is isomorphic to the function space $A^{P_{\langle \ell \rangle}}$ (with open sets being treated as subterminal sheaves), so this constant c is isomorphic to a map $P_{\langle \ell \rangle} \rightarrow B$, so we can complete our diagram to get

$$\begin{array}{ccc}
 P_{\langle \ell \rangle} \times A & \longrightarrow & A \\
 \downarrow & & \downarrow \eta_\bullet \\
 P_{\langle \ell \rangle} & \longrightarrow & \langle \ell \rangle \bullet A
 \end{array}
 \begin{array}{c}
 \searrow f' \\
 \dashrightarrow !f \\
 \searrow c \\
 B
 \end{array}$$

For each $f' : A \rightarrow B$ restricting to a constant function under $\langle \ell \rangle$, there exists a unique $f : \langle \ell \rangle \bullet A \rightarrow B$.

□

The noninterference intuition for the above theorem is that, if we have a map which relies on sealed (secret,private,etc) data, then when we try to take a "less secure" view, the behavior of this map is indistinguishable.

19.4 Solutions

Solution 19.4.1. Solution to Exercise 1:

Let U be open in Y . Then U is an upper set, meaning if $a \in U$ and $a \leq b$, then $b \in U$. Suppose we have $x \in f^{-1}(U)$ and $y \in X$ such that $x \leq y$. We wish to show $y \in f^{-1}(U)$.

Since $x \in f^{-1}(U)$, we know $f(x) \in U$. By monotonicity, we know $f(x) \leq f(y)$, meaning $f(y) \in U$. Then $y \in f^{-1}(U)$, so $f^{-1}(U)$ is an upper set and thus open.

Solution 19.4.2. Solution to Exercise 2:

1. The empty set, along with being an upper set, is also a lower set (for the same vacuous reason). Then we can write the set $\{x \mid x \cap \emptyset \neq \emptyset\} = \emptyset$, so the empty set fits this characterization. Now we look at $\{x \mid x \cap P \neq \emptyset\}$, but this in fact contains every x because each x is nonempty, so this set is equal to P .
2. Arbitrary unions: we want to write $\bigcup_{i \in I} \{x \mid U_i \vdash x\}$ as $\{x \mid ? \vdash x\}$. What does it mean for some x' to be in this union? That there is some U_j such that $x' \cap U_j \neq \emptyset$, which means there exists some y such that $y \in x'$ and $y \in U_j$. So we can reformulate this as saying x is in the arbitrary union if there exists some $j \in I$ and $y \in x$ such that $y \in U_j$. The existence of some j can be rephrased as requiring there exist some $y \in x$ such that $y \in \bigcup_{i \in I} U_i$. Then this is equivalent to the existence of some $y \in x \cap \bigcup_{i \in I} U_i$, so we can rewrite the arbitrary union as

$$\{x \mid x \cap \bigcup_{i \in I} U_i\}.$$

The arbitrary union of lower sets is also a lower set — the proof is very similar to that for upper sets, you can do it if you are not convinced. So, we have that these sets are closed under arbitrary unions.

3. Finite intersections: Similarly, we can rewrite $O_V \cap O_U$ as $\{x \mid (U \cap V) \vdash x\}$, and lower sets are also closed under finite intersections.

Solution 19.4.3. Solution to Exercise 3:

We will compute $\langle M \rangle \bullet A$ as the pushout of $P_{\langle M \rangle} \times A$. We will use Theorem 19.2.10, which was for presheaves. Because this space $Sh(P)$ is particularly nice, we don't have to do any fancy sheafifying after computing the pointwise values. Recall $P_{\langle M \rangle}$ is being treated as a subterminal sheaf, so since $P_{\langle H \rangle} \not\subseteq P_{\langle M \rangle}$ for example, $P_{\langle M \rangle}$ applied to $P_{\langle H \rangle}$ will be \emptyset . On the other hand, both $P_{\langle M \rangle}$ and $P_{\langle L \rangle}$ subset $P_{\langle M \rangle}$, so the application returns 1.

These are the resulting values of applying the monad:

$$\begin{aligned} \langle \langle M \rangle \bullet A \rangle (P_{\langle L \rangle}) &= 1 \\ \langle \langle M \rangle \bullet A \rangle (P_{\langle M \rangle}) &= 1 \\ \langle \langle M \rangle \bullet A \rangle (P_{\langle H \rangle}) &= A(P_{\langle H \rangle}) \\ \langle \langle M \rangle \bullet A \rangle (P_{\langle \top \rangle}) &= A(P_{\langle \top \rangle}) \end{aligned}$$

Solution 19.4.4. Solution to Lemma 19.2.4 (sketch)

Let e be an equalizer map for $A \begin{matrix} \xrightarrow{p} \\ \xrightarrow{q} \end{matrix} B$. We want to show e is a monomorphism, meaning for any $f : C \rightarrow A, g : C \rightarrow A$, if we have $e \circ f = e \circ g$, this implies $f = g$. So, suppose $e \circ f = e \circ g$.

By equalizerness, we know $p \circ e = q \circ e$, so we can set $p \circ e \circ f = q \circ e \circ f$. Then $e \circ f$ is an equalizer of p and q . By the universal property, there exists a unique u such that $e \circ u = e \circ f$, and uniqueness of u gives us $u = f = g$, so we have $f = g$.

The coequalizer \implies epimorphism reasoning is very similar.

Type Refinements as Indexed Inductive Types

In the course of writing code in a simply typed language, one may find oneself wishing to constrain the values of a type based on some “feature” of its values. For instance, when writing functions over lists or trees, it may be desirable to enforce constraints on the length or height of the input, or to ensure that the values it stores conform to some precondition. A classic example is the list zip operation: the author of this function may wish to enforce that its arguments have the same length.

Refinement types enrich simple type systems to allow programmers to express such constraints. For example, the type of lists of length n can be written

$$\{xs : \text{List } A \mid \text{length}(xs) = n\}$$

In dependently typed languages, such refinement types can be naïvely translated as dependent pairs, attaching to lists proofs that they satisfy the length constraint:

$$\sum_{xs : \text{List } A} \text{length}(xs) = n$$

However, dependent types permit a more elegant encoding: such constraints can be directly “baked into” the definition of an inductive type via type indices, ensuring that all values inhabiting the type are correct by construction. Our example above, for instance, yields the type of length-indexed vectors:¹

```

1 inductive Vector (A : Type) : Nat -> Type
2   | vnil  : Vector A 0
3   | vcons : {n : Nat} -> A -> Vector A n -> Vector A (succ n)

```

The type `Vector` naturally internalizes the length refinement from the simply typed setting. In this chapter, we will use a categorical construction to generalize this transformation to arbitrary inductive types and a broad class of refinements. In other words, we answer the following question: how can a refinement of a type in a simply

This chapter was written by Joseph Rotella.

¹ Throughout this chapter, we will write code snippets in a syntax similar to that of Lean 4, though we assume all constructors live in the global namespace. Those familiar with Agda can read `inductive` as `data` and `Type` as `Set`.

typed language be translated into an indexed inductive type in a dependently typed one?

Although this question serves as our motivation, the actual aim of this chapter is threefold:

- To give an account of a categorical semantics of inductive and indexed inductive types;
- To illustrate an application of these tools in the form of the translation discussed above; and
- To provide motivating intuition for the study of *fibrations*, which we regrettably do not have space to treat herein.

The content of this chapter is largely based on the work of Atkey, Johann, and Ghani [2]. The content of the first section on inductive types also draws on Wadler’s “Recursive Types for Free!” [36] and scribed lecture notes from New’s presentation of similar material [28].

20.1 Categorifying Inductive Types

We will begin by adding a new programming-language feature to our categorical repertoire: definitions of inductive datatypes.

20.1.1 Wherefore More Constructions?

Let us first consider why the facilities afforded to us by our existing semantic development are inadequate to express inductive type definitions in full generality. This may not be immediately obvious: for instance, as we have previously seen, the datatype of booleans

```
1 inductive Bool : Type where
2   | true  : Bool
3   | false : Bool
```

has elements isomorphic to those of $\text{Unit} \oplus \text{Unit}$, which we interpret categorically as $1 \oplus 1$. We can even model finite parameterized datatypes: the option datatype

```
1 inductive Option (A : Type) : Type where
2   | some : A -> Option A
3   | none : Option A
```

has elements isomorphic to those of $A \oplus \text{Unit}$, which we interpret as $\llbracket A \rrbracket \oplus 1$ (assuming that we are able to interpret A). More generally, to interpret the elements of a given inductive type, we can treat each constructor argument as a term in a product, then sum together the representations of each constructor. Where lies the problem?

You may notice that our wording above—interpreting “elements of the type” induced by the definition, rather than the “definition”

itself—is somewhat roundabout: this is intentional. An inductive type is characterized not merely by its ground inhabitants, but also by the *maps into* the type (i.e., its constructor forms). An inductive type also comes equipped with a *recursor* (also known as an *iterator* or *fold*) that admits recursive definitions over the type. The account above, which describes only the object denoted by the type itself, does not directly characterize these facets of an inductive definition.

Moreover—and more problematically—we also currently lack the facilities to provide a general categorical account of infinite inductive types, such as the type of natural numbers

```
1 inductive Nat : Type where
2   | zero : Nat
3   | succ : Nat -> Nat
```

or the type of lists

```
1 inductive List (A : Type) : Type where
2   | nil : List A
3   | cons : A -> List A -> List A
```

since the tools we have previously used to give the semantics of lambda calculi—(finite) (co)products, exponential objects, and initial and terminal objects—do not admit the construction of sums and products of unbounded arity. To give a complete account of general inductive definitions, we will need to develop some new (and repurpose some old) tools.

We will motivate our translation of inductive datatypes using our standard approach to translating PL concepts into category theory: translating inference rules. Specifically, we will translate the inference rules for booleans and natural numbers into categorical constructions, then observe some patterns that will allow us to generalize these translations to arbitrary inductive types.

20.1.2 Categorifying Booleans

To begin, we return to the example of the type `Bool`. Its constructors are characterized by the following typing rules:

$$\frac{}{\vdash \text{true} : \text{Bool}} \text{T-TRUE} \qquad \frac{}{\vdash \text{false} : \text{Bool}} \text{T-FALSE}$$

Note that we write these rules with an empty context, rather than an arbitrary Γ . Recalling that the empty context is represented by the terminal object 1 , these rules correspond to the existence of two *global elements* $\text{true}, \text{false} : 1 \rightarrow \llbracket \text{Bool} \rrbracket$.

Definition 101: Global element

In a category with terminal objects, a *global element* of an object A is a morphism $1 \rightarrow A$. It is the categorical interpretation of an “element” of an object A .

Here, our use of global elements corresponds to the fact that `true` and `false` are distinguished members of the “collection” denoted by `Bool`. The rules above translate into the following diagram:

$$1 \xrightarrow{\text{true}} \llbracket \text{Bool} \rrbracket \xleftarrow{\text{false}} 1$$

However, we want to be able to map *out* of our inductive types as well. To do so, we must define the recursor for the Boolean type, canonically known as `if`:²

$$\frac{\vdash t : \tau \quad \vdash e : \tau}{b : \text{Bool} \vdash \text{if}(t, e)(b) : \tau} \text{T-IF}$$

The recursor is operationally characterized by β and η laws. The β laws say that `if` should “case correctly”: it should evaluate to its “then” branch on a `true` input, and to its “else” branch on a `false` one. The η law, meanwhile, says that “things that behave like `ifs` are judgmentally equal to `ifs`,” one can think of this as a restricted form of function extensionality:


$$\frac{\vdash t : \tau \quad \vdash e : \tau}{\vdash \text{if}(t, e)(\text{true}) \equiv t : \tau} \beta_1^{\text{if}} \quad \frac{\vdash t : \tau \quad \vdash e : \tau}{\vdash \text{if}(t, e)(\text{false}) \equiv e : \tau} \beta_2^{\text{if}}$$

$$\frac{x : \text{Bool} \vdash d : \tau \quad \vdash d[\text{true}/x] \equiv t : \tau \quad \vdash d[\text{false}/x] \equiv e : \tau}{x : \text{Bool} \vdash d \equiv \text{if}(t, e)(x) : \tau} \eta^{\text{if}}$$

As in our previous translations of equational laws, the β laws correspond to equalities of morphisms, while the η law is a uniqueness criterion. In particular, the β laws require that if we have morphisms $t, e : 1 \rightarrow \llbracket \tau \rrbracket$, then $\text{if}(t, e) \circ \text{true} = t$ and $\text{if}(t, e) \circ \text{false} = e$. Meanwhile, the η law says that if we have some expression d such that it can be substituted for $\text{if}(t, e)$ in the preceding two equalities³—i.e., $d \circ \text{true} \equiv t$ and $d \circ \text{false} \equiv e$ —then it is equal to the recursor: in other words, the recursor denotes the unique morphism satisfying these equalities. Putting these together, we obtain the following pair of commutative triangles:

² We will adopt the convention of writing recursors with their *major premise* (i.e., the value being “recursed on”) as the final argument. This is opposite the usual notation in functional programming but will better agree with the categorical formulation.

³ Recall that we interpret substitution as composition of morphisms.



For reasons that will become clear shortly, it is helpful to expand these triangles into squares:

$$\begin{array}{ccccc}
 1 & \xrightarrow{\text{true}} & \llbracket \text{Bool} \rrbracket & \xleftarrow{\text{false}} & 1 \\
 \downarrow \text{id}_1 & & \downarrow \text{if}(t,e) & & \downarrow \text{id}_1 \\
 1 & \xrightarrow{t} & \llbracket \tau \rrbracket & \xleftarrow{e} & 1
 \end{array} \quad (20.1)$$

Exercise 20.1

Verify that $1 \oplus 1$ is a suitable choice for $\llbracket \text{Bool} \rrbracket$. What are the denotations of `true`, `false`, and `if(t,e)`?

20.1.3 Categorifying Natural Numbers

We now perform a similar translation for natural numbers. The type `Nat` has two constructors, `zero` and `succ`, and is equipped with a recursor $\text{rec}_{\text{Nat}}(e_z, x.e_s)(n)$. Notice that, unlike `if`, the recursor for `Nat` binds a variable in the successor case: this will be the recursive result of evaluating the recursor on the substructure of n . The typing rules are given below:

$$\frac{}{\vdash \text{zero} : \text{Nat}} \text{T-ZERO} \qquad \frac{}{n : \text{Nat} \vdash \text{succ}(n) : \text{Nat}} \text{T-SUCC}$$

$$\frac{\vdash e_z : \tau \quad x : \text{Nat} \vdash e_s : \tau}{n : \text{Nat} \vdash \text{rec}_{\text{Nat}}(e_z, x.e_s)(n) : \tau} \text{T-NATREC}$$

The β rules for the recursor express its behavior as a fold over the structure of a natural number: applying the recursor to zero evaluates to the first arm, while applying it to a successor binds x to the recursive result of folding over its predecessor, then evaluates e_s :

$$\frac{\vdash e_z : \tau \quad x : \text{Nat} \vdash e_s : \tau}{\vdash \text{rec}_{\text{Nat}}(e_z, x.e_s)(\text{zero}) \equiv e_z : \tau} \beta_z^{\text{rec}_{\text{Nat}}}$$

$$\frac{\vdash e_z : \tau \quad x : \text{Nat} \vdash e_s : \tau \quad \vdash n : \text{Nat}}{\vdash \text{rec}_{\text{Nat}}(e_z, x.e_s)(\text{succ}(n)) \equiv e_s[\text{rec}_{\text{Nat}}(e_z, x.e_s)(n)/x] : \tau} \beta_s^{\text{rec}_{\text{Nat}}}$$

Finally, the η rule says that any expression e that satisfies the above equations is equal to the corresponding recursor application:

$$\frac{y : \text{Nat} \vdash e : \tau \quad \vdash e[\text{zero}/y] \equiv e_z : \tau \quad n : \text{Nat} \vdash e[\text{succ}(n)/y] \equiv e_s[e/x] : \tau}{y : \text{Nat} \vdash e \equiv \text{rec}_{\text{Nat}}(e_z, x.e_s)(y) : \tau} \eta^{\text{rec}_{\text{Nat}}}$$

As with Booleans, categorifying the rules for natural numbers yields an object representing the type Nat , a morphism for each constructor, and, for each pair of “combining-function” morphisms running “parallel to” the constructors, a unique morphism denoted by the recursor.⁴ In the case of natural numbers, however, the resulting construction has a special name: the *natural numbers object*.


⁴ Note that we obtain the uniqueness condition from the η law by recalling that substitution is interpreted as composition. Like the corresponding condition for Bool , this uniqueness condition says that any morphism e satisfying the β -equalities $e \circ \text{zero} = e_z$ and $e \circ \text{succ} = e_s$ is equal to the corresponding recursor morphism.

Definition 102: Natural numbers object

A *natural numbers object* in a Cartesian closed category \mathcal{C} consists of

- An object N of \mathcal{C} , and
- Morphisms $z : 1 \rightarrow N$ and $s : N \rightarrow N$,

such that, for every object A of \mathcal{C} and pair of morphisms $f_z : 1 \rightarrow A$ and $f_s : A \rightarrow A$, there exists a unique morphism u such that the following diagram commutes:



Here, u is $\text{rec}_{\text{Nat}}(e_z, x.e_s)$. As before, it will be helpful to instead view this diagram as a pair of squares:

$$\begin{array}{ccccc} 1 & \xrightarrow{\text{zero}} & \llbracket \text{Nat} \rrbracket & \xleftarrow{\text{succ}} & \llbracket \text{Nat} \rrbracket \\ \downarrow \text{id}_1 & & \downarrow \text{rec}_{\text{Nat}}(e_z, x.e_s) & & \downarrow \text{rec}_{\text{Nat}}(e_z, x.e_s) \\ 1 & \xrightarrow{e_z} & \llbracket \tau \rrbracket & \xleftarrow{e_s} & \llbracket \tau \rrbracket \end{array} \quad (20.2)$$

Exercise 20.2

Give a categorical interpretation of the type of *extended natural numbers*, i.e., the type of natural numbers extended with an extra

element representing infinity, defined as follows:

```

1 inductive ExtNat : Type where
2   | zero : ExtNat
3   | succ : ExtNat -> ExtNat
4   | inf  : ExtNat
  
```

Hint: The resulting diagram will look a bit different from the two preceding examples.

20.1.4 A General Categorification of Inductive Types

Let us now take stock of the preceding examples and generalize what we have seen. In each of our translations above, we did the following:


- First, interpret the constructors of an inductive type as morphisms from the interpretation of the constructor’s domain to the type’s denotation. That is, a constructor $c : D \rightarrow T$ of T becomes a morphism $\llbracket D \rrbracket \xrightarrow{c} \llbracket T \rrbracket$.
- Then, interpret the recursor as a unique morphism $\llbracket T \rrbracket \rightarrow A$ for every collection of morphisms into A “parallel to the constructors,” i.e., morphisms t_i that map into A from the domain of the i th constructor where $\llbracket T \rrbracket$ has been replaced by A . For instance, if T has two constructors

$$c_1 : \text{Bool} \rightarrow T \quad \text{and} \quad c_2 : T \times \text{Nat} \rightarrow T,$$

its recursor is interpreted as a unique morphism $u_{f_1, f_2} : \llbracket T \rrbracket \rightarrow A$ for all objects A and morphisms

$$f_1 : \llbracket \text{Bool} \rrbracket \rightarrow A \quad \text{and} \quad f_2 : A \times \llbracket \text{Nat} \rrbracket \rightarrow A.$$

so that the following squares commute:⁵



⁵ Given morphisms $f : X \rightarrow Z$ and $g : Y \rightarrow W$, we write $f \times g$ to denote the morphism $\langle f \circ \pi_1, g \circ \pi_2 \rangle : X \times Y \rightarrow Z \times W$.

This procedure is a useful heuristic but is not yet a precise categorical account of an inductive type. For instance, the “replacement” performed in the second step is merely a syntactic operation. We will now make this categorical translation more precise.

We start by consolidating our commutative squares. Observe that any inductive type with multiple constructors can be equivalently expressed as having a single constructor whose domain is the sum of the domains of the original constructors. That is, the constructors

$c_1 : A \rightarrow T$, $c_2 : B \rightarrow T$, and $c_3 : C \rightarrow T$ can be equivalently represented by a single constructor $c : A \oplus B \oplus C \rightarrow T$. The recursor is then parameterized by a single morphism running “parallel” to the constructor that maps into A from an analogous coproduct over A .

Applying this rewriting to the `Nat` datatype depicted in Diagram (20.2), we obtain

$$\begin{array}{ccc}
 1 \oplus \llbracket \text{Nat} \rrbracket & \xrightarrow{[z,s]} & \llbracket \text{Nat} \rrbracket \\
 \text{id}_1 \oplus \text{rec}_{\text{Nat}}([f_z, f_s]) \downarrow & & \downarrow \text{rec}_{\text{Nat}}([f_z, f_s]) \\
 1 \oplus A & \xrightarrow{[f_z, f_s]} & A
 \end{array}$$

and a similar transformation occurs for Diagram (20.1).⁶

The top-left corner should look familiar: it is the body of the recursive type equation that defines the natural numbers. That is, using a least-fixed-point operator μ , we would write

$$\text{Nat} \triangleq \mu t. 1 \oplus t,$$

so that `Nat` is the (smallest) type satisfying the type isomorphism

$$\text{Nat} \cong 1 \oplus \text{Nat}.$$

Similarly, Diagram (20.1) is “consolidated” into a single square with a constructor with domain $1 \oplus 1$:

$$\begin{array}{ccc}
 1 \oplus 1 & \xrightarrow{[z,s]} & \llbracket \text{Bool} \rrbracket \\
 \text{id}_1 \oplus \text{id}_1 \downarrow & & \downarrow \text{if}([t,e]) \\
 1 \oplus 1 & \xrightarrow{[t,e]} & A
 \end{array}$$

With this consolidation, every inductive type now translates to a single object (the top-left corner of the square), a single morphism (the constructor), and a universal property that translates the recursor. The “shape” of the objects on the left-hand side of the commutative square are determined by the “shape” of the inductive type. More precisely, using the observation in the preceding paragraph, that “shape” is determined precisely by the expression τ in the fixed-point equation $\mu t. \tau$ for the inductive type.

Adopting a categorical lens, we might observe that an expression $t. \tau$ defines a mapping from a type t to some type τ parameterized thereby. In our categorical interpretation, types form a category; thus, such a mapping from types to types ought to correspond to a functor. More specifically, since we are mapping from a category (of types)

⁶ Above and hereafter, we use the notation $[f, g]$ to denote the unique morphism obtained from the universal property of coproducts. Thinking of morphisms as functions, $[f, g]$ does a case split on its input, mapping elements of the left-hand summand using f and those of the right-hand summand using g . By analogy with $f \times g$, we write $f \oplus g$ for $[\text{inl} \circ f, \text{inr} \circ g]$.

into itself, this is an *endofunctor*. The endofunctor associated with the type of natural numbers is defined by

$$F_{\text{Nat}}(X) = 1 \oplus X$$

$$F_{\text{Nat}}(f : X \rightarrow Y) = \text{id}_1 \oplus f : 1 \oplus X \rightarrow 1 \oplus Y$$

Of course, the mapping $t. \tau$ is not itself a type; rather, we apply the least-fixed-point operator μ to obtain an inductive type. We should expect a corresponding maneuver in the world of categories: we interpret an inductive type as the *least fixed point* μF of an endofunctor F on our semantic domain: i.e., the “least” object μF satisfying $F(\mu F) = \mu F$.

This naturally invites the following question: in the absence of an *a priori* ordering on objects, how do we make precise the fact that μF should be “least” among the fixed points of F ? Appealing to our categorical intuition, we should expect this to be some kind of initiality or colimiting property (recall that, using our order-theoretic lens, we think of initial objects as “minimal”). And, in fact, as we will now demonstrate, we have already seen such a property in our Boolean and natural-number examples!

The key is to recall that an inductive type is defined not only by an object but also by a morphism representing its constructor that tells us how to “map into” the type. Thus, instead of merely considering the “smallest” object among the fixed points of F , we should instead consider the “smallest” *pair* of an F -fixed point *and* its corresponding constructor morphism. Such a pair is an example of an *algebra* over the endofunctor F , also known as an *F-algebra*.

Definition 103: F -Algebra

Let $F : \mathcal{C} \rightarrow \mathcal{C}$ be an endofunctor. An F -algebra, or *algebra over F* , is a pair (A, α) where A is an object in \mathcal{C} and $\alpha : FA \rightarrow A$ is a morphism.

A is called the *carrier* of the F -algebra, and α is the *structure map*. Note that the structure map α alone is sometimes referred to as an F -algebra, leaving A implicit.

We are looking for the “smallest” F -algebra (A, α) such that its carrier A is a fixed point. Let us first, however, consider simply the “smallest” F -algebra, without any further restrictions. We do this because it turns out that there is a natural interpretation of “smallest” in terms of initiality. In particular, F -algebras form a category Alg_F , so the “smallest” F -algebra is simply the initial object in this category.

Definition 104: Category of F -algebras

Fix an endofunctor $F : \mathcal{C} \rightarrow \mathcal{C}$. The *category of F -algebras* (or *category of algebras over F*) is defined by the following:

- Objects are F -algebras $(A, \alpha : FA \rightarrow A)$.
- Morphisms from (A, α) to (B, β) are morphisms $f : A \rightarrow B$ in \mathcal{C} such that the following square commutes:

$$\begin{array}{ccc} FA & \xrightarrow{Ff} & FB \\ \alpha \downarrow & & \downarrow \beta \\ A & \xrightarrow{f} & B \end{array}$$

- Identity and composition are lifted from \mathcal{C} .

Definition 105: Initial F -algebra

Fix an endofunctor $F : \mathcal{C} \rightarrow \mathcal{C}$. The *initial F -algebra* (or *initial algebra over F*), if it exists, is the initial object in the category of F -algebras.

We write in_F for the (structure map of) the initial F -algebra. Given another F -algebra (A, α) , we write $\langle \alpha \rangle_F$ (or just $\langle \alpha \rangle$ if F is clear from context) for the unique morphism from the initial object to (A, α) .

We have thus found a suitable notion of a “smallest” F -algebra. But what about the smallest F -algebra whose carrier is a fixed point of F ? In fact, these are one and the same (up to isomorphism): it will always be the case that the carrier of the initial F -algebra is (isomorphic to) a fixed point of F . In other words, if (A, α) is the initial F -algebra, then $FA \cong A$. This is the content of *Lambek’s lemma*.

Proposition 20.1.1 (Lambek’s lemma). Let $F : \mathcal{C} \rightarrow \mathcal{C}$ be an endofunctor. If $(A, \text{in}_F : FA \rightarrow A)$ is the initial F -algebra, then in_F is an isomorphism.

Proof. Let $(A, \text{in}_F : FA \rightarrow A)$ be the initial F -algebra. We will show that in_F is an isomorphism by showing that its inverse is $\langle F \text{in}_F \rangle : A \rightarrow FA$, the unique F -algebra morphism $(A, \text{in}_F) \rightarrow (FA, F \text{in}_F)$ induced by initiality.

First, observe that in_F is an F -algebra morphism $(FA, F \text{in}_F) \rightarrow (A, \text{in}_F)$, since the following diagram trivially commutes:

$$\begin{array}{ccc}
 F(FA) & \xrightarrow{F \text{ in}_F} & FA \\
 \downarrow F \text{ in}_F & & \downarrow \text{in}_F \\
 FA & \xrightarrow{\text{in}_F} & A
 \end{array}$$

Thus, we have the following diagram in Alg_F :

$$\begin{array}{ccccc}
 (A, \text{in}_F) & \xrightarrow{\langle F \text{ in}_F \rangle} & (FA, F \text{ in}_F) & \xrightarrow{\text{in}_F} & (A, \text{in}_F) \\
 & \searrow & & \nearrow & \\
 & & \text{id}_A & &
 \end{array} \quad (20.3)$$

Since (A, in_F) is initial, there is only one F -algebra morphism from (A, in_F) to itself, so this diagram must commute, i.e.,

$$\text{in}_F \circ \langle F \text{ in}_F \rangle = \text{id}_A. \quad (20.4)$$

Accordingly, $\langle F \text{ in}_F \rangle$ is a right inverse for in_F .

To show that it is also a left inverse, we first expand the F -algebra morphisms in Diagram (20.3) into commutative squares:

$$\begin{array}{ccccc}
 FA & \xrightarrow{F \langle F \text{ in}_F \rangle} & F(FA) & \xrightarrow{F \text{ in}_F} & FA \\
 \downarrow \text{in}_F & & \downarrow F \text{ in}_F & & \downarrow \text{in}_F \\
 A & \xrightarrow{\langle F \text{ in}_F \rangle} & FA & \xrightarrow{\text{in}_F} & A \\
 & \searrow & & \nearrow & \\
 & & \text{id}_A & &
 \end{array}$$

In order to show that $\langle F \text{ in}_F \rangle$ is a left inverse of in_F , we must show that the highlighted lower-left path through the left square is equal to id_A .


Observe, however, that the upper-right path through the same square is $F \text{ in}_F \circ F \langle F \text{ in}_F \rangle$, which we can rewrite as follows by functoriality:

$$\begin{aligned}
 F \text{ in}_F \circ F \langle F \text{ in}_F \rangle &= F (\text{in}_F \circ \langle F \text{ in}_F \rangle) \\
 &= F \text{id}_A \\
 &= \text{id}_{FA}.
 \end{aligned} \quad (\text{Equation 20.4})$$

Adding this fact to our diagram, we can read off the desired equation

$$\langle F \text{ in}_F \rangle \circ \text{in}_F = \text{id}_{FA}$$

by tracing the lower-left path:




Thus, $\text{in}_F : FA \rightarrow A$ is indeed an isomorphism. □

With Lambek’s lemma in hand, we can now give a precise categorification of inductive types: the inductive type represented by $\mu t. \tau$ is interpreted as the initial algebra of the endofunctor $F(t) = \llbracket \tau \rrbracket$, where we extend our previous translation of types to map the free type variable t to the functor variable of the same name.⁷ The algebra’s carrier μF is the object representing the type, while its structure map in_F represents the type’s constructor.

⁷ The action of F on morphisms must also be defined in the “obvious” way (which can be derived from viewing products and coproducts of interpreted types in \mathcal{C} as categorical constructions in the functor category $[\mathcal{C}; \mathcal{C}]$). We elide the details, but one can inspect the definition of F_{Nat} at the beginning of this subsection for an example.

It remains only to precisify our former account of recursors. Recall our previous observation that recursors map out of an inductive type μF into some other type A using a “combining morphism” that runs “parallel” to the constructor morphism—i.e., one that maps into A from an object that “replaces” instances of μF with instances of A . Using our newfound language of functors, we can phrase this much more concisely: a recursor maps from μF to some type A using a “combining morphism” $f_c : FA \rightarrow A$. Or, as a diagram:



In fact, we already saw this diagram (Diagrams (20.1) and (20.2) are variations of it in the particular cases of natural numbers and Booleans). But having now developed the vocabulary of F -algebras, this should look familiar for another reason: recursors are F -algebra morphisms! More precisely, given an inductive type μF , the “unapplied” recursor function is represented by $\llbracket - \rrbracket_F$; given a particular “combining morphism” $f_c : FA \rightarrow A$ —which is precisely an F -algebra over A —the resulting recursor application is $\llbracket f_c \rrbracket : \mu F \rightarrow A$.

As an example, we translate the parity function on natural numbers, which returns `true` if its input is even and `false` if it is odd, as an application of the natural-number recursor. First, we identify the appropriate “combining function” for this fold. Since we are ulti-

mately producing a Boolean, this function should translate to a morphism of type $F_{\text{Nat}} \llbracket \text{Bool} \rrbracket \rightarrow \llbracket \text{Bool} \rrbracket$. We first define its programmatic version below:⁸

$$\begin{aligned} f_{\text{even}} &: 1 \oplus \text{Bool} \rightarrow \text{Bool} \\ f_{\text{even}}(\text{inl } ()) &= \text{true} \\ f_{\text{even}}(\text{inr } b) &= \text{not}(b) \end{aligned}$$

This becomes the F_{Nat} -algebra

$$\alpha_{\text{even}} = [\text{true}, \text{not}]$$

which finally reveals our translation of the parity function to be $\langle \alpha_{\text{even}} \rangle : \llbracket \text{Nat} \rrbracket \rightarrow \llbracket \text{Bool} \rrbracket$.

Note that `not` is, in turn, defined by an algebra over F_{Bool} , the constant- $(1 \oplus 1)$ functor. In particular, $\llbracket \text{not} \rrbracket = \langle \alpha_{\text{not}} \rangle$ where $\alpha_{\text{not}} : 1 \oplus 1 \rightarrow 1 \oplus 1$ is $[\text{inr}, \text{inl}]$.

To put together the developments in this section, let's consider one more example of a categorified inductive datatype, now expressed in terms of initial algebras. Recall the inductive datatype of lists:

```
1 inductive List (A : Type) : Type where
2   | nil : List A
3   | cons : A -> List A -> List A
```

Notice that this type is parameterized over a fixed type A ; we will write A for that type's categorical denotation. We can read the endofunctor of which `List A` is a least fixed point off of its constructors:⁹

$$F_{\text{List}_A}(X) = 1 \oplus A \times X.$$

Lists are thus represented by the initial algebra $(\mu F_{\text{List}_A}, \text{in}_{F_{\text{List}_A}})$.

Returning to our example from earlier, we can define the list length function as the fold induced by the "combining function"

$$\begin{aligned} f_{\text{length}} &: \text{Unit} \oplus A \times \text{Nat} \rightarrow \text{Nat} \\ f_{\text{length}}(\text{inl } ()) &= 0 \\ f_{\text{length}}(\text{inr } (a, \ell)) &= \text{succ}(\ell) \end{aligned}$$

which we translate into the F_{List_A} -algebra

$$\alpha_{\text{length}} = [z, s \circ \pi_1]$$

to obtain the fold $\langle \alpha_{\text{length}} \rangle : \mu F_{\text{List}_A} \rightarrow \llbracket \text{Nat} \rrbracket$.

Exercise 20.3

Below, we define a strange type of so-called "self-trees:" each node can have as many children as the parameterizing type has elements

⁸ It may be helpful to compare this definition to the following equivalent, structurally recursive definition in a more familiar programmatic style:

```
def even : Nat -> Bool
  | 0 => true
  | n + 1 => not (even n)
```

Similarly, `not` would be defined as:

```
def not : Bool -> Bool
  | true => false
  | false => true
```

⁹ We again elide the action on morphisms.

(i.e., a `SelfTree A` node can have, as it were, “`A` many children”).

```

1 inductive SelfTree (A : Type) : Type where
2   | node : A -> (A -> SelfTree A) -> SelfTree A
3   | leaf : SelfTree A

```

1. Specify the functor F_{SelfTree} that characterizes this type.
2. We define the “self-height” of a self-tree to be the number of nodes encountered by successively visiting the a th child of a node containing a . (The self-height of a leaf is 0.) Categorify this operation using a suitable F_{SelfTree} -algebra.

Exercise 20.4

Show that the morphism `app` induced by the universal property of exponentials is the structure map for an algebra whose carrier is the initial object. Use this to give a non-empty definition of an inductive type (i.e., an inductive declaration listing at least one putative constructor) that is isomorphic to the empty type.

Hint: Recall that in a CCC, $0^A \cong 0 \cong 0 \times A$.

20.1.5 On the Existence of Initial Algebras

Up to this point, we have simply assumed that, given an endofunctor F , we can always find an initial F -algebra. This is not, however, always the case. And we should expect as much: after all, not all recursive type equations admit a least fixed point—this is why we insist that no non-positive occurrences of t appear in τ when defining an inductive type $\mu t. \tau$. The failure to obtain a least fixed point in a recursive type equation corresponds to an analogous lack of an initial algebra for the corresponding functor. For instance, $\mu t. (t \rightarrow 2) \rightarrow 2$ (i.e., $F(X) = 2^{2^X}$) does not admit a least fixed point in every Cartesian closed category: for one, Lambek’s lemma says that this would induce an isomorphism $\mu F \cong 2^{2^{\mu F}}$, which by Cantor’s theorem never holds in the CCC `Set`. In fact, we can more generally prove the non-existence of certain initial algebras in arbitrary categories, not just `Set`, but this would take us too far afield. For the purposes of this chapter, it suffices to know that all the usual strictly positive inductive types of interest are characterized by endofunctors that admit initial algebras.

Exercise 20.5

There is another way to think about interpreting inductive types.

While we gave an interpretation by translating the various features of inductive types' structure into categorical terms, we could have alternatively followed our previous method and directly translated the generalized inference rules for isorecursively formulated inductive types (cf. Harper, *Practical Foundations for Programming Languages* (2nd ed.), Ch. 15).

Give a translation of these rules (shown below), and verify that it yields the same construction at which we just arrived.

$$\frac{\vdash t.\tau \text{ pos}}{e : \tau[\mu t.\tau/t] \vdash \text{fold}(e) : \mu t.\tau} \text{ T-FOLD}$$

$$\frac{x : \tau[\rho/t] \vdash e_1 : \rho}{e_2 : \mu t.\tau \vdash \text{rec}(x.e_1)(e_2) : \rho} \text{ T-REC}$$

$$\frac{x : \tau[\rho/t] \vdash e_1 : \rho}{e_2 : \tau[\mu t.\tau/t] \vdash \text{rec}(x.e_1)(\text{fold}(e_2)) \equiv e_1[\text{map}(e_2)(y.\text{rec}(x.e_1)(y))]/x] : \rho} \beta^{\text{rec}}$$

$$\frac{y : \tau[\rho/t] \vdash e : \rho}{e_2 : \tau[\mu t.\tau/t] \vdash e[\text{fold}(e_2)/y] \equiv e_1[\text{map}(e_2)(y.\text{rec}(x.e_1)(y))]/x]} \eta^{\text{rec}}$$

$$\frac{}{y : \mu t.\tau \vdash e \equiv \text{rec}(x.e_1)(y) : \rho}$$

Note that these rules introduce two new auxiliary definitions:

- The judgment $\vdash t.\tau \text{ pos}$ asserts that t has only *positive* occurrences in τ ; i.e., t never appears to the left of an arrow. Note that all of the rules implicitly carry this hypothesis, though we have only explicitly stated it for the first.
- The auxiliary function $\text{map}(e)(e')$ “unpacks” the structure of e based on its type and applies e' to each “element” thereof. Its typing rule is as follows (where t may appear free in τ):

$$\frac{\vdash e : \tau[\rho/t] \quad x : \rho \vdash e' : \rho'}{\vdash \text{map}(e)(x.e') : \tau[\rho'/t]} \text{ T-MAP}$$

Here are two of its equational rules (it has one for every type

former):

$$\frac{\vdash e : (\tau_1 \times \tau_2)[\rho/t] \quad x : \rho \vdash e' : \rho'}{\vdash \text{map}(e)(e') \equiv \langle \text{map}(\text{fst } e)(x.e'), \text{map}(\text{snd } e)(x.e') \rangle : \tau_1[\rho'/t] \times \tau_2[\rho'/t]} \beta_{\times}^{\text{map}}$$

$$\frac{\vdash e : (\tau_1 \oplus \tau_2)[\rho/t] \quad x : \rho \vdash e' : \rho'}{\vdash \text{map}(e)(e') \equiv \text{case } e \text{ of} \begin{array}{l} | \text{inl } x_1 \mapsto \text{inl } (\text{map}(x_1)(x.e')) \\ | \text{inr } x_2 \mapsto \text{inr } (\text{map}(x_2)(x.e')) \end{array} : \tau_1[\rho'/t] \oplus \tau_2[\rho'/t]} \beta_{\oplus}^{\text{map}}$$

Your translation should give an account of what these definitions mean categorically.

Here are some hints for this exercise:

- If parameterized types are translated as functors, to what does type-variable substitution correspond?
- $\text{map}(e)$ will not be translated as a morphism.

20.2 The Families Fibration: Indexing Inductive Types

Having given a categorical translation of simple inductive types, we now turn our attention to *indexed* inductive types. Recall that indexed inductive types enrich inductive types by allowing types to be parameterized by data values at the type level: for instance, in addition to the simple type $\text{List } A$ of arbitrary-length lists with elements of type A , we can also define the indexed type $\text{Vector } A \ n$ of A -containing lists of length n .

While we gave a generalized presentation of inductive types in terms of initial algebras of endofunctors over arbitrary categories, hereafter we will specifically consider a categorical semantics in terms of sets. While all of the ideas presented can be generalized to other categories, focusing on Set will simplify some of the required categorical machinery. (Note that our preceding interpretation of inductive types can indeed be carried out in Set , since it is Cartesian closed and contains finite coproducts—i.e., it is *bicartesian closed*.)

The category Set alone, however, lacks the structure to give an elegant semantics to indexed inductive types. Instead, we will turn to a new category built atop Set that is equipped with a notion of “indexing” that will afford a straightforward interpretation of our inductive types’ indices.¹⁰

¹⁰ Be careful when reading the following definition and elsewhere in this chapter: the notation $\prod_{a \in A} S(a)$ below refers to a *set-theoretic* indexed product, not a Π -type. The symbols \prod , \sum , and \rightarrow are regrettably overloaded (we will see yet another denotation of \sum later on), so the meanings of these notations are highly context-sensitive.

Definition 106: Families fibration

The *category of families (of sets)* $\text{Fam}(\text{Set})$ is defined as follows:

- Objects are pairs $(A, P : A \rightarrow \text{Ob}(\text{Set}))$ of an indexing set A and a *family* of sets P specifying a set $P(a)$ for each index $a \in A$.
- Morphisms from (A, P) to (B, Q) are pairs of functions (f, f^\sim) such that $f : A \rightarrow B$ and $f^\sim : \prod_{a \in A} P(a) \rightarrow Q(f(a))$.

The forgetful functor $U : \text{Fam}(\text{Set}) \rightarrow \text{Set}$ defined by

$$U(A, P) = A$$

$$U(f, f^\sim) = f$$

is called the *families fibration*.

Note that, while we will not treat them in full generality, (*Grothendieck*) *fibrations* are objects of particular interest in category theory and characterize a certain kind of structure (which we will attempt to intuitively motivate at several points hereafter). Throughout this chapter, you will encounter definitions phrased as features “of the families fibration:” these should be read as referring to the entire structure of $\text{Fam}(\text{Set})$, Set , and this mapping between them, rather than particular characteristics of a single forgetful functor.

Let’s now use the category of families to interpret indexed inductive types. We’ll start with the type $\text{BitVec } n$, which is like the type $\text{Vector } A \ n$ we have previously discussed, except that it contains only Booleans:¹¹

```
1 inductive BitVec : Nat -> Type
2   | bvnul  : BitVec 0
3   | bvcons : {n : Nat} -> Bool -> BitVec n -> BitVec (succ n)
```

Drawing inspiration from our preceding development of simple inductive types, we might expect to interpret this type as an endofunctor on $\text{Fam}(\text{Set})$. That endofunctor should produce a family that looks like a representation of the Boolean list type (i.e, the sum of an “empty” object and a Boolean–vector pair object), except that we can only inject into the left-hand side of the sum at index zero and into the right-hand side at successor indices. Concretely, such a functor looks as follows:¹²

¹¹ $\text{Vector } A \ n$ is defined similarly. We choose BitVec for this example to avoid any potential confusion between the parameter A and the index n of Vector .

¹² We will continue our convention of eliding morphism actions. As before, these mappings are “automatic” from our choice of action on objects.

$$F_{\text{BitVec}}(\mathbb{N}, X) = (\mathbb{N}, (n \mapsto \left(\begin{array}{ll} \{\star\} & \text{if } n = 0 \\ \emptyset & \text{else} \end{array} \right) \oplus \left(\begin{array}{ll} \emptyset & \text{if } n = 0 \\ \{n_1\} \times \llbracket \text{Bool} \rrbracket \times X(n_1) & \text{if } n = \text{succ}(n_1) \end{array} \right))) \quad (20.5)$$

But here we encounter a problem: the endofunctor we have just defined only acts on \mathbb{N} -indexed families. Therefore, it is *not* an endofunctor on $\text{Fam}(\text{Set})$, whose objects can have arbitrary indexing sets. Yet this \mathbb{N} -indexing really is essential: our functor cannot perform a case split on zero and successors if our indexing set does not have such elements! Indeed, this ability to analyze the structure of a type index is precisely one of the key attributes that makes indexed inductive types more powerful than mere parametric polymorphism.

Therefore, to represent an indexed inductive type, we take the initial algebra of an endofunctor not on $\text{Fam}(\text{Set})$ itself, but on a restriction of $\text{Fam}(\text{Set})$ where we consider only families indexed by the indexing set of the indexed type of interest. This smaller category is called the *fiber* of the families fibration over A .

Definition 107: Fiber of the families fibration over A

The *fiber* of the families fibration over A , written $\text{Fam}(\text{Set})_A$, is a subcategory of $\text{Fam}(\text{Set})$ in which all objects have indexing set A . Precisely:

- Objects are families $(A, P : A \rightarrow \text{Ob}(\text{Set}))$
- Morphisms are pairs $(\text{id}_A, f^\sim : \prod_{a \in A} P(a) \rightarrow Q(a))$

That is, $\text{Fam}(\text{Set})_A$ is the preimage of id_A under U .

Thus, although it will at times be useful to consider the category $\text{Fam}(\text{Set})$ in its entirety, we will represent A -indexed inductive types by initial algebras of endofunctors on just the fiber of the families fibration over A . As before, the morphisms induced by initiality correspond to folds using the target algebras' structure maps.

Exercise 20.6

Identify the endofunctor on $\text{Fam}(\text{Set})_{\mathbb{N}}$ that characterizes $\text{Fin } n$, the type of “natural numbers less than n ” (or, isomorphically, $\mathbb{Z}/n\mathbb{Z}$):

```
inductive Fin : Nat -> Type where
```

```

2 | finZero : {n : Nat} -> Fin (succ n)
3 | finSucc : {n : Nat} -> Fin n -> Fin (succ n)

```

Hint: Notice that, for a given index, there maybe multiple (or no) applicable constructors.

Notice that we can also interpret refinements of simple types in the category of families. If $T = \mu F$ is an inductive type, then given an A -valued combining function f for T (i.e., an F -algebra f with carrier A) and a value $a : A$, the refinement of T where f yields value a , i.e.,

$$\{t : T \mid \text{rec}_A(f)(t) = a\}$$

can be translated as the family

$$(A, a \mapsto \{x \in \mu F \mid (f)_F(x) = a\}).$$

This is the “naïve” subset interpretation of refinements that we discussed previously. Our objective will be to turn this family into one that is the denotation of an indexed inductive type.

20.3 Constructing a Refinement Translation

Our basic recipe for the translation of $\{t : T \mid \text{rec}_A(f)(t) = a\}$ into an indexed inductive type will be as follows:

1. Lift F , an endofunctor on Set , to an endofunctor \hat{F}_A on $\text{Fam}(\text{Set})_A$.
2. Restrict the family produced by our lifting to those elements which are mapped to the appropriate value by $\text{rec}_A(f)$.

This recipe will require a number of auxiliary constructions, on which we now embark.

20.3.1 Another View of Families

For some of the constructions we give in this chapter, it will be helpful to have in mind an alternative perspective on $\text{Fam}(\text{Set})$. This perspective is inspired by the Curry–Howard correspondence between propositions and types.

Recall that the Curry–Howard correspondence suggests another way to interpret indexed inductive types: as *predicates* on the indexing type. For instance, we can define what it means for a natural number to be even by declaring a type of *proofs* of evenness:

```

1 inductive Even : Nat -> Type where
2 | evZero : Even 0
3 | evSucc : {n : Nat} -> Even n -> Even (succ (succ n))

```

A natural number n is even if the type $\text{Even } n$ is inhabited. Correspondingly, the interpretation of this inductive type as a family of sets $(\mathbb{N}, P_{\text{Even}})$ can be thought of as defining a *predicate* P_{Even} on natural numbers. With this interpretation, we think of a family (A, P) as containing proofs $P(a)$ that the represented predicate holds of elements $a \in A$. In particular, if $P(a)$ is inhabited, then the predicate holds of a ; if it is empty, the predicate does not.

Adopting the logical perspective on families suggested by the Curry–Howard correspondence, it is natural to consider the family that represents the constantly-true predicate on a given set, known as the set’s *truth predicate*.

Definition 108: Truth predicate and truth functor

Given a set A , the *truth predicate* for A is the family $(A, a \mapsto \mathbf{1})$, where $\mathbf{1}$ is the terminal object in Set (i.e., $\{\star\}$). The truth predicate is terminal in $\text{Fam}(\text{Set})_A$.

The functor $\top : \text{Set} \rightarrow \text{Fam}(\text{Set})$ that sends sets to their truth objects (and maps morphisms by pairing them with $_ \mapsto \text{id}_{\mathbf{1}}$) is called the *truth functor* for the families fibration.

20.3.2 Comprehension and Reindexing

Returning to our “recipe” from earlier, our first order of business is, given an endofunctor $F : \text{Set} \rightarrow \text{Set}$ with an initial algebra, to define an endofunctor $\hat{F}_A : \text{Fam}(\text{Set})_A \rightarrow \text{Fam}(\text{Set})_A$ whose initial algebra has a carrier representing an A -indexed type that “looks like” the inductive type characterized by F .

To define a translation between Set and $\text{Fam}(\text{Set})$, we need to be able to move between these two categories. We have already seen one functor in each direction: the truth predicate $\top : \text{Set} \rightarrow \text{Fam}(\text{Set})$ maps sets to the trivial family thereover, while the families fibration $U : \text{Fam}(\text{Set}) \rightarrow \text{Set}$ discards the family associated with an indexing set. We will now introduce one additional functor in the backward direction, which exploits the untypedness of Set to “collapse” a family in $\text{Fam}(\text{Set})$ into a single object in Set .

Definition 109: Comprehension functor

The *comprehension* of a family (A, P) , written $\{(A, P)\}$, is the set $\sum_{a \in A} P(a)$. We correspondingly define the *comprehension functor* $\{-\} : \text{Fam}(\text{Set}) \rightarrow \text{Set}$, which sends a family to its comprehension and a pair of morphisms to a morphism on pairs.

The comprehension functor $\{-\}$ and forgetful functor U are re-

lated by the natural transformation $\pi : \{-\} \Rightarrow U$ whose morphisms project (hence the notation π) the index from an element of a (“comprehended”) indexed family:

$$\pi_{(A,P)}(a, p) = a$$

We will also require means of mapping between $\text{Fam}(\text{Set})_A$ and $\text{Fam}(\text{Set})_B$ for distinct indexing sets A and B . Intuitively, this is because we interpret our indexed inductive types as endofunctors only over a specific fiber, so we will need to move between fibers in order to, for instance, fold over an indexed inductive type into a distinct output type. This will be accomplished using another functor, known as the *op-reindexing functor*.

Definition 110: Indexed coproduct

Let $f : A \rightarrow B$ be a function, and let (A, P) be a family of sets. The *indexed coproduct of (A, P) along f* , written $\Sigma_f(A, P)$, is defined as

$$\left(B, b \mapsto \sum_{a \in A} (b = f(a)) \times P(a) \right)$$

where $(b = f(a))$ is a set that is inhabited if the equality holds and uninhabited if not.

This definition induces the *op-reindexing functor* $\Sigma_f : \text{Fam}(\text{Set})_A \rightarrow \text{Fam}(\text{Set})_B$, whose action on morphisms lifts functions to act on the relevant component of tuples:

$$\Sigma_f(\text{id}_A, g^\sim) = (\text{id}_B, b \mapsto (a, h, p) \mapsto (a, h, g^\sim(a)(p)))$$

In other words, $\Sigma_f(A, P)$ maps a given b in the image of f to the members of the family P at each element of b ’s preimage. Indeed, an isomorphic (and perhaps more intuitive) characterization of $\Sigma_f(A, P)$ is given by

$$\left(B, b \mapsto \sum_{a \in f^{-1}\{b\}} P(a) \right). \quad (20.6)$$

We summarize the developments of this section in the diagram below: the category of families sits “above” the category of sets, with each fiber situated atop its corresponding index set. As discussed previously, we have functors running in “orthogonal” directions, with U , \top , and $\{-\}$ mapping between $\text{Fam}(\text{Set})$ and Set while Σ_f reindexes between fibers. We will have more to say about this diagram later.

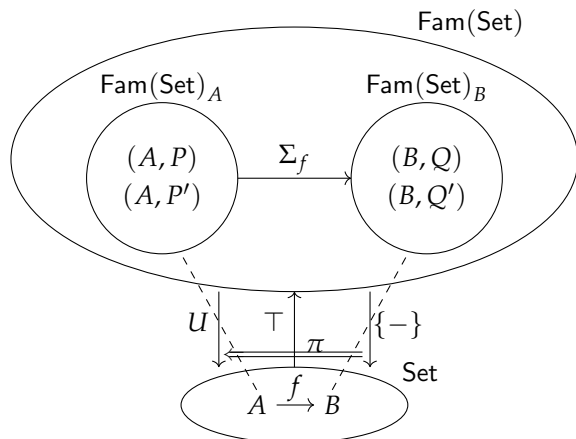


Figure 20.1: A summary of functors to be used in our translation.

Exercise 20.7

Fix a family (A, P) . Let $f : A \rightarrow B$ be a function. Consider the following diagram, where $\psi : (A, P) \rightarrow \Sigma_f(A, P)$ acts in the straightforward way:

$$\begin{array}{ccc}
 \{(A, P)\} & \xrightarrow{\{\psi\}} & \{\Sigma_f(A, P)\} \\
 \pi_{(A,P)} \downarrow & & \downarrow \pi_{\Sigma_f(A,P)} \\
 A & \xrightarrow{f} & B
 \end{array}$$

1. Show that this diagram commutes.
2. Show that $\{\psi\}$ is an isomorphism. (This shows that the families fibration has *very strong coproducts*.)

20.3.3 *Liftings*

We now have all the tools required to define a lifting of an endofunctor F on Set to an endofunctor \hat{F} on $\text{Fam}(\text{Set})$, which we may then restrict to $\text{Fam}(\text{Set})_A$. (It makes sense that we can lift “all the way” to $\text{Fam}(\text{Set})$, rather than just a particular fiber $\text{Fam}(\text{Set})_A$, since the type whose endofunctor we are lifting does not depend on a particular index.) Before we attempt to define such a lifting, however, we should first set out the desiderata we expect such a construction to satisfy.

As a basic criterion, we ought to expect that the indexing set produced by our “lifted” endofunctor \hat{F} is the set obtained by simply applying F to the original indexing set. Indeed, we can now drop

the scare quotes: this is precisely the definition of the lifting of an endofunctor to Fam(Set).

Definition 111: Lifting

A *lifting* of an endofunctor $F : \text{Set} \rightarrow \text{Set}$ is an endofunctor $\hat{F} : \text{Fam}(\text{Set}) \rightarrow \text{Fam}(\text{Set})$ such that $F \circ U = U \circ \hat{F}$.

However, this criterion says nothing about the behavior of the *family* produced by \hat{F} itself. Adopting the logical view again, we can formulate one such criterion: \hat{F} should “respect truth.” That is, if we think of $F(A)$ (given some A) as implicitly defining a predicate (namely, one which holds of all its elements), then \hat{F} should yield the same predicate—i.e., it should yield a predicate that holds of the same values—when “applied to” A . Of course, this intuitive phrasing does not type-check: $F(A)$ is not a predicate (it is a set), and \hat{F} cannot be applied to A (since A is not a family). But we can fill in some straightforward coercions to make this definition—which characterizes a so-called *truth-preserving* lifting—precise: to treat the objects in the codomain of F as predicates, we post-compose the truth functor; to allow \hat{F} to act on sets, we pre-compose the truth functor; and the notion that these functors always yield the “same predicate” as each other is captured by natural isomorphism.

Definition 112: Truth-preserving lifting

Let $F : \text{Set} \rightarrow \text{Set}$ be an endofunctor. A lifting of F is *truth-preserving* if there exists a natural isomorphism $\top \circ F \cong \hat{F} \circ \top$.

Exercise 20.8

Verify that

$$\begin{aligned} \hat{F}(A, P) &= (FA, _ \mapsto 1) \\ \hat{F}(f, f^\sim) &= (Ff, _ \mapsto \text{id}_1) \end{aligned}$$

defines a trivial truth-preserving lifting of $F : \text{Set} \rightarrow \text{Set}$.

The truth-preserving lifting above, however, is degenerate in the sense that it entirely forgets the structure of the original family. We therefore instead define the following truth-preserving lifting:¹³

$$\hat{F}(A, P) = \Sigma_{F\pi_{(A,P)}} \top(F\{(A, P)\})$$

Unlike in the exercise, we incorporate the “structure of P under F ” by applying F to the *comprehension* of the original family. We then

¹³ One can find a description of the lifting’s action on morphisms, along with a proof that it is truth-preserving, in [13].

re-separate the indexing set FA and associated family “under F ” by taking an indexed coproduct along $F\pi_{(A,P)}$. Thus, the resulting family can be seen as containing those elements which “wrap the structure of P in F ” (or, from a logical perspective, the resulting predicate holds only of those elements of FA whose “contents” in A were validated by P). Unfolding definitions makes these intuitions more apparent:

$$\begin{aligned}\hat{F}(A, P) &= \Sigma_{F\pi_{(A,P)}} \top(F\{(A, P)\}) \\ &= \Sigma_{F\pi_{(A,P)}} (F\{(A, P)\}, - \mapsto 1) \\ &= \left(FA, x \mapsto \sum_{y \in F\{(A, P)\}} (F\pi_{(A,P)}(y) = x) \times 1 \right) \\ &\cong (FA, x \mapsto \{y \in F\{(A, P)\} \mid (F\pi_{(A,P)})(y) = x\})\end{aligned}$$

As a concrete example, consider the truth-preserving lifting of F_{Nat} :

$$\begin{aligned}\hat{F}_{\text{Nat}}(A, P) &= \Sigma_{F_{\text{Nat}}} \top(\hat{F}_{\text{Nat}}\{(A, P)\}) \\ &= \Sigma_{F_{\text{Nat}}\pi_{(A,P)}} (1 \oplus \{(A, P)\}, - \mapsto 1) \\ &\cong \left(1 \oplus A, b \mapsto \sum_{a \in F_{\text{Nat}}\pi_{(A,P)}^{-1}[\{b\}]} 1 \right) \\ &\cong \left(1 \oplus A, \begin{cases} \text{inl } () \mapsto \text{inl } () \\ \text{inr } a \mapsto \{\text{inr } ((a, p), \star) \mid p \in P(a)\} \end{cases} \right) \\ &\cong \left(1 \oplus A, b \mapsto \begin{pmatrix} 1 & \text{if } b = \text{inl } () \\ \emptyset & \text{if } b = \text{inr } a \end{pmatrix} \oplus \begin{pmatrix} \emptyset & \text{if } b = \text{inl } () \\ \{a\} \times P(a) & \text{if } b = \text{inr } a \end{pmatrix} \right)\end{aligned}$$

Finally, since our objective is ultimately to obtain an endofunctor on $\text{Fam}(\text{Set})_A$ —not $\text{Fam}(\text{Set})$ —we define the functor \hat{F}_A to be the restriction of \hat{F} to $\text{Fam}(\text{Set})_A$. The attentive reader may notice that this is not an endofunctor: \hat{F}_A maps from the fiber $\text{Fam}(\text{Set})_A$ to the fiber $\text{Fam}(\text{Set})_{FA}$. However, it is the first half of a construction that will yield the endofunctor we desire; we will fill in the remaining piece that takes us back to $\text{Fam}(\text{Set})_A$ in the following section.

20.3.4 Translating Refinements

Recalling our original objective—to translate type refinements into indexed inductive types—we can see that truth-preserving liftings get us halfway to our goal: they transform the endofunctors that

characterize simple inductive types into analogous endofunctors on families. But our aim is not to produce an indexed inductive type that corresponds to the unrestricted simple one, but rather some refinement thereof. This refinement translation is the last piece we must develop.

In fact, we already have the tool we need to do this: op-reindexing. Recall that indexed coproducts (or, more generally, the op-reindexing functor) let us take a family and reindex it, such that the resulting family contains only those elements which are mapped by the reindexing function to a given value in its codomain. This is precisely what we do when we refine a datatype: given $P : T \rightarrow A$, the type $\{t : T \mid P(t) = a\}$ is an A -indexed type that, at a given index a , consists of the elements of T on which P produces value a . In other words, op-reindexing by f corresponds precisely to refinement by f .

Therefore, the construction to which we have building all this time can finally be written as follows:

Definition 113: Refinement translation

Let $F : \text{Set} \rightarrow \text{Set}$ be an endofunctor characterizing the inductive type μF . Let $(A, \alpha : FA \rightarrow A)$ be an F -algebra. We define the translation of the refinement type characterized by the family

$$(A, a \mapsto \{x \in \mu F \mid (\alpha)_F(x) = a\})$$

as the indexed inductive type μF^α characterized by the endofunctor $F^\alpha : \text{Fam}(\text{Set})_A \rightarrow \text{Fam}(\text{Set})_A$ such that

$$F^\alpha = \Sigma_\alpha \circ \hat{F}_A.$$

Once again, it is illuminating to expand this functor's action on objects:

$$\begin{aligned} F^\alpha(A, P) &= (\Sigma_A \circ \hat{F}_A)(A, P) \\ &\cong \Sigma_\alpha(FA, x \mapsto \{y \in F\{(A, P)\} \mid (F\pi_{(A, P)})(y) = x\}) \\ &\hspace{15em} \text{(As before)} \\ &= \left(A, a \mapsto \sum_{x \in FA} (a = \alpha(x)) \times \{y \in F\{(A, P)\} \mid (F\pi_{(A, P)})(y) = x\} \right) \\ &\cong (A, a \mapsto \{y \in F\{(A, P)\} \mid \alpha((F\pi_{(A, P)})(y)) = a\}) \quad (20.7) \end{aligned}$$

Upon cursory inspection, at least, the general shape of this family should seem reasonable: it is A -indexed, with elements drawn from a set of “ F -shaped objects” restricted to those “on which” (after a certain projection) α attains value a . To grasp the significance of this

final expression more completely, it may be helpful to return to it after observing the example below.

We illustrate this construction in practice with our example from the beginning: deriving the type of length-indexed vectors from the refinement of lists by their length. To simplify, we will consider just lists of Booleans, and attempt to derive the type `BitVec` from before.

First, letting $2 \cong 1 \oplus 1$ be the object denoted by the Boolean type, we recall the definition of $F_{\text{BitList}} : \text{Set} \rightarrow \text{Set}$ and the algebra $(\mathbb{N}, \text{length})$ thereover:

$$\begin{aligned} F_{\text{BitList}}(X) &= 1 \oplus 2 \times X \\ F_{\text{BitList}}(f) &= \text{id}_1 \oplus (\text{id}_2 \times f) \\ \text{length}(\text{inl } ()) &= 0 \\ \text{length}(\text{inr } (b, \ell)) &= \text{succ}(\ell) \end{aligned}$$

We now compute the corresponding indexed inductive type. Note that we use the isomorphic expansion of F^α from Equation (20.7) to avoid re-unfolding definitions:

$$\begin{aligned} &F_{\text{BitList}}^{\text{length}}(\mathbb{N}, P) \\ &\cong (\mathbb{N}, n \mapsto \{y \in F_{\text{BitList}}\{(\mathbb{N}, P)\} \mid \text{length}((F_{\text{BitList}}\pi_{(\mathbb{N}, P)})(y)) = n\}) \\ &= (\mathbb{N}, n \mapsto \{y \in 1 \oplus 2 \times (\mathbb{N}, P) \mid (y = \text{inl } () \text{ and } 0 = n) \text{ or} \\ &\quad (y = \text{inr } (b, (n_1, p)) \text{ and } \text{succ}(n_1) = n)\}) \\ &= \left(\mathbb{N}, \left\{ \begin{array}{l} 0 \mapsto \{\text{inl } ()\} \\ \text{succ}(n_1) \mapsto \{\text{inr } (b, (n_1, p)) \mid b \in 2, p \in P(n_1)\} \end{array} \right. \right) \\ &= (\mathbb{N}, (n \mapsto \left(\begin{array}{cc} 1 & \text{if } n = 0 \\ \emptyset & \text{else} \end{array} \right) \oplus \\ &\quad \left(\begin{array}{cc} \emptyset & \text{if } n = 0 \\ 2 \times (\{n_1\} \times P(n_1)) & \text{if } n = \text{succ}(n_1) \end{array} \right) \!)) \end{aligned}$$

Up to isomorphism (specifically, with a minor rearrangement and reassociation in the right-hand summand), this is precisely the functor F_{BitVec} derived in Equation (20.5)!

Exercise 20.9

Consider the following inductive type of trees with colored nodes (we refrain from calling these “red-black trees” in prose because they need not maintain the relevant invariants):

```

1 inductive RBList (A : Type) where
2   | leaf : RBTREE
3   | red  : A -> RBTREE -> RBTREE -> RBTREE

```

4 | black : A -> RBTree -> RBTree -> RBTree

1. Define the endofunctor $F_{\text{RBTree}} : \text{Set} \rightarrow \text{Set}$ such that $\llbracket \text{RBTree} \rrbracket \cong \mu F_{\text{RBTree}}$.
2. Define an F_{RBTree} -algebra bh that computes the *black height* of an RBTree , i.e., the maximum number of black nodes on a path from the root to a leaf. You may assume the existence of a function $\text{max} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$.
3. Using the above, compute $F_{\text{RBTree}}^{\text{bh}} : \text{Fam}(\text{Set})_{\mathbb{N}} \rightarrow \text{Fam}(\text{Set})_{\mathbb{N}}$, and read off of this expression an indexed inductive type of black height-indexed colored trees. Verify that your definition matches how you would intuitively define this type from scratch.

20.4 Correctness

These illustrative examples are compelling, but we have yet to prove that our foregoing construction always characterizes the indexed inductive type of interest. To begin with, we do not even know that the functor F^α admits an initial algebra in the first place, nor—if it does—whether the carrier of this algebra is in fact isomorphic to our naïve subset formulation of the refinement of μF by α . In this section, we will sketch the proof of these facts, referring the reader to the paper by Atkey et al. [2] for further details.

20.4.1 Preliminaries

We begin with a few auxiliary lemmas and definitions we will require for our proofs. First, recall our observation from the previous section that the functors $\top : \text{Set} \rightarrow \text{Fam}(\text{Set})$ and $U : \text{Fam}(\text{Set}) \rightarrow \text{Set}$ map in opposite directions. In fact, we can say more: these form an adjoint pair.

Lemma 20.4.1. The functors $U \dashv \top$ form an adjoint pair.

Proof. We must show $\text{Set}(U(A, P), B) \cong \text{Fam}(\text{Set})((A, P), \top B)$, i.e.,

$$\text{Set}(A, B) \cong \text{Fam}(\text{Set})((A, P), (B, _ \mapsto 1)).$$

Elements of the right-hand set are pairs of morphisms $(f : A \rightarrow B, f^\sim : \prod_{a \in A} P(a) \rightarrow 1)$. But by terminality, there is only ever one choice of f^\sim , so the right-hand side is isomorphic to the set of functions $f : A \rightarrow B$, i.e., $\text{Set}(A, B)$. Naturality is straightforward. \square

It turns out that \top is also *left* adjoint, in particular to the comprehension functor $\{-\}$, but we will not make use of this fact.

As before, we now turn from functors between $\text{Fam}(\text{Set})$ and Set (i.e., between families and their index sets) to those between fibers $\text{Fam}(\text{Set})_A$ and $\text{Fam}(\text{Set})_B$ (i.e., between families with distinct index sets). While our construction in the previous section required only the op-reindexing functor Σ_f , the following alternative means of reindexing will be useful for our proof.

Definition 114: Reindexing functor

Let $f : A \rightarrow B$ be a function (of sets). The *reindexing functor* $f^* : \text{Fam}(\text{Set})_B \rightarrow \text{Fam}(\text{Set})_A$ is defined by

$$\begin{aligned} f^*(B, Q) &= (A, Q \circ f) \\ f^*(\text{id}_B, g^\sim) &= (\text{id}_A, g^\sim \circ f) \end{aligned}$$

Notice that this functor, in contrast with Σ_f , is contravariant: the domain of f^* is the fiber over the codomain of f , and vice versa. (Hence the “op” prefix in “op-reindexing.”) Thus, given a function $f : A \rightarrow B$, we can produce functors running in both directions between $\text{Fam}(\text{Set})_B$ and $\text{Fam}(\text{Set})_A$. Moreover, these functors are adjoint.

Lemma 20.4.2. Given a function $f : A \rightarrow B$, the functors $\Sigma_f \dashv f^*$ form an adjoint pair.

Proof. We show

$$\text{Fam}(\text{Set})_B(\Sigma_f(A, P), (B, Q)) \cong \text{Fam}(\text{Set})_A((A, P), f^*(B, Q)).$$

Recall that morphisms in these fibers are fully characterized by their second component (i.e., their mapping on families), so, unfolding definitions (and using the isomorphic representation of the indexed coproduct from Equation (20.6)), it suffices to show

$$\left\{ g : \prod_{b \in B} \sum_{a \in f^{-1}[\{b\}]} P(a) \rightarrow Q(b) \right\} \cong \left\{ h : \prod_{a \in A} P(a) \rightarrow Q(f(a)) \right\}$$

satisfying naturality. It is straightforward to check that ι (mapping from left to right) and its putative inverse defined below indeed form such an isomorphism between these sets:

$$\begin{aligned} \iota(g) &= a \mapsto p \mapsto g(f(a))(a, p) \\ \iota^{-1}(h) &= b \mapsto (a, p) \mapsto h(a)(p) \end{aligned} \quad \square$$

We will also make use of a lemma that characterizes the morphisms into a reindexed family: they are isomorphic to those morphisms into the original family that mapped indexing sets according to f .

Lemma 20.4.3. Let (A, P) and (B, Q) be objects in $\text{Fam}(\text{Set})$. Let $f : A \rightarrow B$ be a function. Then there is an isomorphism between the set of $\text{Fam}(\text{Set})$ -morphisms $h : (A, P) \rightarrow (B, Q)$ such that $Uh = f$ and $\text{Fam}(\text{Set})_A$ morphisms $(A, P) \rightarrow f^*(B, Q)$:

$$\{h \in \text{Fam}(\text{Set})((A, P), (B, Q)) \mid Uh = f\} \cong \text{Fam}(\text{Set})_A((A, P), f^*(B, Q))$$

Exercise 20.10

Prove Lemma 20.4.3.

We can now update Figure 20.1 to summarize these additional results:

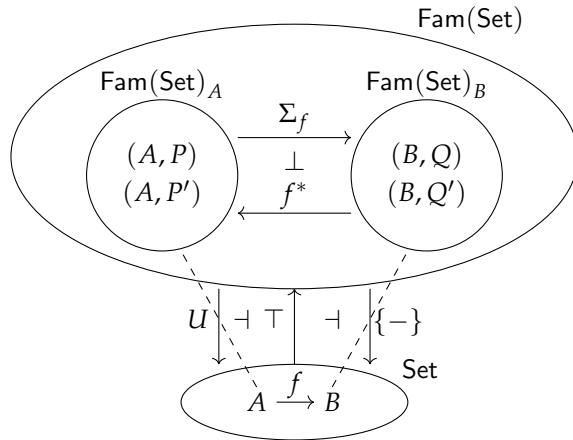


Figure 20.2: A revision of Figure 20.1, now with adjoint relationships and the reindexing functor (π has been removed for clarity).

20.4.2 A Proof of Correctness

We want to show that, given an endofunctor F on Set with an initial algebra, the endofunctor F^α on $\text{Fam}(\text{Set})_A$ has one as well: that is, we must find an initial object in Alg_{F^α} . Since we already have an initial object in Alg_F , it will be useful to relate Alg_{F^α} to this category. Mirroring our construction of F^α itself, we will do this in two stages: first, we will consider algebras over $\hat{F} : \text{Fam}(\text{Set}) \rightarrow \text{Fam}(\text{Set})$, which we relate to algebras over F ; then, we will limit our attention to an “ α -based restriction” of \hat{F} -algebras.

Algebras over \hat{F} are elements of the category $\text{Alg}_{\hat{F}}$. Since $\text{Fam}(\text{Set})$ “sits atop” Set (in the sense that a family P is “atop” its indexing set UP), the category $\text{Alg}_{\hat{F}}$ likewise “sits atop” Alg_F , a fact witnessed by the following *algebras fibration*.

Definition 115: Algebras fibration

The *algebras fibration* U^{Alg} induced by the families fibration U is a functor $U^{\text{Alg}} : \text{Alg}_{\hat{F}} \rightarrow \text{Alg}_F$ defined by

$$U^{\text{Alg}}(P, k : \hat{F}P \rightarrow P) = (UP, Uk)$$

$$U^{\text{Alg}}(f : (P, k) \rightarrow (Q, k')) = Uf : (UP, Uk) \rightarrow (UQ, Uk')$$

It is worth pausing to type-check this definition, in which we have taken some notational shortcuts for concision. Recall that $U : \text{Fam}(\text{Set}) \rightarrow \text{Set}$ maps families to their indexing sets. Since \hat{F} is an endofunctor on *families*, P comprises a family indexed by some set; k is a morphism of families, which is a pair of functions; and f is an algebra morphism (as defined in Definition 104) over an endofunctor on families and thus also a morphism of families (satisfying a certain commutativity condition). Thus, U has an action on each of these.

The algebras over F inherit an additional useful property from those over \hat{F} : in particular, if F has an initial algebra, then so does \hat{F} :

Theorem 20.4.4. Suppose F has an initial algebra. Then \hat{F} has an initial algebra too, and its carrier is $\top(\mu F)$.

Proof. The key insight for this proof is that F -algebras with carrier $\{(A, P)\}$ are interderivable with \hat{F} -algebras with carrier (A, P) . A proof of this fact is provided by Ghani et al. [13]; we will not give one here. Using this fact, the result follows almost immediately: the carrier μF of the initial F -algebra is isomorphic to $\sum_{t \in \mu F} \{(- \mapsto \star)\}$ as a set, since the summand is a singleton. This set is precisely $\{\top(\mu F)\}$, so that $\top(\mu F)$ is the carrier of the initial \hat{F} -algebra by the aforementioned interderivability fact. \square

However, since F^α is a restriction of \hat{F} by values of an algebra (A, α) , we want to look at a correspondingly restricted category of algebras, not $\text{Alg}_{\hat{F}}$ in its entirety. We thus restrict the category $\text{Alg}_{\hat{F}}$ to an appropriate *fiber* over our refining algebra. This definition, given below, is directly analogous to the definition of fibers over $\text{Fam}(\text{Set})$; however, we phrase it slightly differently to suggest the generality of such a construction. (One may find it worthwhile to “pattern-match” the tuple-valued variables below and unfold the definitions of U^{Alg} and U to see this correspondence more directly.) Note that we are also now adopting the convention of abbreviating (A, α) as simply α .

Definition 116: Fiber of the algebras fibration over α

The fiber of the algebras fibration over α , or equivalently the category of \hat{F} -algebras over α , denoted by $(\text{Alg}_{\hat{F}})_{\alpha}$, is defined as follows:

- Objects are \hat{F} -algebras (P, k) such that $U^{\text{Alg}}(P, k) = \alpha$.
- Morphisms are \hat{F} -algebra morphisms $f : (P, k) \rightarrow (Q, k')$ such that $Uf = \text{id}_{\alpha}$.

Once again, it is worth briefly type-checking this definition. For instance, the criterion on objects is well defined: $k = (k_1, k_2)$ is a morphism $\hat{F}P \rightarrow P$, which is an abbreviation of $\hat{F}(A, P) \rightarrow (A, P)$ with the indexing set explicit. By the definition of $\text{Fam}(\text{Set})$, k_1 maps between the carriers; since $\hat{F}(A, P)$ has carrier FA by definition, we thus have $k_1 : FA \rightarrow A$. And since $U^{\text{Alg}}(P, k) = U^{\text{Alg}}((A, P), (k_1, k_2)) = (A, k_1)$, this is of the same type as α , i.e., as $(A, \alpha : FA \rightarrow A)$.

Moreover, it turns out that $(\text{Alg}_{\hat{F}})_{\alpha}$ is the same category as $\text{Alg}_{F\alpha}$, just described as a “decomposition” in terms of \hat{F} and α .

Lemma 20.4.5. Let (A, α) be an F -algebra. Then there is an isomorphism of categories

$$(\text{Alg}_{\hat{F}})_{\alpha} \cong \text{Alg}_{F\alpha}$$

Proof. Observe that an object of $(\text{Alg}_{\hat{F}})_{\alpha}$ is a pair of a family and a $\text{Fam}(\text{Set})$ -morphism—i.e., a tuple $((A, P), k : \hat{F}(A, P) \rightarrow (A, P))$ —such that $Uk = \alpha$. By Lemma 20.4.3, we can characterize such a $\text{Fam}(\text{Set})$ -morphism as a $\text{Fam}(\text{Set})_{FA}$ -morphism via the reindexing functor, so the object is equivalently expressed as $((A, P), k' : \hat{F}(A, P) \rightarrow \alpha^*(A, P))$. In turn, we observe that morphisms $k' : \hat{F}(A, P) \rightarrow \alpha^*(A, P)$ are themselves in isomorphism with $\text{Fam}(\text{Set})_A$ -morphisms $k'' : \Sigma_{\alpha}\hat{F}(A, P) \rightarrow (A, P)$ by Lemma 20.4.2. Recalling the definition of F^{α} , we can rewrite this as $k'' : F^{\alpha}(A, P) \rightarrow (A, P)$. Thus, we can again equivalently rewrite our object as $((A, P), k'' : F^{\alpha}(A, P) \rightarrow (A, P))$, which is precisely an object in $\text{Alg}_{F\alpha}$. Transporting a morphism along the corresponding sequence of isomorphic transformations yields the required action on morphisms. \square

In order to work with this new formulation of the category, we will reproduce for $(\text{Alg}_{\hat{F}})_{\alpha}$ several results analogous to those we showed for the families fibration. The following lemmas are analogues of Lemmas 20.4.2 and 20.4.3; their proofs are technical but not exceedingly difficult, so we elide them.

Lemma 20.4.6. Let f be an F -algebra morphism $(A, \alpha) \rightarrow (B, \beta)$. Then there are functors $\Sigma_f^{\text{Alg}} : (\text{Alg}_{\hat{F}})_{\alpha} \rightarrow (\text{Alg}_{\hat{F}})_{\beta}$ and $f^{*\text{Alg}} : (\text{Alg}_{\hat{F}})_{\beta} \rightarrow (\text{Alg}_{\hat{F}})_{\alpha}$ such that $\Sigma_f^{\text{Alg}} \dashv f^{*\text{Alg}}$ and for any \hat{F} -algebras $((A, P), k)$ and $((B, Q), k')$,

- The \hat{F} -algebra $\Sigma_f^{\text{Alg}}((A, P), k)$ has carrier $\Sigma_f(A, P)$, and
- The \hat{F} -algebra $f^*\text{Alg}((B, Q), k')$ has carrier $f^*(B, Q)$.

Lemma 20.4.7. Let (A, α) and (B, β) be \hat{F} -algebras. Fix an object $((A, P), k)$ in $(\text{Alg}_{\hat{F}})_{\alpha}$ and $((B, Q), k')$ in $(\text{Alg}_{\hat{F}})_{\beta}$. Let $f : (A, \alpha) \rightarrow (B, \beta)$ be an \hat{F} -algebra morphism. Then there is an isomorphism between the set of $\text{Alg}_{\hat{F}}$ -morphisms $h : k \rightarrow k'$ such that $U^{\text{Alg}}h = f$ and the set of morphisms $k \rightarrow f^*\text{Alg}(k')$:

$$\{h \in \text{Alg}_{\hat{F}}(k, k') \mid U^{\text{Alg}}h = f\} \cong \text{Alg}_{\hat{F}}(k, f^*\text{Alg}(k')).$$

We can now draw a diagram analogous to Figure 20.2 for the algebras fibration. Note that we temporarily drop our abbreviated notation for families and algebras to better analogize this diagram to the one for the families fibration.

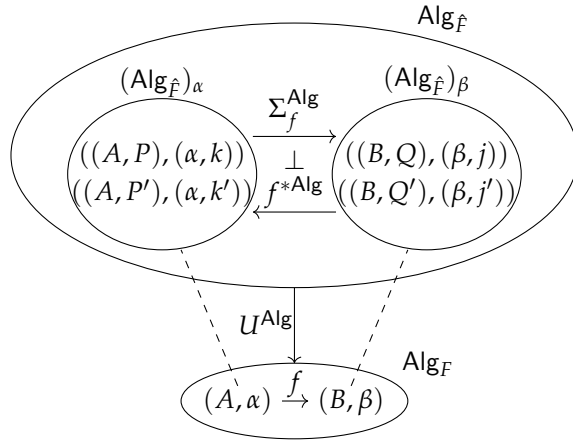


Figure 20.3: A summary of constructions for the algebras fibration.

Equipped with these intermediate results, we can now finally prove what we set out to show: that F^{α} admits an initial algebra, and that its carrier is isomorphic to the “naïve” refinement interpretation. We make one brief observation: we can rewrite our previous “naïve” interpretation of a refinement by $(\alpha) : \mu F \rightarrow A$ using the truth and op-reindexing functors as follows:

$$(A, a \mapsto \{x \in \mu F \mid (\alpha)(x) = a\}) \cong \Sigma_{(\alpha)} \top(\mu F).$$

Using this formulation, we now state our main theorem.

Theorem 20.4.8. The endofunctor F^{α} has an initial algebra, and its carrier is $\Sigma_{(\alpha)} \top(\mu F)$.

Proof. We will first construct a candidate algebra in $(\text{Alg}_{\hat{F}})_{\alpha}$, then show that it is initial. Note that, by Lemma 20.4.5, this is equivalent to identifying an F^{α} -algebra and showing its initiality in $\text{Alg}_{F^{\alpha}}$.

From Theorem 20.4.4, we obtain an initial algebra $(\top(\mu F), \text{in}_{\hat{F}})$ in $\text{Alg}_{\hat{F}}$. We will use this to construct an initial object in the fiber over α . Mirroring the construction of F^α , we apply the op-reindexing functor $\Sigma_{\langle \alpha \rangle}^{\text{Alg}}$ for the algebras fibration to obtain

$$\left(\Sigma_{\langle \alpha \rangle}^{\text{Alg}} \top(\mu F), \Sigma_{\langle \alpha \rangle}^{\text{Alg}} \text{in}_{\hat{F}} : \hat{F}(\top(\mu F)) \rightarrow \top(\mu F) \right).$$

We will aim to show that this object is our desired initial object in $(\text{Alg}_{\hat{F}})_\alpha$.

First, observe that by Lemma 20.4.6, we have

$$\Sigma_{\langle \alpha \rangle}^{\text{Alg}} \top(\mu F) \cong \Sigma_{\langle \alpha \rangle} \top(\mu F),$$

i.e., the carrier of this algebra is our naïve refinement encoding. So it remains only to show that the algebra is indeed initial.

We show initiality directly from the definition: we show that there is exactly one morphism between $\Sigma_{\langle \alpha \rangle}^{\text{Alg}} \text{in}_{\hat{F}}$ and an arbitrary \hat{F} -algebra $k : \hat{F}(A, P) \rightarrow (A, P)$ (hereafter we will denote each algebra under discussion by its structure map alone). We proceed by considering a series of isomorphisms of hom-sets:

$$\begin{aligned} & (\text{Alg}_{\hat{F}})_\alpha \left(\Sigma_{\langle \alpha \rangle}^{\text{Alg}} \text{in}_{\hat{F}}, k \right) \\ & \cong (\text{Alg}_{\hat{F}})_{\text{in}_F} (\text{in}_{\hat{F}}, \langle \alpha \rangle^*(k)) \quad (\text{Adjointness portion of Lemma 20.4.6}) \\ & \cong \{h \in \text{Alg}_{\hat{F}}(\text{in}_{\hat{F}}, k) \mid U^{\text{Alg}} h = \langle \alpha \rangle\} \quad (\text{Lemma 20.4.7}) \end{aligned}$$

Notice that this final set is a subset of morphisms out of an initial object, so it can have at most one element (namely, $\langle k \rangle$). If this is to be a singleton set, then it must indeed contain $\langle k \rangle$, so in particular it must be that

$$U^{\text{Alg}} \langle k \rangle = \langle \alpha \rangle.$$

Recall that U^{Alg} is left adjoint, so it preserves colimits: the above equation (of morphisms from initial objects) thus holds.

Accordingly, the algebra $\Sigma_{\langle \alpha \rangle}^{\text{Alg}} \text{in}_{\hat{F}}$ is indeed initial in $(\text{Alg}_{\hat{F}})_\alpha$, completing the proof. \square

Before concluding this section, it is worth noting that the construction given here can be extended in various ways, such as to support refinements defined in terms of functions that are not directly expressed as folds. We direct the reader to Atkey et al. [2] for details.

Let us finally take stock of what we have done. We began by interpreting simple inductive types as initial algebras $(\mu F, \text{in}_{\mu F})$ of endofunctors F . On this interpretation, maps out of the datatype defined in terms of folds are the F -algebra homomorphisms $\langle \alpha \rangle$ guaranteed by initiality. We then gave a categorical semantics of indexed

inductive types in terms of initial algebras of endofunctors on *families*, in which we gave a “naïve” subset interpretation of the refinement $\{x : \mu F \mid (\alpha)(x) = a\}$. To identify the indexed inductive type corresponding to this semantic interpretation, we constructed the endofunctor on families F^a using a *lifting* of the functor F characterizing the unrefined simple inductive type. Lastly, we showed that this construction yielded an initial algebra with the carrier we identified in our naïve translation.

20.5 *A (Literal) Invitation to Fibered Category Theory*

Throughout this chapter, we have followed Atkey et al. [2] in carefully presenting the material so as to admit a generalization to the abstract language of *fibrations*, and have gradually adjusted our presentation over the course of the chapter toward a more abstract framing in terms of fibrational constructions. Regrettably, we do not have the space here to present the full generalization—a construction that can be carried out in any *full Cartesian Lawvere category with very strong coproducts*—or to discuss fibrations in their own right.

Instead, we direct the reader to the following sources. The original development of this material by Atkey et al. [2] gestures further toward its generalization, in addition to providing several further constructions that translate broader classes of refinements into indexed inductive types. Ghani et al.’s paper [13] includes the original, abstract derivation of the truth-preserving lifting \hat{F} we employ in this paper. Some relevant background can also be found in Jacobs’ paper on comprehension categories [16], while the same author’s textbook *Categorical Logic and Type Theory* [17] provides an exposition of fibered category theory and several applications. A selection of material from among these sources should provide sufficient background to interpret this chapter in a general fibrational setting, thereby characterizing the precise categorical structure required to unify refinement and indexed types.

Solutions to Exercises

Solution 20.1:

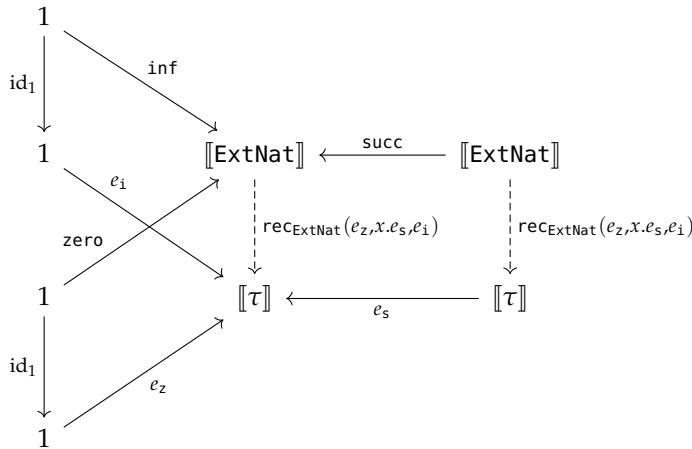
Pick $\llbracket \text{true} \rrbracket = \text{inl}$ and $\llbracket \text{false} \rrbracket = \text{inr}$, respectively the left and right injections into the coproduct, and take $\llbracket \text{if}(t, e) \rrbracket = \langle t, e \rangle$. Commutativity and recursor uniqueness are immediate from the universal property for coproducts.

Solution 20.2:

Since the recursor for ExtNat is not explicitly given, we first derive its typing rule (its equational laws are analogous to those for the regular natural-number recursor):

$$\frac{\vdash e_z : \tau \quad x : \tau \vdash e_s : \tau \quad \vdash e_i : \tau}{\vdash \text{rec}(e_z, x.e_s, e_i) : \tau}$$

As the hint alludes to, this example differs from the preceding ones because there are three constructors rather than two, yielding a third square in our diagram; it is a “revolving door” rather than a pair of squares. It looks as follows:


Solution 20.3:

1. The actions of F_{SelfTree} on objects and morphisms are given below:

$$F_{\text{SelfTree}}(X) = 1 \oplus A \times X^A$$

$$F_{\text{SelfTree}}(e) = \text{id}_1 \oplus (\text{id}_A \times \lambda(e \circ \text{app}))$$

2. The relevant algebra is

$$\alpha_{\text{selfHeight}} : F_{\text{SelfTree}} \llbracket \text{Nat} \rrbracket \rightarrow \llbracket \text{Nat} \rrbracket = [0, \text{succ} \circ \text{app} \circ \text{flip}].$$

We denote by $\text{flip} : A \times X^A \rightarrow X^A \times A$ the isomorphism that witnesses the commutativity of products.

Solution 20.4:

Fix objects A, B and consider their exponential. Since $\text{app} : B^A \times A \rightarrow B$, we are led to define the endofunctor $F(X) = X^A \times A$, whose action on morphisms is given by $F(e) = \lambda(e \circ \text{app}) \times \text{id}_A$. (Notice that, despite the type of app mentioning two objects, there is only one possible choice of endofunctor: if we tried to define the functor in A , it would be contravariant.)

We claim that the identity morphism for the initial object 0 is the initial F -algebra. To see why, first observe that $F(0) = 0^A \times A \cong 0 \times A \cong 0$, so the choice of initial algebra $\text{id}_0 : F(0) \rightarrow 0$ (transported across the relevant isomorphisms) type-checks. Moreover, its initiality is immediate from the fact that 0 is initial.

By Lambek's lemma, we thus have that $\mu F \cong 0$, so the inductive type that this functor characterizes is empty. We can read off the corresponding inductive definition:

```
1 inductive Empty' (A : Type) : Type where
2   | mk : (A -> Empty' A) -> Empty' A -> Empty' A
```

Solution 20.5:

We translate each rule in turn.

- **T-FOLD:** As previously discussed, the positivity hypothesis of this rule is analogous to the requirement that the endofunctor F denoted by $t.\tau$ has an initial algebra. The substitution $\tau[\mu t.\tau/t]$ corresponds to plugging in $\llbracket \tau \rrbracket$ for the functor variable t in $F(t)$: it is the functor action on objects. Thus, the conclusion of the rule induces the following morphism:

$$\llbracket F(\mu F) \rrbracket \xrightarrow{\text{fold}} \llbracket \mu F \rrbracket$$

- **T-REC:** Now that we know how to interpret type-variable substitution, this translation proceeds mechanically: the hypothesis says we have a morphism $e_1 : \llbracket \rho \rrbracket \rightarrow F \llbracket \rho \rrbracket$, and the conclusion uses this to produce a morphism $\mu F \rightarrow \llbracket \rho \rrbracket$:

$$\begin{array}{ccc} & & \llbracket \mu F \rrbracket \\ & & \downarrow \\ & & \text{rec}(x.e_1) \\ & & \downarrow \\ F \llbracket \rho \rrbracket & \xrightarrow{e_1} & \llbracket \rho \rrbracket \end{array}$$

- β^{rec} : Now we must contend with the function $\text{map}(e)$. Observe that $e : (t.\tau)[\rho/t]$ from rule T-MAP, so we know that the type of e denotes a functor application $F \llbracket \rho \rrbracket$. The key insight—which is most easily observed by inspecting the sample β rules for this function—is that the function $\text{map}(e)$ denotes precisely the action of F on *morphisms*, i.e., terms. Thus, we translate $\text{map}(e)$ as the morphism action $F(-)$.

We can now interpret the β^{rec} rule. The left-hand side corresponds to the morphism $\text{rec}(x.e_1) \circ \text{fold}$. The right-hand side becomes (removing the explicit point y for clarity, and recalling that substitution becomes morphism composition) $e_1 \circ \llbracket \text{map}(e_2)(\text{rec}(x.e_1)) \rrbracket$, which unfolds to $e_1 \circ F(\text{rec}(x.e_1))$. Since this rule is equational, we have that these two morphisms are equal, i.e.,

$$\text{rec}(x.e_1) \circ \text{fold} = e_1 \circ F(\text{rec}(x.e_1))$$

This just says that if we stitch together the diagrams from the preceding two rules, the resulting square commutes:

$$\begin{array}{ccc} F(\mu F) & \xrightarrow{\text{fold}} & \llbracket \mu F \rrbracket \\ \downarrow F(\text{rec}(x.e_1)) & & \downarrow \text{rec}(x.e_1) \\ F \llbracket \rho \rrbracket & \xrightarrow{e_1} & \llbracket \rho \rrbracket \end{array}$$

- η^{rec} : The hypothesis of this rule translates to

$$e \circ \text{fold} = e_1 \circ F(\text{rec}(x.e_1))$$

and its conclusion to

$$e = \text{rec}(x.e_1).$$

In other words, any morphism e satisfying the equality from the β rule must be equal to the recursor; that is, the recursor is unique.

We reflect this in our diagram from above:

$$\begin{array}{ccc} F(\mu F) & \xrightarrow{\text{fold}} & \llbracket \mu F \rrbracket \\ \downarrow F(\text{rec}(x.e_1)) & & \downarrow \text{rec}(x.e_1) \\ F \llbracket \rho \rrbracket & \xrightarrow{e_1} & \llbracket \rho \rrbracket \end{array}$$

Observe that (renaming fold to c , e_1 to f_c , and ρ to A) we have now arrived at the same diagram we derived in our previous translation!

Solution 20.6:

$$F_{\text{Fin}}(\mathbb{N}, P) = \left(\mathbb{N}, \begin{cases} 0 \mapsto \emptyset \\ \text{succ}(n) \mapsto 1 \oplus P(n) \end{cases} \right)$$

Solution 20.7:

First, note that the “straightforward” morphism ψ is concretely

$$\psi(a, p) = (f, a \mapsto p \mapsto (a, p)).$$

(Note that we are using the isomorphic formulation of indexed coproducts given in Equation (20.6) to minimize bureaucracy.)

1. Unfolding definitions, we find that

$$(f \circ \pi_{(A,P)})(a, p) = f(a)$$

and

$$(\pi_{\Sigma_f(A,P)} \circ \{\psi\})(a, p) = \pi_{\Sigma_f(A,P)}(f(a), a, p) = f(a)$$

so that

$$f \circ \pi_{(A,P)} = \pi_{\Sigma_f(A,P)} \circ \{\psi\}$$

as required.

2. Define the putative inverse of $\{\psi\}$ to be $\iota : \{\Sigma_f(A, P)\} \rightarrow \{(A, P)\}$ given by

$$\iota(b, (a, p)) = (a, p)$$

Then observe that

$$\{\psi\}(\iota(b, (a, p))) = \{\psi\}(a, p) = (f(a), (a, p)) = (b, (a, p))$$

where the last equality follows by the definition of $\Sigma_f(A, P)$, and moreover

$$\iota(\{\psi\}(a, p)) = \iota(f(a), (a, p)) = (a, p)$$

so that $\{\psi\}$ is indeed an isomorphism with inverse ι .

Solution 20.8:

This is clearly a lifting, since

$$F(U(A, P)) = FA = U(FA, _ \mapsto 1) = U(\hat{F}(A, P))$$

and

$$F(U(f, f^\sim)) = Ff = U(f, _ \mapsto \text{id}_1) = U(\hat{F}(f, f^\sim)).$$

Moreover, that it is truth preserving is witnessed by the natural isomorphism

$$\eta_{(A,P)} = \text{id}_{(A,P)}$$

since, noting that

$$\top(F(A)) = (FA, _ \mapsto 1) = \hat{F}(\top(A))$$

and

$$\top(F(f)) = (f, _ \mapsto \text{id}_1) = \hat{F}(\top(f))$$

the following square trivially commutes:

$$\begin{array}{ccc} (\top \circ F)(A) & \xrightarrow{\eta_{(A, _ \mapsto 1)}} & (\hat{F} \circ \top)(A) \\ \downarrow & & \downarrow \\ (\top \circ F)(f) & & (\hat{F} \circ \top)(f) \\ \downarrow & & \downarrow \\ (\top \circ F)(B) & \xrightarrow{\eta_{(B, _ \mapsto 1)}} & (\hat{F} \circ \top)(B) \end{array}$$

Solution 20.9:

1. The endofunctor F_{RBTree} has the following actions on objects and morphisms:

$$\begin{aligned} F_{\text{RBTree}}(X) &= 1 \oplus ((A \times (X \times X)) \oplus (A \times (X \times X))) \\ F_{\text{RBTree}}(e) &= \text{id}_1 \oplus ((\text{id}_A \times (e \times e)) \oplus (\text{id}_A \times (e \times e))) \end{aligned}$$

2. The desired F_{RBTree} -algebra is

$$\alpha_{\text{bh}} = [0, [\text{succ} \circ \max \circ \pi_2, \max \circ \pi_2]]$$

3. We compute as follows:

$$\begin{aligned} &F_{\text{RBTree}}^{\text{bh}}(\mathbb{N}, P) \\ &\cong (\mathbb{N}, n \mapsto \{y \in F_{\text{RBTree}}(\mathbb{N}, P) \mid \text{bh}((F_{\text{RBTree}} \pi_{(\mathbb{N}, P)})(y)) = n\}) \\ &= (\mathbb{N}, n \mapsto \{y \in 1 \oplus ((A \times ((\mathbb{N}, P) \times (\mathbb{N}, P))) \oplus (A \times ((\mathbb{N}, P) \times (\mathbb{N}, P)))) \mid \\ &\quad (y = \text{inl} () \text{ and } 0 = n) \text{ or} \\ &\quad (y = \text{inr} (\text{inl} (a, ((m_\ell, p_\ell), (m_r, p_r)))) \text{ and } n = \text{succ}(\max(n_\ell, n_r))) \\ &\quad (y = \text{inr} (\text{inr} (a, ((m_\ell, p_\ell), (m_r, p_r)))) \text{ and } n = \max(n_\ell, n_r))\}) \\ &= \left(\mathbb{N}, \left(n \mapsto \begin{pmatrix} 1 & \text{if } n = 0 \\ \emptyset & \text{else} \end{pmatrix} \right) \oplus \right. \\ &\quad \left(\left(\begin{pmatrix} A \times ((\{m_\ell\} \times P(m_\ell)) \times (\{m_r\} \times P(m_r))) & \text{if } n = \text{succ}(\max(m_\ell, m_r)) \\ \emptyset & \text{else} \end{pmatrix} \right) \oplus \right. \\ &\quad \left. \left(\begin{pmatrix} A \times ((\{m_\ell\} \times P(m_\ell)) \times (\{m_r\} \times P(m_r))) & \text{if } n = \max(m_\ell, m_r) \\ \emptyset & \text{else} \end{pmatrix} \right) \right) \right) \end{aligned}$$

One can verify that this is isomorphic to the translation of the following inductive type:

```

1 inductive BHRBTree : Nat -> Type where
2   | leaf : BHRBTree 0
3   | red  : {m n : Nat} -> A -> BHRBTree m -> BHRBTree n -> BHRBTree (succ
      (max m n))
4   | red  : {m n : Nat} -> A -> BHRBTree m -> BHRBTree n -> BHRBTree (max
      m n)

```

Solution 20.10:

Recall that $\text{Fam}(\text{Set})$ -morphisms from (A, P) to (B, Q) are pairs

$$\left(h : A \rightarrow B, h^\sim : \prod_{a \in A} P(a) \rightarrow Q(h(a)) \right).$$

The left-hand set in the lemma statement is the restriction of such pairs to those whose first element is f , so they are of the form

$$\left(f, h^\sim : \prod_{a \in A} P(a) \rightarrow Q(f(a)) \right).$$

By the definitions of $\text{Fam}(\text{Set})_A$ and f^* , morphisms in the right-hand set are pairs

$$\left(\text{id}_A, f^\sim : \prod_{a \in A} P(a) \rightarrow Q(f(a)) \right).$$

Since the first component of every pair in either set is fixed, both sets are isomorphic to the set of functions of type

$$\prod_{a \in A} P(a) \rightarrow Q(f(a))$$

and thus to each other.

Algebras and automata

Michael Zhang

1 Algebras

In algebra, there are many interesting structures to study, such as groups, rings, vector spaces, etc. If we look at the definitions for some common algebraic structures:

- A group is a tuple $(G, \cdot, (-)^{-1})$ satisfying some axioms.
- A ring is a tuple $(R, +, \cdot)$ satisfying some axioms.
- A boolean algebra is a tuple (B, \wedge, \vee, \neg) satisfying some axioms.

They all seem to look kind of familiar – a carrier, usually a set, followed by some operations on that carrier, satisfying some axioms. It may be useful to look for an abstraction over all of these.

In order to make it easier to think about abstractions for a universal algebra, we'll start by focusing our work in the category of sets **Set**. So our carrier will be some object in **Set**.

Now, consider an endofunctor $F : \mathbf{Set} \rightarrow \mathbf{Set}$. This will have two parts: an action on objects, which is simply a set function, and an action on morphisms. If we take X to be the input carrier set, then $F(X)$ is essentially the representation of *inputs* to operations you can take with X . Then, we have a morphism $\alpha : F(X) \rightarrow X$, known as the *structure* map, which defines how the operations take place.

For example:

- the identity element of a group is not dependent on anything, so we say $F(X) = 1$, and so the morphism $\alpha : 1 \rightarrow X$ simply selects the identity element with no further input.
- the inverse operation of a group is a set-function $X \rightarrow X$. This translates to our framework as $F(X) = X$ and so $\alpha : X \rightarrow X$ in this case.
- the group operation is a binary set-function $(\cdot) : X \times X \rightarrow X$. We have $F(X) = X \times X$, so the morphism $\alpha : X \times X \rightarrow X$ has the correct type.

But each of these are still independent pieces of a group definition. We must combine them somehow. A single morphism $\alpha : F(X) \rightarrow X$ must represent all actions you can possibly take on a group. One answer would be to take $F(X) = 1 \uplus X \uplus X \times X$, where \uplus is the disjoint set-union operation. This is usually written instead as $F(X) = 1 + X + X \times X$, as we will soon generalize disjoint unions in **Set** to coproducts in an arbitrary category.

So our combined morphism $\alpha : F(X) \rightarrow X$ essentially gives us a choice of which operation we would like to take, as well as how each is implemented. It should be clear at this point that each functor F can possibly have multiple $\alpha : F(X) \rightarrow X$. This corresponds to the notion that there are multiple groups, multiple rings, multiple Boolean algebras, etc.

Of course, since F is a functor, it must have an action on morphisms as well. In the case of $F(X) = A \times X + B$, a functor application to a morphism $F\left(X \xrightarrow{f} Y\right)$ would have type $A \times X + B \xrightarrow{F(f)} A \times Y + B$. It may be useful to think of $F(f)$ as being a coproduct of morphisms $A \times X \xrightarrow{f} A \times Y$ and $B \xrightarrow{1} B$. This is often true for the cases we will be studying, as many of the functors for the common algebras do exhibit this property. However, it may not be true in general, and the categorical

framework for algebra allows for functors that mix between the different components. Still, we will use the notation $F(f) = f_1 + f_2$ in the remainder of the notes.

Definition 1.1 (*F*-algebra)

Let \mathcal{C} be an ambient category. An *F*-algebra is a pair $(X, \alpha : F(X) \rightarrow X)$, where $X \in \text{Ob}(\mathcal{C})$, F is an endofunctor $\mathcal{C} \rightarrow \mathcal{C}$, and $\alpha : F(X) \rightarrow X$.

This doesn't cover all of the group properties yet. We are still missing associativity and the inverse laws. Unfortunately, the functor F primarily deals with the types of operations in an algebra, also known as the "signature", so the extra laws must be separately required.¹ For example, we will define groups as follows:

A group is an *F*-algebra defined by the signature

$$F(X) = 1 + X + X \times X \tag{1}$$

where the components are named $[e, \text{inv}, (\cdot)]$, such that the following diagrams commute:

$$\begin{array}{ccc} X \times (X \times X) & \xlongequal{\text{assoc}} & (X \times X) \times X \\ \downarrow (\text{id}, \cdot) & & \downarrow (\cdot, \text{id}) \\ X \times X & \xrightarrow{(\cdot)} X \xleftarrow{(\cdot)} & X \times X \end{array}$$

$$\begin{array}{ccccc} X & \xrightarrow{(\text{id}, \text{inv}_R)} & X \times X & \xleftarrow{(\text{inv}_L, \text{id})} & X \\ \downarrow ! & & \downarrow (\cdot) & & \downarrow ! \\ 1 & \xrightarrow{\text{id}} & X & \xleftarrow{\text{id}} & 1 \end{array}$$

$$\begin{array}{ccccc} 1 \times X & \xrightarrow{(e, \text{id})} & X \times X & \xleftarrow{(\text{id}, e)} & X \times 1 \\ & \searrow & \downarrow (\cdot) & \swarrow & \\ & & X & & \end{array}$$

It turns out there are some interesting *F*-algebras when we consider the category of *F*-algebras themselves. But before we get there, we must first define what it means to have a morphism of *F*-algebras. As is standard in category theory, an *F*-algebra is simply an underlying carrier plus some extra data. We will thus define morphisms as a function between the carrier along with proof of coherence between the data:

¹There is a different construction which generates algebras from a monad, that can encapsulate having these extra axioms, but we won't really explore that here.

Definition 1.2 (F -algebra morphism)

Let (X, α_X) and (Y, α_Y) be two F -algebras in a category \mathcal{C} . An F -algebra morphism is then a morphism $f : X \rightarrow Y$ such that the following diagram commutes:

$$\begin{array}{ccc} F(X) & \xrightarrow{\alpha_X} & X \\ F(f) \downarrow & & \downarrow f \\ F(Y) & \xrightarrow{\alpha_Y} & Y \end{array}$$

Exercise 1.1. Verify that F -algebras form a category.

Let us now interpret this commutative square in the context of the F -algebras generated by each of our group operations above.

- In the case of identity, the component of $F(X)$ we are focusing on is $F(X) = 1$. So both $F(X)$ and $F(Y)$ evaluate to 1, so $F(X) \rightarrow F(Y)$ can only be the unique identity morphism in 1. Thus, the left side collapses and we have a triangle:

$$\begin{array}{ccc} 1 & \xrightarrow{\alpha_X} & X \\ & \searrow \alpha_Y & \downarrow f \\ & & Y \end{array}$$

Translating this to group theory, this means our choice of α_Y , the function that picks out the group identity of Y , must agree with the result of mapping f to the group identity of X . In other words, group homomorphisms preserve identities.

- In the case of inverse, $F(X) = X$, so we have the following diagram:

$$\begin{array}{ccc} X & \xrightarrow{\text{inv}_X} & X \\ f \downarrow & & \downarrow f \\ Y & \xrightarrow{\text{inv}_Y} & Y \end{array}$$

This translates to $\text{inv}_Y \circ f = f \circ \text{inv}_X$, which written algebraically says that for any $x \in X$, it must hold that $f(x)^{-1} = f(x^{-1})$. In other words, group homomorphisms preserve inverses.

- In the case of the binary operation (\cdot) , we have $F(X) = X \times X$, so we have the following diagram (recall the definition of the action of F on morphisms):

$$\begin{array}{ccc}
 X \times X & \xrightarrow{(\cdot)_X} & X \\
 (f, f) \downarrow & & \downarrow f \\
 Y \times Y & \xrightarrow{(\cdot)_Y} & Y
 \end{array}$$

Translating this to algebra, we have $f(x \cdot_X y) = f(x) \cdot_Y f(y)$, which is the main group homomorphism property.

As you can see, the group homomorphism properties fell purely out of just instantiating our F -algebra morphism law with our specific group properties.

1.1 Initial F -algebras and induction

We said before that $F(X)$ builds up the input data to be used by α . But that's a rather abstract description. To see an example of how this is used, consider a functor $F(X) = 1 + X$. We can define an F -algebra by providing an $\alpha : F(X) \rightarrow X$. Expanding the type gives us $\alpha : (1 + X) \rightarrow X$, which we can define by cases:

- $\alpha_1 : 1 \rightarrow X$
- $\alpha_X : X \rightarrow X$

Suppose we are working with an (X, α) that is initial in the category of F -algebras for $F(X) = 1 + X$. The defining property of an initial object is that there exists a [morphism](#) out of this F -algebra to any other F -algebra, such that the square commutes. This morphism is commonly known as the **catamorphism**.

For example, consider any other F -algebra $\beta : (1 + Y) \rightarrow Y$, with corresponding $[\beta_1, \beta_Y]$. Then, there exists a *unique* $f : X \rightarrow Y$ such that this square commutes:

$$\begin{array}{ccc}
 1 + X & \xrightarrow{\alpha} & X \\
 F(f) \downarrow & & \downarrow f \\
 1 + Y & \xrightarrow{\beta} & Y
 \end{array}$$

where $F(f)$ is defined on cases:

- for the left case of the sum, there is a morphism $1_X \rightarrow 1_Y$
- for the right case of the sum, there is a morphism $X \rightarrow Y$

The square forces f to respect the F -algebra structure.

There is a sense where the initial (F, α) shown above represents the natural numbers \mathbb{N} , and that the unique morphism f is the natural number recursor. Recall that the natural number recursor looks something like this:

$$\text{natrec} : (y_0 : Y) \rightarrow (y_s : Y \rightarrow Y) \rightarrow (\mathbb{N} \rightarrow Y) \tag{2}$$

If we change y_0 to be of type $1_Y \rightarrow Y$ (this is equivalent, it's just a function that ignores the input), it should be clear that the inputs y_0 and y_s represents the two cases of β . So really, we can reframe this as: given $\beta : 1 + Y \rightarrow Y$, which is the definition of the F -algebra you are trying to produce, you get the unique morphism $\mathbb{N} \rightarrow Y$ back, since it's forced by the square.

Exercise 1.1.1. *Prove that `natrec` is the unique morphism that makes the above square commute.*

We defined `natrec` for naturals here, but the idea that the unique catamorphism is the recursor holds for arbitrary F -algebras.

Definition 1.1.1 (Recursor)

Given some functor F , and an initial F -algebra (X, α) , the recursor is the unique F -algebra morphism to any other F -algebra (Y, β) .

This generalizes to other algebra-shaped structures as well. In fact, F -algebras generally build inductive data types. For example, lists are $F(X) = 1 + A \times X$ for A -typed lists. The initial object in this category is then the pair (X, α) defined by:

- X is the set of all finite lists of type A .
- $\alpha : 1 + A \times X \rightarrow X$ is then defined by cases on the input:
 - In the case of `inl(★)`, just output the empty list.
 - In the case of `inr(a, x)`, output the list whose head is a and tail is x .

Exercise 1.1.2. *Define some other common algebraic structures using functors.*

So without initiality, what we have is only a formula for building some generic objects with this shape. But there is nothing really forcing the objects to take a coproduct shape, or to behave well. In fact, the initiality enforces the self-similarity of the structure map α , which we will see in the next section.

1.2 Lambek's Theorem

There seems to be some parallel between the structure of $F(X)$ and X itself. For example, the structure map $\alpha : F(X) \rightarrow X$ seemed to be perfectly split into the same cases as X itself. It turns out that this structure map itself is an isomorphism.

Theorem 1.2.1 (Lambek's theorem)

Given some functor F and some F -algebra (X, α) . If (X, α) is initial, then α is an isomorphism.

Proof. We begin with our structure map

$$F(X) \xrightarrow{\alpha} X$$

which is initial in the category of F -algebras. Let us “use” this initiality by coming up with another F -algebra to map this to. We will use $!$ to denote the catamorphism – the unique morphism out of the initial F -algebra.

$$\begin{array}{ccc}
 F(X) & \xrightarrow{\alpha} & X \\
 F(!) \downarrow & & \downarrow ! \\
 ? & \xrightarrow{\quad} & ? \\
 & & ?
 \end{array}$$

We will choose the particular F -algebra defined by applying F again:

$$\begin{array}{ccc}
 F(X) & \xrightarrow{\alpha} & X \\
 F(!) \downarrow & & \downarrow ! \\
 F(F(X)) & \xrightarrow{F(\alpha)} & F(X)
 \end{array}$$

If we tack on a tautological square to the bottom, we get:

$$\begin{array}{ccc}
 F(X) & \xrightarrow{\alpha} & X \\
 F(!) \downarrow & & \downarrow ! \\
 F(F(X)) & \xrightarrow{F(\alpha)} & F(X) \\
 F(\alpha) \downarrow & & \downarrow \alpha \\
 F(X) & \xrightarrow{\alpha} & X
 \end{array}$$

We can then “squish” the square using composition, and also using the functoriality of F to simplify $F(\alpha) \circ F(!)$ into $F(\alpha \circ !)$:

$$\begin{array}{ccc}
 F(X) & \xrightarrow{\alpha} & X \\
 F(\alpha \circ !) \downarrow & & \downarrow \alpha \circ ! \\
 F(X) & \xrightarrow{\alpha} & X
 \end{array}$$

This is exactly the [F-algebra morphism diagram](#) showing that $\alpha \circ !$ is a morphism from the F -algebra (X, α) to itself. Since the morphism out of an initial object is unique, we know that it must coincide with the identity morphism, so $\alpha \circ ! = \text{id}_X$.

This is the first half of our proof. It remains to be shown that $! \circ \alpha = \text{id}_{F(X)}$. Using the result that $\alpha \circ ! = \text{id}_X$, apply F to both sides, getting $F(\alpha \circ !) = F(\text{id}_X)$, which by functoriality yields $F(\alpha) \circ F(!) = \text{id}_{F(X)}$. But looking at this square from above:

$$\begin{array}{ccc}
 F(X) & \xrightarrow{\alpha} & X \\
 F(!) \downarrow & & \downarrow ! \\
 F(F(X)) & \xrightarrow{F(\alpha)} & F(X)
 \end{array}$$

We know that $F(\alpha) \circ F(!) = ! \circ \alpha$. Therefore, $! \circ \alpha = \text{id}_{F(X)}$, and α is an isomorphism with inverse $!$. \square

Since α is essentially defining the “constructors” of the inductive type defined by $F(X)$, one consequence of Lambek’s theorem is that all constructors are invertible. For example, we can derive an inverse suc function that takes any $\text{suc}(n)$ and produces n .

Exercise 1.2.1. *Prove that there is no non-empty initial algebra for the functor $F(X) = A \times X$ in \mathbf{Set} , for any arbitrary object A . What does this mean in English?*

Applying Lambek’s theorem. Consider the list example again, $F(X) = 1 + A \times X$. Let us derive the initial F -algebra using Lambek’s theorem. First, Lambek’s theorem tells us that for initial algebras, $X \cong F(X)$ necessarily. Thus, $X \cong 1 + A \times X$. Since this is a fixpoint, we can start unfolding X and see what happens:

$$\begin{aligned}
 X &\cong 1 + A \times X \\
 &\cong 1 + A \times (1 + A \times X) \\
 &\cong 1 + A \times (1 + A \times (1 + A \times X)) \\
 &\cong \dots \\
 &\cong 1 + (A \times 1) + (A \times A \times 1) + (A \times A \times A \times 1) + \dots \\
 &\cong 1 + A + A^2 + A^3 + \dots
 \end{aligned} \tag{3}$$

This is a union of n -products of A for all n , which is exactly the set of all n -length sequences in A .

2 Coalgebra

As we know, we can get dual concepts in category theory for free by reversing the arrows. By dualizing algebras, we get coalgebras. Let’s start with the definition and then try to extract some intuitions from it.

Definition 2.1 (F -coalgebra)

Let \mathcal{C} be an ambient category. An F -coalgebra is a pair (X, ν) where $X \in \text{Ob}(\mathcal{C})$, F is an endofunctor $\mathcal{C} \rightarrow \mathcal{C}$ and $\nu : X \rightarrow F(X)$.

And as usual, we must also define morphisms:

Definition 2.2 (F -coalgebra morphism)

Let (X, ν_X) and (Y, ν_Y) be two F -coalgebras in a category \mathcal{C} . An F -coalgebra morphism is then a morphism $f : X \rightarrow Y$ such that the following diagram commutes:

$$\begin{array}{ccc}
 X & \xrightarrow{\nu_X} & F(X) \\
 f \downarrow & & \downarrow F(f) \\
 Y & \xrightarrow{\nu_Y} & F(Y)
 \end{array}$$

Similar to how in algebras, we had the intuition of building up some kind of “input data” to our α morphism which then represented the constructors for our X , we would like to think of coalgebras as functions out of X , giving us some *one-step* observation data that may involve X .

Let’s work through an example that might solidify this intuition. Once again, we are working temporarily in the category of sets, so let X be a set. Consider the endofunctor $F(X) = 1 + A \times X$. In algebras, this corresponds to the signature for list-y structures.

We will see now what the dual to list-y things are. To complete the definition of an F -coalgebra, we need a structure map of type $\nu : X \rightarrow F(X)$, which when unfolded, is $\nu : X \rightarrow 1 + A \times X$. Whereas in algebras, the structure map served a purpose of consuming the structure into a single X , here we have ν expanding the X into its observable data.

Our intuition tells us that X is supposed to be list-y. So let’s say our ν does case analysis to see if our list is cons or nil. In the case of cons, we return $A \times X$, which is the head and tail separately. In the case of nil, we return 1. So whereas our F -algebra from before was building *up* a list, now our F -coalgebra is essentially mapping *out* of a list.

Compared to the algebraic version, this one sort of axiomatizes observations out of the list rather than building up the exact list structure. We assume the list exists as a mystery object, and are able to decompose it into an element and “the rest of the stream,” which continues being the mystery object.

We also have a version of Lambek’s theorem for coalgebra:

Theorem 2.3 (Lambek’s theorem for coalgebra)

Given some functor F and some F -coalgebra (X, ν) . If (X, ν) is terminal, then ν is an isomorphism.

Exercise 2.1. Prove [Theorem 2.3](#).

Of course, with coalgebraic lists, we could also imagine a situation where we have a definition that continuously returns an element of the $A \times X$ side, and never returns the 1. In this case, the list would never end. We cannot formulate this type using algebraically defined data ([Exercise 1.2.1](#)), but with coalgebraically defined codata, this is completely in bounds. In fact, we may even specify that the stream *cannot* ever end. In this case, we only need to use signature functor $F(X) = A \times X$, opting not to include the extra 1. This is useful for reasoning about applications such as servers which are designed to run and continue responding to input forever. In the next section we will discuss the terminal F -coalgebra and coinduction.

2.1 Coinduction and bisimulation

Recall that a terminal object has morphisms into it from every other object. These morphisms are typically called **anamorphisms**. To see an example, let's work in **Set** again, with the stream functor, where $F(X) = A \times X$, for some type A . This represents streams of elements with type A .

We can use [Theorem 2.3](#) to derive what the terminal coalgebra should look like. Starting with $X \cong A \times X$, we can unfold several times to get $X \cong A \times A \times A \times \dots$. This is an unending product of A by itself, which is the same as the function space $\mathbb{N} \rightarrow A$. We can also write this as $A^{\mathbb{N}}$, using exponential notation for functions.

Next is to define the structure map, $\nu_X : X \rightarrow A \times X$. Of course, we don't really have a choice in this matter either – it must follow from the terminality conditions. An obvious choice would be to just pick the head and tail and return that as a pair.

Definition 2.1.1

The structure map for the terminal F -coalgebra defined by $F(X) = A \times X$ is

$$\begin{aligned} \nu_X : X &\rightarrow A \times X \\ \nu_X((a_0, a_1, a_2, \dots)) &= a_0, (a_1, a_2, \dots) \end{aligned} \tag{4}$$

But does this choice work? We must check that all the rules are satisfied.

Lemma 2.1.2

For any F -coalgebra (Y, ν_Y) , there exists a morphism $f : (Y, \nu_Y) \rightarrow (X, \nu_X)$.

Proof. First, we must show that for any (Y, ν_Y) , that an $f : Y \rightarrow X$ even exists. We will do so by generating a sequence of a 's from ν_Y . First, begin with $y \in Y$, and run $\nu_Y(y)$ to get $y_1 : A \times Y$. We'll take the first element $\text{head}(y_1) : A$ as the first observation, and keep the rest which is $\text{tail}(y_1) : Y$. Then, we can repeat the process, running $\nu_Y(y_1)$ to get $y_2 : A \times Y$. By continuing to repeat this process, we'll just end up with an infinite series of a 's, which we can write as a_0, a_1, a_2, \dots . We'll refer to this as the *behavior* of Y . Notice that we didn't need to depend on what the structure of Y was.

Now we need to verify that our choice of f is well-behaved, with regards to all of the conditions we defined earlier. Let us revisit the commuting square from [Definition 2.2](#):

$$\begin{array}{ccc} Y & \xrightarrow{\nu_Y} & A \times Y \\ f \downarrow & & \downarrow F(f) \\ A^{\mathbb{N}} & \xrightarrow{\nu_X} & A \times A^{\mathbb{N}} \end{array}$$

Translating this to an equation, we get:

$$\nu_X \circ f = F(f) \circ \nu_Y \tag{5}$$

First, note that it is necessary that $F(f) = \text{id} \times f$. This is because of the functoriality of F , similar to how our F 's earlier had to be a coproduct. So really, we have:

$$\nu_X \circ f = (\text{id} \times f) \circ \nu_Y \quad (6)$$

This requires us to show that ν_Y indeed produces the head element of the list that we are manipulating in ν_X . More literally, if we have $y \in Y$, then $f(\text{tail}(\nu_Y(y)))$ must produce the same thing as just taking the tail end of $f(y)$. This is just true by definition. To spell it out with elements, we have:

$$\begin{aligned} \nu_X(f(y)) &= (\text{id} \times f)(\nu_Y(y)) \\ \nu_X((a_0, a_1, a_2, \dots)) &= \\ a_0, (a_1, a_2, \dots) &= \\ &= (\text{id} \times f)(\nu_Y(y)) \\ &= (\text{id} \times f)(a_0, y_1) \\ &= a_0, f(y_1) \end{aligned} \quad (7)$$

As noted above, y_1 is just the next sequence after y , so it will produce a_1, a_2, \dots by f . So this commuting square checks out. \square

Lemma 2.1.3

The morphism f defined in [Lemma 2.1.2](#) is unique.

For this, we'll need to use a different kind of reasoning than before. We will have to use **coinduction**, which is dual to induction and is perfect for working with codata such as streams.

Proof. Suppose we have $g : Y \rightarrow A^{\mathbb{N}}$ that respects the square. We would like to show that $f = g$. Extensionally, this means f and g produce the same output on all inputs, as functions. Now, g is an unknown function, but it has the shape $Y \rightarrow A^{\mathbb{N}}$, so we know it produces an infinite stream of a 's. Furthermore, we know that:

$$\begin{array}{ccc} Y & \xrightarrow{\nu_Y} & A \times Y \\ g \downarrow & & \downarrow \text{id} \times g \\ A^{\mathbb{N}} & \xrightarrow{\nu_X} & A \times A^{\mathbb{N}} \end{array}$$

which says that $\nu_X \circ g = (\text{id} \times g) \circ \nu_Y$. This means at least the first element of the sequence produced by g must be picked out by ν_Y . Otherwise, $\text{id}(\text{head}(\nu_Y(y)))$ would not produce the same result as $\text{head}(\nu_X(g(y)))$. Thus, the first element of the sequence must be $\text{head}(\nu_Y(y))$.

Then, regarding the tail end of the sequence, we would like to argue coinductively that it must follow the same shape. Specifically, let $\nu_Y(y) = (a_0, y')$. Following the top-right path, we can see that the tail part of the result is $g(y')$. Since g is unknown, we can't really say what this is yet, but we can look at the left-bottom path, which tells us that the tail part of the result must be $\text{tail}(g(y))$, since as a

reminder, our ν_X specifically splits a sequence $A^{\mathbb{N}}$ into its head and tail parts. This tells us that $g(y') = \text{tail}(g(y))$, which is exactly what we expect.

So putting this together with f , we know that both $\text{head}(f(y))$ and $\text{head}(g(y))$ result in $\text{fst}(\nu_Y(y))$. So $f(y)$ and $g(y)$ agree on that part. For the tail segment, we know that $\text{tail}(g(y)) = g(y')$. It turns out this is also true of f , due to how we defined it: $\text{tail}(f(y)) = f(y')$. Thus, to determine if $\text{tail}(f(y)) = \text{tail}(g(y))$, it suffices to show that $f(y') = g(y')$. But since y' is exactly self-similar to y , by coinduction, we know that $f(y') = g(y')$. \square

In the above proof, we used function extensionality to show that f and g are equal. But for showing two coalgebra are “the same,” it is often more convenient to use a weaker notion of equivalence. A classic example is of automata, which we will see more about in a later section. It is completely sensical to have two automata that capture the same language (yielding the same transitions), yet have different non-reconcilable states. For this, we will introduce bisimulation.

Definition 2.1.4 (Bisimulation of streams)

Let (X, ν_X) and (Y, ν_Y) be F -coalgebras, where $F(X) = A \times X$. Let R be a relation between X and Y – specifically, $R \subset X \times Y$. R is a **bisimulation** iff for any $(x, y) \in R$:

$$\text{head}(x) = \text{head}(y) \wedge (\text{tail}(x), \text{tail}(y)) \in R \quad (8)$$

Then, it is said that X and Y are **bisimilar**. This sort of looks like a more generalized version of the reasoning we did as a part of the uniqueness proof, but with equality replaced with the relation R . But the reasoning is still stream-specific, and relatively ad-hoc. We’d like to not have to make arbitrary bisimulation conditions for every type we’d like to analyze.

It turns out if we just take everything that has specifically to do with streams and abstract it, we can arrive at a more general notion of bisimulation.

Definition 2.1.5 (Aczel-Mendler bisimulation)

Let \mathcal{C} be a category with products, and $F : \mathcal{C} \rightarrow \mathcal{C}$ be an endofunctor. Let (X, ν_X) and (Y, ν_Y) be F -coalgebras. Let R be a subobject of the product object $X \times Y$ with projection morphisms π_X and π_Y such that there exists a morphism $\nu_R : R \rightarrow F(R)$ such that the following diagram commutes:

$$\begin{array}{ccccc} X & \xleftarrow{\pi_1} & R & \xrightarrow{\pi_2} & Y \\ \nu_X \downarrow & & \downarrow \nu_R & & \downarrow \nu_Y \\ F(X) & \xleftarrow{F(\pi_1)} & F(R) & \xrightarrow{F(\pi_2)} & F(Y) \end{array}$$

Let’s quickly make sense of this in the context of our previous [Definition 2.1.4](#) again. We have some relation R which is a subobject of the product. This corresponds exactly with $R \subset X \times Y$. The rough interpretation of this is that at the current state, X and Y are related through R .

Then, we have the object $F(R)$, which is a lift of the product through F . So for streams, since $F(X) = A \times X$, we have $F(R) \subset (A \times X) \times (A \times Y)$. Notice how this part is completely determined by the specific F 's interaction with the relation. The ν_R lifting operation simply calls head and tail component-wise on both the left and right sides. Then, the $F(\pi_1)$ projection operator would grab the components separately and turn them into a product, and $F(\pi_2)$ does the corresponding thing for $F(Y)$.

The commuting squares, like earlier, essentially force us to take this component-wise motion.

Exercise 2.1.1. *Does the category of F -algebras have an initial object for any F ? Does the category of F -coalgebras have a terminal object for any F ?*

3 Automata

We have been working a lot with streams so far, so from here on we will be introducing some richer coalgebras that have applications in other fields. Automata are a classic example of coalgebras in practice. At a high level, an automaton is a state machine that has states and transitions between them. It has a notion of “accept” states which allow the machine to respond with a success response. This is often used in recognizing languages, such as in parsing for compiler development, but also for network protocols and hardware state machines.


We’ll look at how automata can be abstracted using algebras and coalgebras, and then work through a generalization of Brzozowski’s minimization algorithm using our definition, in order to see how to use the duality of algebras and coalgebras in order to achieve minimization. Let’s start with a formal definition of automata.

Definition 3.1 (Deterministic automaton)

For some alphabet Σ , a **deterministic automaton (DA)** is a 4-tuple (Q, δ, i, F) :

- Q , a set of states
- $\delta : Q \times \Sigma \rightarrow Q$, a transition map
- $i \in Q$, an initial state
- $F \subseteq Q$, final states

Visually, we can represent an automaton like this:



This visualization represents the formal automaton defined as $(\{q_0, q_1, q_2\}, \{a, b, c\}, \delta, q_0, \{q_2\})$ where δ would map:

- $(a, q_0) \mapsto q_2$

- $(b, q_0) \mapsto q_1$
- $(b, q_1) \mapsto q_1$
- $(c, q_1) \mapsto q_2$
- all other combinations of states would resolve to an implicit “error” state (which would also be in Q), which only has transitions back to itself. It does not appear in F , the set of final states. For all intents and purposes in these notes, we can safely just ignore the fact that it exists, although we should formally make a mental note of the erroring behavior, since δ is not a partial function.

The double-circle means that q_2 is a final state, in other words, that $q_2 \in F$. This automaton will recognize a language consisting of either the exact string a or any non-zero number of b s followed by a c . This is typically written using the terse regular expression syntax as $a | (b^+)c$.

Formally, a language \mathcal{L} is just a series of strings consisting of symbols in the alphabet Σ . An automaton **recognizes** a language if for every string s in the language, running s through the automaton would result in a final state.

There is also a variant which allows non-deterministic transitions – for example, it would be allowed to have multiple arrows from the same state with the same transition. This machine is known as a *non-deterministic* automaton.

Definition 3.2 (Non-deterministic automaton)


For some alphabet Σ , a **non-deterministic automaton** is a 4-tuple (Q, δ, i, F) :

- Q , a set of states
- $\delta : Q \times \Sigma \rightarrow \mathcal{P}(Q)$, a transition map
- $i \in Q$, an initial state
- $F \subseteq Q$, final states

The δ is the only difference from the DA definition. Instead of only transitioning to one state, it could possibly transition to any number of states. A non-deterministic automata would accept a string from a language if taking any path using the symbols of the string would result in acceptance.

Traditionally, non-deterministic automata would be converted deterministic using an algorithm such as the powerset construction, in which you would consider the set of all states that could possibly result following a particular transition. There are also some methods of evaluating non-deterministic automata directly without first doing the conversion, as the conversion may produce an explosion ($O(2^n)$) of states in the resulting DA.

For this section, we will consider a particular problem concerning automata: there is a sense in which an NFA might not be minimal. For example, consider this automaton:



You could easily imagine collapsing all of these down to a single node. Due to this, there are several minimization algorithms. A well-known algorithm, due to Brzozowski, involves reversing the NFA, determinizing it using the powerset construction, and then reversing it again. In this next section, we will investigate a category-theoretic approach by (Bonchi et al. 2014) to show why this works, specifically for the case of deterministic automata.

3.1 The category of automata

In order to categorify automata, we must first identify a category. We want the objects to be automata, so we must define some morphisms.

Definition 3.1.1 (Automaton morphism)

For some alphabet Σ , a morphism between two automata $(Q_1, \delta_1, i_1, F_1)$ and $(Q_2, \delta_2, i_2, F_2)$ is defined as a function $f : Q_1 \rightarrow Q_2$ such that the following hold:

- f commutes with the transition map, or in other words, for any $s \in \Sigma$, this square commutes:

$$\begin{array}{ccc} Q_1 & \xrightarrow{f} & Q_2 \\ \delta_1(-, s) \downarrow & & \downarrow \delta_2(-, s) \\ Q_1 & \xrightarrow{f} & Q_2 \end{array}$$

- f preserves the initial state, or in other words $f(i_1) = i_2$
- f preserves the final states, or in other words $f(q) \in F_2$ for all $q \in F_1$

Thus, we can form a category of DAs:

Definition 3.1.2 (Category of deterministic automata, $\mathbf{DA}_{\mathcal{L}}$)

For some alphabet Σ and language over that alphabet \mathcal{L} , the category $\mathbf{DA}_{\mathcal{L}}$ is defined as a category where:

- objects are DAs that recognize the language \mathcal{L}
- morphisms are defined as in [Definition 3.1.1](#)

It's important that the category itself is parameterized by a language \mathcal{L} , since morphisms need to preserve the language. If this was not the case, then we could not have initial or final objects at all – it would essentially require that all automata at all recognize the same language, which is not the case. Such a category would still exist, and in fact $\mathbf{DA}_{\mathcal{L}}$ is a subcategory of that category, but it is not useful to us.

Remark

We can characterize a DA categorically as *both* an algebra and as a coalgebra. As an algebra, it uses the signature functor

$$F(Q) = 1 + \Sigma \times Q \tag{9}$$

In order to represent an automata as an algebra, we would need a structure map of type $\alpha : 1 + \Sigma \times Q \rightarrow Q$. The $1 \rightarrow Q$ part represents the initial state, since it takes no input, and $\Sigma \times Q \rightarrow Q$ is the type of a transition map – it takes a transition symbol, a state, and outputs the next state.

As a coalgebra, it can be characterized with the signature functor:

$$F(Q) = 2 \times Q^\Sigma \quad (10)$$

For the coalgebra case, we need a structure map $\nu : Q \rightarrow 2 \times Q^\Sigma$. For any input state $q \in Q$, the Q^Σ part of the output represents all possible states that the input state could transition to, and 2 refers to a tag on each state that is assigned the value 1 if that state is a final state, or the value 0 if it isn't.

3.2 Minimality of automata

Next, let us consider the initial and terminal objects of $\text{DA}_{\mathcal{L}}$. We'd like to determine some automaton that recognizes \mathcal{L} and has a morphism to every other automaton. A pretty trivial “maximal” object that lets us do this is to add every single possible word as a state.

Formally, let's define this automaton:

Definition 3.2.1 (Initial object of $\text{DA}_{\mathcal{L}}$)

Define the initial automaton In_{it} as:

- $Q_{\text{In}_{\text{it}}}$ is defined as Σ^* , the set of all strings over the alphabet Σ
- $\delta_{\text{In}_{\text{it}}} : Q \times \Sigma \rightarrow Q$ simply concatenates the currently running state's string with the transition symbol
- $i_{\text{In}_{\text{it}}}$ is the state corresponding to the empty string
- $F_{\text{In}_{\text{it}}}$ are the states corresponding to all strings in \mathcal{L}

Of course, we must show that it is in fact initial.

Theorem 3.2.2

The automaton described above is initial in $\text{DA}_{\mathcal{L}}$.

Proof. Consider some other automaton (Q', δ', i', F') . We must define a morphism from In_{it} to this automaton.

First, define $f : Q_{\text{In}_{\text{it}}} \rightarrow Q'$. For every state $q \in Q_{\text{In}_{\text{it}}}$, there is a corresponding string $s \in \Sigma^*$. We will define f to map s to a state in Q' by running Q' – in other words, starting with i' and transforming it using δ' on each symbol of s , and then returning the final state that it landed in.

We must show that f commutes with the transition map, or in other words, for any symbol $t \in \Sigma$:

$$f \circ \delta_{\text{In}_{\text{it}}}(-, t) = \delta'(-, t) \circ f \quad (11)$$

Let $q \in Q_{\text{In}_{\text{it}}}$ be an arbitrary state in In_{it} , and $s \in \Sigma^*$ be the corresponding string. Then, we are trying to show that $f(\delta_{\text{In}_{\text{it}}}(q, t)) = \delta'(f(q), t)$. Since $\delta_{\text{In}_{\text{it}}}$ is defined by concatenation, we have $f(\delta_{\text{In}_{\text{it}}}(q, t)) = f(s + t)$ on the left side. Since f simply runs δ' iteratively, this is actually true *by the definition* of f .

It is trivial to show that f also preserves the initial and final states.

In order to show f is the unique such morphism, assume $g : Q_{\text{init}} \rightarrow Q'$. We need to show for any state q , that $f(q) = g(q)$. Since the states map one-to-one with all strings Σ^* , perform induction on strings. For the empty string, the corresponding state is the initial state, which is preserved by the requirement of the automaton morphism. Thus, $f(i_{\text{init}}) = g(i_{\text{init}}) = i'$.

For the inductive case, assume $f(q) = g(q)$. Then, for any symbol $t \in \Sigma$, identify the state corresponding to concatenating $s + t$ (once again letting s be the string corresponding to q). We must show $f(s + t) = g(s + t)$. By the commutative square, we can rewrite this property to $\delta'(f(q), t) = \delta'(g(q), t)$. Since $f(q) = g(q)$ by our inductive hypothesis, the two sides are identical. \square

Alternate proof. Note that automata behave algebraically. Use [Theorem 1.2.1](#) to deduce that the structure map $1 + \Sigma \times Q \rightarrow Q$ must be an isomorphism, so:

$$Q \cong 1 + \Sigma \times Q \quad (12)$$

Expanding this, we get: $Q \cong 1 + \Sigma \times (1 + \Sigma \times (1 + \Sigma \times \dots))$, which simplifies to $1 + \Sigma + \Sigma^2 + \Sigma^3 \dots$. This is exactly Σ^* – all strings of all lengths, sourced from the alphabet.² \square

The rough design of this morphism is that for some arbitrary automaton Q , we're running every possible string in Σ^* to see which states are able to be reached via Q , and capturing that in a single function f .

We're interested specifically in automata where every state is reachable. This means there's no redundancies in the states – if you had some redundant state that isn't final, then only running δ until completion would completely miss that state, excluding it from the image of the morphism out of the initial automaton. We capture this with the following definition:

Definition 3.2.3 (Reachable automaton)

An automaton Q is **reachable** if the unique morphism out of the initial automaton $\text{init} \rightarrow Q$ is an epimorphism.

In sets, epimorphisms (set functions) capture the notion of surjectivity, which means every element of the codomain has a pre-image. Since automata are essentially sets with more structure, a morphism being epic captures the same notion, which in automata corresponds to every state being reached by a morphism. If this property is true of the morphism out of the initial automaton (which remember, runs every possible string through the automaton), we know every state is being used, which captures our desired property of reachability.

There is a similar construction on the coalgebraic side. We can imagine a terminal automaton, which has morphisms going into it from any other automaton. For some arbitrary automaton Q , we'd imagine that the result of a morphism should be some *observation* of the states in Q . So the terminal automaton should be a collective observation about all possible automata in $\text{DA}_{\mathcal{L}}$.

Definition 3.2.4 (Terminal object in $\text{DA}_{\mathcal{L}}$)

The terminal automaton Term is defined as:

²In fact, this approach can be used to derive the initial automaton. I didn't realize this until after I wrote the previous proof... However, I left the original proof in, because it actually defines the function needed for the morphism.

- Q_{Term} is the set of all possible subsets of Σ^* . This represents all languages that could possibly be recognized for this alphabet.
- $\delta_{\text{Term}} : Q \times \Sigma \rightarrow Q$ maps every $(q, s) \in Q \times \Sigma$ to the state representing the remainder of the language after consuming the prefix s . This is also known as the *Brzowski derivative*.
- i_{Term} is the language \mathcal{L} .
- F_{Term} is the set of all languages that contain the empty string ϵ .

Staring at this definition, it may kind of make sense why this could capture all observations, but let's make this concrete by actually proving terminality.

Theorem 3.2.5

The automaton described above is final in $\text{DA}_{\mathcal{L}}$.

Proof. This proof largely follows in the same vein as the initial object proof. As such, we will skip some parts of it for brevity. Just trust me 😊.

Consider some other automaton (Q', δ', i', F') . We'll now define $f : Q' \rightarrow Q_{\text{Term}}$. For any state $q \in Q'$, we will consider the automaton that has the same states and transitions as Q' , but with initial state q . The language that this automaton recognizes will be a set of strings, which is a subset of 2^{Σ^*} . This language is the output of f .

The exercise of verifying that f preserves the initial and final states as well as the transition map is left to the reader. □

Exercise 3.2.1. Verify that f satisfies the requirements for being a $\text{DA}_{\mathcal{L}}$ -morphism.

Alternative proof. Same as before, notice that $\text{DA}_{\mathcal{L}}$ is coalgebraic and use [Theorem 2.3](#). The structure map $\nu : Q \rightarrow 2 \times Q^{\Sigma}$ must be an isomorphism, so:

$$Q \cong 2 \times Q^{\Sigma} \tag{13}$$

Expanding this, we get

$$\begin{aligned} Q &\cong 2 \times \left(2 \times \left(2 \times \left(\dots \right)^{\Sigma} \right)^{\Sigma} \right)^{\Sigma} \\ &= 2 \times 2^{\Sigma} \times \left(2 \times \left(\dots \right)^{\Sigma} \right)^{\Sigma^2} \\ &= 2 \times 2^{\Sigma} \times 2^{\Sigma^2} \times \left(\dots \right)^{\Sigma^3} \\ &= 2 \times 2^{\Sigma} \times 2^{\Sigma^2} \times 2^{\Sigma^3} \times \dots \\ &= 2^{1+\Sigma+\Sigma^2+\Sigma^3+\dots} \\ &= 2^{\Sigma^*} \end{aligned} \tag{14}$$

Thus, 2^{Σ^*} is the state space for the terminal automaton. □

The rough design for this morphism is that we're mapping the automaton to what language it recognizes, inside of Term . In that sense, this morphism would capture the behavior of all the strings that could possibly be recognized by the automaton.

This captures another important property, which is that every state is distinct. Without this property, you could have two states which recognize the same suffix language. That wouldn't be very minimal. For minimality, we care specifically that there are no duplicate states that can just be merged. This can be captured by the definition of observability, which is dual to [Definition 3.2.3](#):

Definition 3.2.6 (Observable automaton)

An automaton Q is **observable** if the unique morphism into the terminal automaton $Q \rightarrow \text{Term}$ is a monomorphism.

Again, to make a comparison to sets, a set function is monic if it is injective, or that distinct inputs map to distinct outputs. Thus, a morphism in $\text{DA}_{\mathcal{L}}$ is monic if it maps distinct input states to distinct output states. If our output state is the terminal automaton, then two states mapping into the same output state means that they recognized the same language, which would mean they were duplicated.

Putting these together, we can finally define what a minimal automaton is:

Definition 3.2.7 (Minimal automaton)

An automaton Q is **minimal** if it is both [reachable](#) and [observable](#).

The reachability property enforces that no states can be deleted, and the observability property enforces that no states can be merged, at least without changing the language being recognized. These are the only two ways to reduce an automaton without changing the language. This claim could be justified by the fact that the definition of minimization used above coincides with the minimization definition for a coalgebra defined through epi-mono factorization (Bezhanishvili et al. 2012).

On the computation side of things, reachability is pretty trivial to compute algorithmically – just traverse the state machine and grab all the states visited. On the other hand, observability requires quotienting by equivalence classes of a hard-to-compute property of basically *every* state. In practice, algorithms such as Hopcroft's algorithm optimize this by incrementally comparing states and quotienting as you go.

3.3 Reversing automata

The crux of Brzozowski's minimization algorithm involves reversing the automaton. The intuition here is, if reachability is more straightforward, we should aim to do it. Then, we can find some dualization property that allows us to somehow transfer reachability of an automaton into observability of the reverse language, and then use reachability *again* to get both desired properties on the same language.

The actual reversal process amounts to reversing the arrows of the transition map, as well as exchanging the initial and final states. Note that normally, when you do this to a deterministic automaton, it becomes a *non-deterministic* one. This is because there could be multiple final states, but only one initial state. When you reverse it, you have a number of initial states. The algorithm requires that you then make the non-deterministic automaton deterministic, through the powerset construction.

We'll explore a construction which combines the reversing and powersetting in one go. First off, consider what the shape of output we want is, since this will guide our type-checking as we hone in towards a solution. We are starting off with a category $\text{DA}_{\mathcal{L}}$ of deterministic automata, and we are trying to get some way of taking each object and obtaining the reverse automata.

Notice that the reverse automata will necessarily recognize the reverse language. So the output will be an automaton which recognizes $\text{rev}(\mathcal{L})$, in other words an object in $\text{DA}_{\text{rev}(\mathcal{L})}$! This makes it necessary for us to provide a *functor* with this signature:

$$\text{DA}_{\mathcal{L}} \rightarrow \text{DA}_{\text{rev}(\mathcal{L})} \quad (15)$$

As always, we need to find a functor, because the rules governing morphisms will restrict the way the objects can even be mapped to only well-behaved maps. Otherwise, we could produce some degenerate maps. This will also give us more benefits, as we will see later.

Time to define action on objects. For some arbitrary object $Q \in \text{DA}_{\mathcal{L}}$, we'd like to produce the reverse automaton. We can start with what we expect $\delta_{Q'}$, the δ of the transformed automaton to be. It should be some sort of reversing function, so we should be taking the pre-image of δ_Q :

$$\delta_{Q'} := (q, s) \mapsto \{q' \in Q \mid \delta(q', s) = q\} \quad (16)$$

But this doesn't type-check. The result is a set, since multiple q' s could map to the same q under the old δ . This means our state space fundamentally needs to be a powerset of Q . In that case, $\delta_{Q'}$'s type would need to be $\mathcal{P}(Q) \times \Sigma \rightarrow \mathcal{P}(Q)$. We can update the function to look like this:

$$\delta_{Q'} := (q, s) \mapsto \{q' \in Q \mid \delta(q', s) \in q\} \quad (17)$$

We can essentially fill in the rest of the definitions now:

- Q' 's state space is $\mathcal{P}(Q)$
- $\delta_{Q'}$ was defined above
- $i_{Q'}$ is the set containing all the final states in Q
- $F_{Q'}$ is the set containing the initial state in Q

Here is where the magic happens. The function $\delta_{Q'}$ is essentially a function over sets. Even though it's doing $Q \times \Sigma \rightarrow Q$, this is equivalent to $Q \rightarrow (\Sigma \rightarrow Q)$, which is $Q \rightarrow Q^\Sigma$. Over in our reverse language, we have $\delta_{Q'}$ having type $\mathcal{P}(Q) \times \Sigma \rightarrow \mathcal{P}(Q)$, which is the same as $\mathcal{P}(Q) \rightarrow \mathcal{P}(Q)^\Sigma$.

This is exactly the behavior of the contravariant power functor!

Let's step back one step and build up to this insight. If we just considered the category of sets \mathbf{Set} , the *powerset functor* $\mathcal{P} : \mathbf{Set} \rightarrow \mathbf{Set}$ maps every set X to its power set $\mathcal{P}(X)$. In particular, for morphisms, it has type $\mathcal{P}\left(X \xrightarrow{f} Y\right) : \mathcal{P}(X) \rightarrow \mathcal{P}(Y)$. For any subset of X , it maps f restricted to that subset to obtain an image that is in the subset of Y .

There is also a contravariant version, which we will denote $\mathcal{P}^* : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$. In particular, the action on morphisms has type $\mathcal{P}^*\left(X \xrightarrow{f} Y\right) : \mathcal{P}(Y) \rightarrow \mathcal{P}(X)$. This action maps a subset of Y into the *pre-image* under f , which is a subset of X .

This is exactly what our $\delta_{Q'}$ is doing! In essence, our desired transformation is performing the covariant powerset functor on each of the components of the automata. We can concretize this idea by constructing a forgetful functor $U : \text{DA}_{\mathcal{L}} \rightarrow \mathbf{Set}$ which simply takes the set of states and forgets all the other structure. Then, we can see clearly that:

$$U(Q) \xrightarrow{\mathcal{P}^*} U(Q') \quad (18)$$

Here we introduce a powerful property of the contravariant powerset functor:

Lemma 3.3.1

The contravariant powerset functor \mathcal{P}^* is left adjoint to its own opposite functor $\mathcal{P}^{*\text{op}}$.

Proof. The contravariant powerset functor is $\mathcal{P}^* : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$. Its opposite functor would be $\mathcal{P}^{*\text{op}} : \mathbf{Set} \rightarrow \mathbf{Set}^{\text{op}}$. This amounts to showing that for any $X, Y \in \mathbf{Set}$:

$$\begin{aligned} \mathbf{Hom}_{\mathbf{Set}}(\mathcal{P}^*(X), Y) &\cong \mathbf{Hom}_{\mathbf{Set}}(X, \mathcal{P}^{*\text{op}}(Y)) \\ \mathbf{Hom}_{\mathbf{Set}}(Y, \mathcal{P}^*(X)) &\cong \mathbf{Hom}_{\mathbf{Set}}(X, \mathcal{P}^{*\text{op}}(Y)) \end{aligned} \quad (19)$$

and really, both \mathcal{P}^* and $\mathcal{P}^{*\text{op}}$ have the same action on sets, which is to turn them into their power sets. So really, we are trying to show:

$$\mathbf{Hom}(Y, \mathcal{P}(X)) \cong \mathbf{Hom}(X, \mathcal{P}(Y)) \quad (20)$$

We can show this via creating a bijection. First, the forward direction: define $\text{fwd} : \mathbf{Hom}(Y, \mathcal{P}(X)) \rightarrow \mathbf{Hom}(X, \mathcal{P}(Y))$ with the function $\lambda(f : Y \rightarrow \mathcal{P}(X)).\lambda(x : X).\{y \in Y \mid x \in f(y)\}$.³

Then, define the inverse direction $\text{bwd} : \mathbf{Hom}(X, \mathcal{P}(Y)) \rightarrow \mathbf{Hom}(Y, \mathcal{P}(X))$ with the function $\lambda(g : X \rightarrow \mathcal{P}(Y)).\lambda(y : Y).\{x \in X \mid y \in g(x)\}$.

It's easily to show the inverses compose to the identity:

$$\begin{aligned} \forall f. \quad f &= \text{bwd}(\text{fwd}(f)) \\ &= \lambda(y : Y).\{x \in X \mid y \in \text{fwd}(f)(x)\} \\ &= \lambda(y : Y).\{x \in X \mid y \in (\lambda(x : X).\{y \in Y \mid x \in f(y)\})(x)\} \\ &= \lambda(y : Y).\{x \in X \mid y \in \{y \in Y \mid x \in f(y)\}\} \\ &= \lambda(y : Y).\{x \in X \mid x \in f(y)\} \\ &= \lambda(y : Y).f(y)\} \\ &= f \end{aligned} \quad (21)$$

and the other way:

$$\begin{aligned} \forall g. \quad g &= \text{fwd}(\text{bwd}(g)) \\ &= \lambda(x : X).\{y \in Y \mid x \in \text{bwd}(g)(y)\} \\ &= \lambda(x : X).\{y \in Y \mid x \in (\lambda(y : Y).\{x \in X \mid y \in g(x)\})(y)\} \\ &= \lambda(x : X).\{y \in Y \mid x \in \{x \in X \mid y \in g(x)\}\} \\ &= \lambda(x : X).\{y \in Y \mid y \in g(x)\} \\ &= \lambda(x : X).g(x) \\ &= g \end{aligned} \quad (22)$$

□

³Sorry about the sudden lambda notation, I needed to see the type of the argument inline for a second...

[Lemma 3.3.1](#) is essentially saying that the contravariant powerset functor is self-dual. The strategy now is: if we can somehow lift this self-duality into $DA_{\mathcal{L}}$, we will be able to get an operation that turns reachable automata into observable automata and vice versa.

We will devise a functor called $\bar{2}$, which is self-dual, that makes this diagram commute:

$$\begin{array}{ccc}
 & \bar{2} & \\
 DA_{\mathcal{L}} & \xrightarrow{\quad \perp \quad} & DA_{\text{rev}(\mathcal{L})}^{\text{op}} \\
 \downarrow U & \xleftarrow{\quad \bar{2}^{\text{op}} \quad} & \downarrow U \\
 \text{Set} & \xrightarrow{\quad \mathcal{P}^* \quad} & \text{Set}^{\text{op}} \\
 & \xleftarrow{\quad \perp \quad} & \\
 & \mathcal{P}^{*\text{op}} &
 \end{array}$$

This essentially lifts the self-duality of \mathcal{P}^* into $DA_{\mathcal{L}}$. The paper (Bonchi et al. 2014) notes that this lifting is true generically

Definition 3.3.2 (Self-dual powerset functor lift)

Define the functor $\bar{2} : DA_{\mathcal{L}} \rightarrow DA_{\text{rev}(\mathcal{L})}^{\text{op}}$ as:

- action on objects: $\bar{2}(Q)$ applies the contravariant powerset functor to all components of the automaton:
 - Q is transformed into the power set of states $\mathcal{P}(Q)$
 - δ_Q is transformed into the pre-image map as described in Equation 17. Note that this recognizes the reverse language.
 - i_Q is transformed into the set of all final states
 - F_Q is transformed into the set of all initial states
- action on morphisms maps a morphism $Q_1 \xrightarrow{f} Q_2$ to its inverse mapping $\mathcal{P}(Q_2) \xrightarrow{\bar{\mathcal{P}}(f)} \mathcal{P}(Q_1)$


Proof. Omitting for brevity. See (Bonchi et al. 2014, Proposition 9.1) for the full proof. □

3.4 (Co-)reflective subcategories

In order to see how all of this machinery comes together, we will introduce some subcategories of $DA_{\mathcal{L}}$:

- $rDA_{\mathcal{L}}$ is the category of *only* reachable automata in $DA_{\mathcal{L}}$
- $oDA_{\mathcal{L}}$ is the category of *only* observable automata in $DA_{\mathcal{L}}$
- $mDA_{\mathcal{L}}$ is the category of *only* minimal automata in $DA_{\mathcal{L}}$

Since minimal automata are ones that are both reachable and observable, it must also be true that $mDA_{\mathcal{L}}$ is a subcategory of both $rDA_{\mathcal{L}}$ and $oDA_{\mathcal{L}}$. This means the subcategory relationships can be represented as a sort of diamond diagram:



The arrows between them represent inclusion functors. In fact, these are a special kind of subcategory known as a *reflective* (and *co-reflective*) subcategory. For reflective subcategories, this means that the inclusion functor has a left adjoint. Dually, for co-reflective subcategories, this means that the inclusion functor has a right adjoint. Let us see why this holds.

Theorem 3.4.1

$rDA_{\mathcal{L}}$ is a co-reflective subcategory of $DA_{\mathcal{L}}$. In other words, supposing $I : rDA_{\mathcal{L}} \hookrightarrow DA_{\mathcal{L}}$ is the inclusion functor, there exists a functor $R : DA_{\mathcal{L}} \rightarrow rDA_{\mathcal{L}}$ such that

$$I \dashv R \tag{23}$$


Proof. For our action on objects, assume we are dealing with an object $Q : DA_{\mathcal{L}}$. We are looking to construct a functor R that is right-adjoint to I .

To do this, we must first define R , then construct a unit $\eta : \text{id}_{rDA_{\mathcal{L}}} \Rightarrow R \circ I$ and a counit $\epsilon : I \circ R \Rightarrow \text{id}_{DA_{\mathcal{L}}}$ that satisfies the triangle inequalities below. Since we are looking for a functor that takes an automaton with possibly unreachable states into its closest reachable approximation, the most obvious thing to do is simply discarding all unreachable states. Using the machinery we used above to define reachability, we can say that $R(Q)$ takes the codomain of the morphism $\text{Init} \rightarrow Q$. Thus, R is obviously reachable, and lies in $rDA_{\mathcal{L}}$.

Next, define the unit: for some $Q' : rDA_{\mathcal{L}}$, we must define $\eta_{Q'} : Q' \rightarrow R(I(Q'))$. We can define this as the *identity* morphism, since this round trip is a no-op both ways.

For the counit, for some $Q : DA_{\mathcal{L}}$, we must define $\epsilon_Q : I(R(Q)) \rightarrow Q$. Since $R(Q)$ contains the subset of Q 's states that are reachable, it's actually a subset. So ϵ_Q can just be a subset inclusion (notice that unlike above, $I(R(Q))$ may be a different object than just Q).

Let's check that our definitions for unit and counit satisfy the triangle identities:



For the first triangle, we can write this equationally as:

$$\epsilon I \circ I\eta = 1_I \tag{24}$$

Instantiating this with our $Q' : \text{rDA}_{\mathcal{L}}$:

$$\begin{aligned} (\epsilon I)_{Q'} \circ (I\eta)_{Q'} &= (1_I)_{Q'} \\ I(Q') \xrightarrow{(I\eta)_{Q'}} I(R(I(Q'))) &\xrightarrow{(\epsilon I)_{Q'}} I(Q') = I(Q') \xrightarrow{(1_I)_{Q'}} I(Q') \end{aligned} \quad (25)$$

Well, we have defined η to be a no-op, so the left side $(I\eta)_{Q'}$ is the identity. Since ϵ is a subset inclusion, for reachable automata, $(\epsilon I)_{Q'}$ is *also* the identity. Thus, identity compose identity is the identity.

For the second triangle, we can write this equationally as:

$$\epsilon R \circ R\eta = 1_R \quad (26)$$

Instantiating this with $Q : \text{DA}_{\mathcal{L}}$:

$$\begin{aligned} (\epsilon R)_Q \circ (R\eta)_Q &= (1_R)_Q \\ R(Q) \xrightarrow{(R\eta)_Q} R(I(R(Q))) &\xrightarrow{(\epsilon R)_Q} R(Q) = R(Q) \xrightarrow{(1_R)_Q} R(Q) \end{aligned} \quad (27)$$

Again, we have η is a no-op, so the left side is the identity. The right side goes from reachable to reachable, so the subset inclusion ϵ is simply restricted to the identity by default here.

Both triangle inequalities check out, so our adjunction is defined. □

Similarly, for $\text{oDA}_{\mathcal{L}}$:

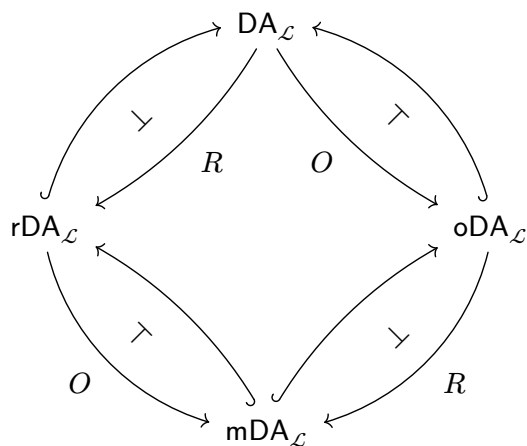
Theorem 3.4.2

$\text{oDA}_{\mathcal{L}}$ is a reflective subcategory of $\text{DA}_{\mathcal{L}}$. In other words, supposing $I : \text{oDA}_{\mathcal{L}} \hookrightarrow \text{DA}_{\mathcal{L}}$ is the inclusion functor, there exists a functor $O : \text{DA}_{\mathcal{L}} \rightarrow \text{oDA}_{\mathcal{L}}$ such that


$$O \dashv I \quad (28)$$

Exercise 3.4.1. Prove [Theorem 3.4.2](#). The proof should look mostly like a dualization of the proof of [Theorem 3.4.1](#).

The corresponding adjoint functor for $\text{rDA}_{\mathcal{L}}$, called R , then restricts automata in the input category to their reachable subset. Similarly, the corresponding adjoint functor for $\text{oDA}_{\mathcal{L}}$, called O , restricts automata in *its* input category to their observable subset. And then at the bottom, $\text{mDA}_{\mathcal{L}}$ is both a reflective subcategory and co-reflective subcategory of $\text{oDA}_{\mathcal{L}}$ and $\text{rDA}_{\mathcal{L}}$, respectively. Let us update the diamond diagram to see these new functors:



Finally, due to our self-dual functor $\bar{2}$, we have a functor between $rDA_{\mathcal{L}}$ and $oDA_{\mathcal{L}}$, that reflects reachable automata into observable automata in the *reverse* language: Thus, our final minimization morphism is simply a concatenation of a series of functors:



Theorem 3.4.3 (Brzozowski minimization algorithm)

The functor:

$$(\bar{2}^{op} \circ R^{op} \circ \bar{2} \circ R) : DA_{\mathcal{L}} \rightarrow mDA_{\mathcal{L}} \tag{29}$$

minimizes deterministic automata.

Proof. Due to the way the functors were constructed, we get the following two properties “for free” (ignoring all the work that went into constructing the functors):

- the series of functors preserves the language being recognized (notice the subscript \mathcal{L})
- the resulting automaton is minimal (by definition of $mDA_{\mathcal{L}}$)

Thus, the theorem holds by construction. □

4 Exercises

Here is a table of some of the exercises that have been distributed throughout the explainer.

[Exercise 1.1](#) 3

Exercise 1.1.1	5
Exercise 1.1.2	5
Exercise 1.2.1	7
Exercise 2.1	8
Exercise 2.1.1	12
Exercise 3.2.1	17
Exercise 3.4.1	23

Bibliography

- Bezhanishvili, Nick, Clemens Kupke, and Prakash Panangaden. 2012. “Minimization via Duality.” In *Logic, Language, Information and Computation*, edited by Luke Ong and Ruy de Queiroz, *Logic, Language, Information and Computation*. Springer. https://doi.org/10.1007/978-3-642-32621-9_14.
- Bonchi, Filippo, Marcello M. Bonsangue, Helle H. Hansen, Prakash Panangaden, Jan J. M. M. Rutten, and Alexandra Silva. 2014. “Algebra-Coalgebra Duality in Brzowski's Minimization Algorithm.” *ACM Transactions on Computational Logic* 15 (1): 1–29. <https://doi.org/10.1145/2490818>.

Bibliography

- [1] Julian Asilis. *Probability Monads*. PhD thesis, Ph. D. thesis, PhD thesis, Harvard University Cambridge, Massachusetts, 2020.
- [2] Robert Atkey, Patricia Johann, and Neil Ghani. Refining inductive types. *Logical Methods in Computer Science*, Volume 8, Issue 2, Jun 2012.
- [3] Robert J Aumann. Borel structures for function spaces. *Illinois Journal of Mathematics*, 5(4):614–630, 1961.
- [4] Andrej Bauer. Notes on realizability. February 2025.
- [5] Nick Benton and Chung-Kil Hur. Realizability and compositional compiler correctness for a polymorphic language. Technical Report MSR-TR-2010-62, April 2010.
- [6] Bodil Biering, Lars Birkedal, and Noah Torp-Smith. Bi-hyperdoctrines, higher-order separation logic, and abstraction. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 29(5):24–es, 2007.
- [7] Rahul Chhabra. Modest Sets are Equivalent to PERs, 2024.
- [8] Guy Cousineau, P-L Curien, and Michel Mauny. The categorical abstract machine. *Science of computer programming*, 8(2):173–202, 1987.
- [9] Roy L Crole. *Categories for types*. Cambridge University Press, 1993.
- [10] Haskell B. Curry and Robert Feys. *Combinatory logic*, volume 1 of *Studies in logic and the foundations of mathematics*. North-Holland, Amsterdam, 1958.
- [11] Brijesh Dongol, Ian J Hayes, and Georg Struth. Convolution as a unifying concept: Applications in separation logic, interval

- calculi, and concurrency. *ACM Transactions on Computational Logic (TOCL)*, 17(3):1–25, 2016.
- [12] Derek Dreyer, Amal Ahmed, and Lars Birkedal. Logical step-indexed logical relations. In *2009 24th Annual IEEE Symposium on Logic In Computer Science*, pages 71–80. IEEE, 2009.
- [13] Neil Ghani, Patricia Johann, and Clément Fumex. Fibrational induction rules for initial algebras. In Anuj Dawar and Helmut Veith, editors, *Computer Science Logic*, pages 336–350, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [14] Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. A convenient category for higher-order probability theory. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12. IEEE, 2017.
- [15] Pieter J W Hofstra. Partial Combinatory Algebras and Realizability Toposes.
- [16] Bart Jacobs. Comprehension categories and the semantics of type dependency. *Theor. Comput. Sci.*, 107(2):169–207, January 1993.
- [17] Bart Jacobs. *Categorical Logic and Type Theory*. Number 141 in Studies in Logic and the Foundations of Mathematics. North Holland, Amsterdam, 1999.
- [18] S. C. Kleene. On the interpretation of intuitionistic number theory. *Journal of Symbolic Logic*, 10(4):109–124, December 1945.
- [19] Stephen Cole Kleene. *Introduction to metamathematics*, volume 1 of *Bibliotheca mathematica*. North-Holland, Amsterdam, 1 edition, 1952.
- [20] Joachim Lambek and Philip J Scott. *Introduction to higher-order categorical logic*, volume 7. Cambridge University Press, 1988.
- [21] F William Lawvere and Stephen H Schanuel. *Conceptual mathematics: a first introduction to categories*. Cambridge University Press, 2009.
- [22] John M Li, Amal Ahmed, and Steven Holtzen. Lilac: a modal separation logic for conditional probability. *Proceedings of the ACM on Programming Languages*, 7(PLDI):148–171, 2023.
- [23] John Longley. Matching typed and untyped realizability. *Electronic Notes in Theoretical Computer Science*, 23(1):74–100, 1999.

- [24] John Longley. Unifying typed and untyped realizability. May 1999.
- [25] John R. Longley. *Realizability Toposes and Language Semantics*. PhD thesis, University of Edinburgh, 1994.
- [26] Saunders Mac Lane. *Categories for the working mathematician*, volume 5. Springer Science & Business Media, 1998.
- [27] Per Martin-Löf. On the meanings of the logical constants and the justifications of the logical laws. *Nordic journal of philosophical logic*, 1(1):11–60, 1996.
- [28] Max S. New. Category theory for computer scientists. Lecture notes, 2023. Scribed by Pranav Srinivasan and Kevin Wang.
- [29] Daniel Baker Patterson. *Interoperability through realizability: expressing high-level abstractions using low-level code*. PhD thesis, Northeastern University, 2022.
- [30] Andrew M Pitts. *Nominal sets: Names and symmetry in computer science*. Cambridge University Press, 2013.
- [31] Gordon Plotkin and John Power. Notions of computation determine monads. In *International Conference on Foundations of Software Science and Computation Structures*, pages 342–356. Springer, 2002.
- [32] John C Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing 83, Proceedings of the IFIP 9th World Computer Congress*, pages 513–523, 1983.
- [33] John C Reynolds. The coherence of languages with intersection types. In *International Symposium on Theoretical Aspects of Computer Software*, pages 675–700. Springer, 1991.
- [34] Michael B Smyth and Gordon D Plotkin. The category-theoretic solution of recursive domain equations. *SIAM Journal on Computing*, 11(4):761–783, 1982.
- [35] Sam Staton. Commutative semantics for probabilistic programming. In *European Symposium on Programming*, pages 855–879. Springer, 2017.
- [36] Philip Wadler. Recursive types for free! Unpublished draft, 1990.