

Making Discrete Decisions Based on Continuous Values

by

Benjamin Sherman

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degree of
Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2017

© Benjamin Sherman, MMXVII. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Author
Department of Electrical Engineering and Computer Science
May 19, 2017

Certified by
Adam Chlipala
Associate Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Certified by
Michael Carbin
Jamieson Career Development Assistant Professor
of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by
Leslie A. Kolodziejcki
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Theses

Making Discrete Decisions Based on Continuous Values

by

Benjamin Sherman

Submitted to the Department of Electrical Engineering and Computer Science
on May 19, 2017, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

Many safety-critical software systems are cyber-physical systems that compute with continuous values; confirming their safety requires guaranteeing the accuracy of their computations. It is impossible for these systems to compute (total and deterministic) discrete computations (e.g., decisions) based on connected input spaces such as \mathbb{R} . We propose a programming language based on constructive topology, whose types are spaces and programs are executable continuous maps, that facilitates making formal guarantees of accuracy of computed results. We demonstrate that discrete decisions can be made based on continuous values by permitting nondeterminism. This thesis describes variants of the programming language allowing nondeterminism and/or partiality, and introduces two tools for creating nondeterministic programs on spaces. *Overlapping pattern matching* is a generalization of pattern matching in functional programming, where patterns need not represent decidable predicates and also may overlap, allowing potentially nondeterministic behavior in overlapping regions. *Binary covers*, which are pairs of predicates such that at least one of them holds, yield a formal logic for constructing approximate decision procedures.

Thesis Supervisor: Adam Chlipala

Title: Associate Professor of Electrical Engineering and Computer Science

Thesis Supervisor: Michael Carbin

Title: Jamieson Career Development Assistant Professor
of Electrical Engineering and Computer Science

Acknowledgments

I thank my advisors, Adam Chlipala and Michael Carbin, who contributed to this thesis in many ways, by helping me to focus, justify, and explain my research, and by reviewing drafts of this work and similar ones. I am grateful for their support of my exploration of a subject that is relatively far from their interests.

I was fortunate to work with Luke Sciarappa during the summer of 2016 on related subjects that led us to the subject of this thesis. Luke taught me about categories and toposes, and we worked together to understand constructive topology.

Adam, Michael, Luke, and I submitted a draft on similar material to the Logic in Computer Science (LICS) conference in 2017. We thank the four anonymous reviewers for their helpful feedback. One review was particularly helpful, noting the possibility of using open maps as patterns and noticing Remark 4.17.

Tej Chajed, Gabriel Scherer, and Muralidaran Vijayaraghavan read drafts of this thesis or prior related drafts and provided helpful feedback.

I am grateful for the financial support for my graduate research, which I have received from an MIT SMA fellowship as well as NSF grant CCF-1512611.

I thank my parents for their love and support.

Contents

1	Introduction	11
1.1	Example application: approximate root-finding	12
1.2	Why formal topology?	12
1.2.1	Reasoning about programs as mathematical expressions	12
1.2.2	Computational content of compactness	14
1.3	Discrete decision-making on connected spaces	14
1.4	Terminology and foundations	16
2	Constructive topology	17
2.1	Introduction via the real numbers	17
2.2	Formal topology	20
2.2.1	Inductively generated formal spaces	23
2.2.2	Locales	24
2.2.3	The computational content of formal topology	26
2.3	Spaces	27
2.3.1	Discrete spaces	27
2.3.2	Open subspaces	28
2.3.3	The Sierpinski space	29
2.3.4	Disjoint unions (sums)	29
2.3.5	Products	31
2.3.6	Metric spaces, including \mathbb{R}	32
2.4	As a programming language	34
2.4.1	Opens and Σ -valued maps	35
3	Partiality and nondeterminism	37
3.1	Partiality	39
3.1.1	Lifted spaces	41
3.2	Nondeterminism	44
3.2.1	Nondeterministic powerspaces	45
3.3	Both partiality and nondeterminism	46
3.3.1	Lower powerspaces	47
3.4	Open maps and open embeddings	47
3.4.1	Open maps	48
3.4.2	Open embeddings	50
3.5	Summary	54

3.6	Variants of the programming language	54
4	Overlapping pattern matching	57
4.1	Example uses of overlapping pattern matching	58
4.1.1	A familiar example	58
4.1.2	Approximate computation and decision-making	58
4.1.3	Manipulating partial data	59
4.1.4	“Sheafification” of constructions	61
4.2	Pattern families	61
4.2.1	Totality	63
4.2.2	Determinism	64
4.2.3	Totality and determinism	66
4.2.4	Syntax of patterns	67
4.3	Pattern matching with “gluing” conditions	69
4.4	Properties of overlapping pattern matching	75
4.4.1	Partial patterns and lifted spaces	75
4.4.2	Related to Grothendieck pretopologies	76
4.4.3	Evaluation of patterns	77
5	Binary covers: nondeterministic decision procedures	79
5.1	Binary covers as nondeterministic decisions	79
5.2	Quantification over compact/overt spaces	81
5.2.1	On compact/overt spaces	81
5.2.2	On compact/overt subspaces	84
5.3	Binary covers on \mathbb{R}	87
5.4	Approximate root-finding	88
6	Related work	91
6.1	Alternative theories of constructive topology	91
6.2	Related programming formalisms involving continuity, partiality, or nondeterminism	91
6.3	Overlapping pattern matching	92
6.4	Binary covers	92
6.5	Implementation	93
7	Discussion and conclusion	95
7.1	Coq implementation	95
7.2	Programming with a theorem prover	96
7.3	Future work	96
	Bibliography	99
	List of symbols	103

List of Figures

1-1	Car approaching a yellow light.	15
2-1	Rules that a formal cover relation must satisfy.	21
2-2	Rules that a point of a formal space must satisfy.	22
2-3	Rules that a continuous map of formal spaces must satisfy.	23
3-1	The lattice of categories representing potentially nondeterministic or partial maps on spaces.	38
3-2	The restriction of the categories of nondeterministic and partial maps of spaces to the discrete spaces.	39
3-3	A schematic diagram of the lifted space A_{\perp}	41
3-4	Summary of categories of morphisms on formal spaces and functors relating them.	54
4-1	Syntax of patterns	67

Chapter 1

Introduction

Many safety-critical software systems are cyber-physical systems that compute with continuous values, such as space, time, magnitude, and probability; confirming their safety requires guaranteeing the accuracy of their computations. We claim that *formal topology*, a constructive theory of topology, provides a suitable theory for building programs that compute with continuous values. We propose a programming language based on formal topology, whose types are spaces and programs are executable continuous maps, that allows execution of programs that manipulate continuous values, while at the same time allowing reasoning in terms of their mathematical descriptions and guaranteeing accuracy of computed results. The interfaces defined in constructive topology give insight into both how to specify such programs and how to compute with continuous spaces [36, 10, 9, 34].

This thesis examines the problem of making discrete decisions based on continuous values from the perspective of constructive topology. By topological arguments, any program $f : \mathbb{R} \rightarrow_c \mathbb{B}$ that makes a (total and deterministic) discrete decision based on a real number must be a constant map (i.e., always **true** or always **false**). We demonstrate that decisions *can* be made based on continuous values by permitting nondeterminism.

This thesis characterizes two programming language constructs for constructive topology that make it easier to program with continuous values (particularly involving nondeterminism):

- *Overlapping pattern matching* (Chapter 4) generalizes pattern matching on inductive types in functional programming. It differs in that patterns need not correspond to decidable predicates, and patterns are allowed to overlap; an input satisfying multiple patterns may nondeterministically follow any such branch. Because spaces often have few decidable predicates, nondeterminism is often essential for decision-making, and overlapping pattern matching facilitates construction of such programs.
- *Binary covers* (Chapter 5), which are pairs of opens that cover a space, provide an expressive formal logic of “approximately decidable” properties. Most significantly, the logic admits quantification over compact spaces, demonstrating

“approximate decidability” of questions such as whether a real-valued function on a compact space has any roots.

We intend to explain in the remainder of this chapter how this thesis helps to make sense of making decisions based on continuous values, and how the above techniques make it easier to devise such programs. Section 1.1 describes how a rather general approximate-root finding function can be written with the above techniques. Section 1.2 justifies the choice of formal topology as the appropriate formalism for programming with continuous values, and 1.3 explains informally why it is often difficult to make decisions based on continuous values, and why nondeterminism may be necessary.

1.1 Example application: approximate root-finding

Consider the following computational task. Suppose we have an arbitrary continuous function $f : K \rightarrow_c \mathbb{R}$, where K is some compact space (such as the unit interval $[0, 1]$), and want to determine if it has any roots. In general, it is not decidable whether f has any roots [36], but an approximation is possible: fix a tolerance $\varepsilon > 0$. Then either, or both, of the following statements must hold:

- There is some $x \in K$ such that $|f(x)| < \varepsilon$.
- For every $x \in K$, $f(x) \neq 0$.

Though each of the above statements is in general undecidable, we can use the techniques devised in this thesis to define a function that nondeterministically computes that one of the above two statements holds, in the former case *computing* such an x that is “almost a root.” We will observe in Section 5.4 that the pair of bulleted statements in fact comprise a binary cover, and that an overlapping pattern match can be used to construct the program that solves the task.

1.2 Why formal topology?

This section motivates formal topology as a language theory for programming with continuous values, explaining some desirable properties of such a theory and how these properties are absent in alternative methods of specifying and computing with continuous values.

1.2.1 Reasoning about programs as mathematical expressions

A principal desire is to be able to soundly reason about the behavior of programs that manipulate continuous values as the mathematical expressions that they denote. For instance, we might hope that the real-valued programs $(1 + x) - x$ and 1 would be considered equivalent (since the corresponding real-valued mathematical expressions are equivalent).

Notably, floating point lacks this property: in particular, those two programs are not considered equivalent, and for the value $x \triangleq 10^{50}$, for instance, the expression $(1 + x) - x$ returns 0 rather than 1.

Systems lacking such a connection between the programs and the mathematical expressions they are supposed to represent are difficult to program with: a programmer cannot necessarily use a mathematical identity of real-valued expressions to replace one program with another (which may be faster, simpler, or otherwise more efficient) without worrying that it might change the program’s behavior. Similarly, compilers cannot use mathematical identities, either, for the purpose of optimization.

Comparison with floating point

A common approach for computing with real numbers is to use floating-point numbers as a surrogate. Floating point is intended behave like the real numbers: for instance, most floating-point operations return the floating-point value closest to the “ideal” real number result. Unfortunately, this guarantee does not compose well: it is easy to conceive floating-point programs whose results are arbitrarily erroneous compared to the ideal real-valued result. For instance, consider the real-valued function [31, Chapter 9.1]

$$f : \{x : \mathbb{R} \mid x \neq 0\} \rightarrow \mathbb{R}$$

$$f(x) \triangleq \frac{1 - \cos(x)}{x^2},$$

and its double-precision floating-point approximation using a cosine operation that returns an answer to within 1 ULP of the ideal answer. Then $f(1.1 \times 10^{-8})$ is very close to $1/2$ (easily within 10^{-8}), while the floating-point code returns approximately 0.9175. In practice, it is very difficult to bound the error of floating-point functions.

Issues such as overflow, roundoff error, and catastrophic cancellation make it very difficult to reason about the behavior of floating-point programs. Mathematical properties of real numbers do not necessarily hold for floating-point programs. For instance, addition/multiplication of floating-point values is not associative. Accordingly, replacing one floating-point program with another one which is equivalent when viewed as an expression on real numbers can drastically change behavior. Moreover, there may not be a single such program with the least error for all inputs [27].

The lack of the connection between floating point and real numbers makes programming with floating point a difficult task. A programmer cannot necessarily use a mathematical identity of real-valued expressions to replace a floating-point program with another (which may be faster, simpler, or otherwise more efficient) without worrying that it might negatively impact its accuracy. Nor can programmers (without serious additional verification effort) be reasonably assured that their floating-point program is at all accurate.

Compilers also face difficulty in handling floating-point programs. Since the source language of a compiler will generally reflect all details about the low-level floating point, compilers must preserve *exact* behavior of floating-point programs, even

though often the programmer does not rely on such exactness. Accordingly, compilers are consigned to perform only very limited optimization of floating-point code that doesn't change floating point behavior. For instance, they cannot rewrite a program $(1 + x) - x$ to the constant 1 because floating-point addition is not associative.

So floating point unfortunately lacks a key desideratum for a theory of programming with continuous values.

1.2.2 Computational content of compactness

There are several computational frameworks for computing with real numbers that *do* allow reasoning about programs as mathematical expressions. Most can be labeled a form of *exact real arithmetic*, meaning that terms in their language faithfully represent real numbers themselves. A result from computability theory regarding exact real arithmetic underscores the fundamental importance of topology:

Theorem 1.1. *Any computable function¹ $f : \mathbb{R} \rightarrow \mathbb{R}$ is necessarily continuous [49].*

So one is not really missing out by restricting attention to continuous maps.

Constructive topology also has an advantage over other frameworks for exact real arithmetic: the computational content in its definition of compactness. *Compactness* is a fundamental topological property of spaces that generalizes the concept of (Kuratowski) finiteness of sets. Compactness is important for many results in analysis. One such result is that a real-valued continuous function on a compact space has a minimum value. Unlike other frameworks for exact real arithmetic, constructive topology allows this value to be computed.

Comparison with Bishop-style constructive analysis

One way to approach programming with real numbers is with Bishop-style constructive analysis, where one develops the theory of metric spaces in a constructive manner. The c-CoRN library in fact implements Bishop-style constructive analysis within Coq.

However, there are known difficulties with Bishop-style analysis. For instance, the notions of continuity, compactness, and subspaces do not interact well, unless one accepts additional axioms². In particular, Waaldijk [48] proves that in Bishop-style analysis, if one desires certain basic properties of continuous maps to hold, then one must accept an axiom called the fan theorem.

1.3 Discrete decision-making on connected spaces

Lamport makes a bold claim with severe implications [18]:

¹ In the theory of type II computability.

² It is theoretically possible to give a computational model of type theory admitting such axioms via Type II computability theory [2]. However, to compile systems like Coq to such a computational model would be a difficult task.




Figure 1-1: An autonomous car approaches a traffic light that has just turned yellow, and must decide whether to stop before the intersection or proceed through it.

Buridan's Principle. A discrete decision based upon an input having a continuous range of values cannot be made in a bounded length of time.

Lamport's name "Buridan's Principle" for the above claim comes from the famous philosophical problem of Buridan's Ass, which has many variations, but in one form finds a donkey unable to decide between two equally distant and equally delicious bails of hay, finally starving to death.

Why must it be so? Lamport argues that ultimately the decision arises by physical mechanisms, and physics is necessarily continuous. The same argument could be made from the perspective of computability, given that any computable function also is necessarily continuous (theorem 1.1). Buridan's principle can be rephrased as the topological theorem:

Theorem 1.2. *Any continuous map $f : A \rightarrow_c \mathbb{B}$ from a connected space A to the Booleans is necessarily constant.*

Since the connected spaces include those such as \mathbb{R} and \mathbb{R}^n , this theorem precludes the possibility of computing total, deterministic decisions based on real-valued inputs.

Consider an analogous situation, depicted in Figure 1.3, in which an autonomous car approaches a yellow light and must decide whether to stop before the intersection or proceed through it. How can we program the car to safely navigate the intersection, given the impossibility result of Theorem 1.2?

This thesis argues that *nondeterminism* makes it possible to make decisions such as these in bounded time, providing a loophole to work around Buridan's Principle. We claim that in each case there is some (open) region, perhaps small, where both decisions may be made, perhaps from the *exact same* input. In the case of floating point, though a floating-point comparison operation is indeed deterministic, there is effectively nondeterminism in the way measurements are approximated to floating point, and in the rounding in floating-point arithmetic operations.

This thesis explores how to program nondeterministic functions on spaces, developing programming-language constructions for computing approximate decisions that make it easy to reason about what behaviors may occur.

1.4 Terminology and foundations

We intend mathematical statements to be interpreted within a constructive and predicative metatheory, potentially formalizable within, for instance, the predicative fragment of Coq (i.e., without use of the `Prop` universe).

Since all topological notions in this article are pointfree, we will coopt terminology from classical topology without fear of confusion. For instance, we use the word “space” rather than “locale” when describing the pointfree analogue of spaces. Because we want to *compute* with the mathematical structures, we use the propositions-as-types correspondence [30] to encode all propositions as types, such that everything is formalizable within Martin-Löf type theory (as well as the predicative fragment of Coq). Let \mathcal{U} denote the universe of types, where the universe level is left implicit, but we use “small” and “large” to relative sizes of universes.

We use the term “type” to refer to a type, and the term “set” to refer to what is often called a setoid or a Bishop set: a type together with a distinguished equivalence relation on it. We use the symbol \equiv to denote intensional equality on members of any type, reserving $=$ for the equality relation on sets. The reader should assume that if A and B are both sets, then the notation $f : A \rightarrow B$ means that f is a morphism in the category of sets (i.e., it maps equivalent elements of A to equivalent elements of B).

Let Ω the (large) set of propositions: \mathcal{U} endowed with the equivalence relation of bi-implication. Ω defines a completely distributive lattice with small meets and joins, and we use lattice notation (\top , \perp , \wedge , \vee , \leq) accordingly. Propositional functions $P : A \rightarrow \Omega$ represent the subsets of A , and for $a : A$ we use the familiar notation $a \in P$ to denote the proposition $P(a) : \Omega$. We use \cup and \cap to denote union and intersection of subsets, respectively.

Chapter 2

Constructive topology

This chapter introduces formal topology as a theory of computation with continuous values. We will develop the idea using the real numbers (or metric spaces, more generally) as an example. We will identify the computational interface that the real numbers provide (that is, what observations we are allowed to make on real numbers) in terms of their “ideal” (implementation-independent) observable properties, and discover that this interface has the logical structure of a *formal topology* and is shared by all formal spaces.

2.1 Introduction via the real numbers

This section introduces formal topology via the example of the real numbers, \mathbb{R} . We begin by defining the real numbers by more conventional means (as Cauchy approximations) and attempting to define a program that makes a nondeterministic decision based on a real number in this framework. It is difficult to reason spatially about which real numbers may be mapped to **true** or **false** in this way, because one must reason about the many possible implementations of a given real number. We then adjust the definition of the real numbers so that they can be specified in an implementation-independent way, which leads to their definition as a formal space. In this framework, we first *specify* how the nondeterministic behavior should behave spatially, and then derive an *implementation* (formal proof) that computes the behavior. Because the specification completely describes the desired behavior, one need not reason about the form of the implementation beyond the fact that some implementation exists.

The real numbers can be viewed as arbitrarily fine rational-valued approximations to some well-defined quantity. For instance, given a square with sides of length 1, the diagonal has length $\sqrt{2}$, which is a real number but not a rational number. That is, one can produce rational numbers which come ever closer to the length of the diagonal, though no rational number alone suffices to characterize it. The ability to produce arbitrarily fine approximations is a useful and well-behaved one: it is this idea that the real numbers encapsulate.

One standard construction of \mathbb{R} is as the set of Cauchy sequences of rational numbers. A sequence $x : \mathbb{N} \rightarrow \mathbb{Q}$ is *Cauchy* if for all $\varepsilon : \mathbb{Q}^+$, there is some $N : \mathbb{N}$

such that for all $m, n : \mathbb{N}$ where $N \leq n$ and $N \leq m$, we have $d(x_m, x_n) < \varepsilon$, where $d : \mathbb{Q} \times \mathbb{Q} \rightarrow \mathbb{Q}$ denotes the distance between two rational numbers. Two sequences x and y are equivalent if the distance between their approximations $d(x_n, y_n)$ tends to 0 as n approaches ∞ .

This standard definition can be reduced to a simpler one [24]. We can instead define *regular functions* **RegFunc**, functions $x : \mathbb{Q}^+ \rightarrow \mathbb{Q}$ such that for all $\delta, \varepsilon : \mathbb{Q}^+$, $d(x_\delta, x_\varepsilon) \leq \delta + \varepsilon$. Intuitively, we imagine that x_δ is at most a distance δ away from the “true” (but unrepresentable as a rational number) value of x , such that $d(x_\delta, x_\varepsilon) \leq \delta + \varepsilon$ represents that each approximation is consistent with each other in approximating the same limiting value, by imagining a notional application of the triangle inequality,

$$d(x_\delta, x_\varepsilon) \leq d(x_\delta, x) + d(x, x_\varepsilon) \leq \delta + \varepsilon.$$

The relational characterization of consistent approximations of a real number also suffices to characterize when two regular functions implement the same real number: the real numbers \mathbb{R} are the regular functions where two regular functions x and y are considered equivalent if for all $\delta, \varepsilon : \mathbb{Q}^+$, $d(x_\delta, y_\varepsilon) \leq \delta + \varepsilon$. For a regular function $x : \mathbf{RegFunc}$, let $[x] : \mathbb{R}$ be the real number that it implements.

Suppose we want to approximately compute whether a given real number is positive. It is impossible to do so exactly, so suppose we pick some error tolerance $\varepsilon : \mathbb{Q}^+$, and we wish to define a function $f : \mathbf{RegFunc} \rightarrow \mathbb{B}$ with the following behavior on real numbers:

1. If $f(x) = \mathbf{true}$ then $[x] > 0$.
2. If $f(x) = \mathbf{false}$, then $[x] \leq \varepsilon$.

We observe that for any $\delta : \mathbb{Q}^+$, $[x] - \delta \leq x_\delta \leq [x] + \delta$, so with a bit of reasoning we find that a valid implementation of f is $f(x) \triangleq x_{\varepsilon/2} > \varepsilon/2$ (recalling that comparison of rational numbers is decidable). It is far from obvious that this definition of f satisfies the expectations defined above, because one must reason about all implementations of a given real number.

This is an issue with the definition of \mathbb{R} , as it does not clearly describe what a real number *is* independently of some implementation of that real number as a regular function, so it is unclear what are properties of real numbers rather than of particular regular function. We will tease out those properties that suffice to characterize a real number that

1. are implementation-independent and
2. can be observed by querying *any* implementation (meaning that if the property holds, we can observe it to hold by interacting with the implementation).

We call these properties (subsets) *open*.

Such a characterization would allow us to decouple a real number into a specification that says which of these such properties a real number satisfies, and an implementation which allows computational observation of *only* those properties (and no

additional information), such that we could not care *which* implementation we are given, only that we are given any implementation at all (which also serves as a model to confirm that the specification is consistent).

Examples of open properties of \mathbb{R} are $\cdot > 0$ (i.e., positivity) and $\cdot < \varepsilon$. The property $\cdot \leq \varepsilon$ is not open; its interior $\cdot < \varepsilon$ is the largest open included within it (every subset has an interior). Once we have characterized \mathbb{R} as a formal space, defining the approximate comparison function will reduce to simply finding a formal proof $\top \leq (\cdot > 0) \vee (\cdot < \varepsilon)$ that every number is either positive or less than ε ¹.

The regular functions come close to explaining \mathbb{R} in the terms as above, but don't quite work. When a regular function $x : \mathbb{Q}^+ \rightarrow \mathbb{Q}$ is queried for some tolerance $\varepsilon : \mathbb{Q}^+$ to produce an answer $x_\varepsilon = q$, we learn that $d([x], q) \leq \varepsilon$, but we may not necessarily learn this same fact from a different representative of the same real number. For instance, we can define the regular functions 0^+ and 0^- by $0_\varepsilon^+ \triangleq \varepsilon$ and $0_\varepsilon^- \triangleq -\varepsilon$. Both of these regular functions implement the real number 0, from 0^+ we easily learn that 0 is nonnegative, which we can never learn by observing 0^- . So these types of observations are not necessarily observable given any implementation.

We can ensure that the regular functions only produce observable properties by making a small change. Whenever a regular function $x : \text{RegFunc}$ allows us to observe that $d([x], q) \leq \varepsilon$, we expect that it should then be able to make tighter this estimate, that is, it should be able to produce a $\varepsilon' > 0$ such that $\varepsilon' < \varepsilon$ but still $d([x], q) \leq \varepsilon'$. Therefore, any estimate that says $d([x], q) \leq \varepsilon$ in fact tells us that x is *strictly less than* ε away from q .

We will claim that after making this adjustment, any observation made with one implementation can be made with another. First, let's make the adjustment precise. To allow an implementation x of a real number $[x]$ to make previous approximations more precise, we need to make a datatype representing previous approximations. For each $q : \mathbb{Q}$ and $\varepsilon : \mathbb{Q}^+$, let $[x] \Vdash B_\varepsilon(q)$ denote the type of evidence that $[x]$ lies in the open ball of radius ε around q (or equivalently, $d(x, q) < \varepsilon$). Then any implementation x offers two ways by which we can learn more about $[x]$,

$$\frac{\varepsilon : \mathbb{Q}^+}{\sum_{q:\mathbb{Q}} [x] \Vdash B_\varepsilon(q)} \text{ APPROX} \qquad \frac{[x] \Vdash B_\varepsilon(q)}{\sum_{\varepsilon':\mathbb{Q}^+ | \varepsilon' < \varepsilon} [x] \Vdash B_{\varepsilon'}(q)} \text{ SHRINK} ,$$

as well as a rule that allows us to “forget” information,

$$\frac{[x] \Vdash B_\varepsilon(q) \quad B_\varepsilon(q) \subseteq B_\delta(r)}{[x] \Vdash B_\delta(r)} \text{ WEAKEN} .$$

Proposition 2.1. *If $[x] = [y]$, then if $[x] \Vdash B_\varepsilon(q)$, we can observe the same of y by applications of the APPROX and SHRINK rules.*

Proof sketch. By SHRINK, there is some $\varepsilon' : \mathbb{Q}^+$ such that $\varepsilon' < \varepsilon$ and $[x] \Vdash B_{\varepsilon'}(q)$.

¹ It will satisfy an even tighter specification, that if $f(x) = \text{false}$, then $[x]$ is *strictly less than* ε .

Define $\delta : \mathbb{Q}^+$ by $\delta \triangleq \varepsilon - \varepsilon'$. Then by APPROX there is some $r : \mathbb{Q}$ such that $y \Vdash B_\delta(r)$. Since $[x] = [y]$, $B_{\varepsilon'}(q)$ and $B_\delta(r)$ must overlap, and therefore by WEAKEN, $y \Vdash B_\varepsilon(q)$. \square

Since any open subset is a union of open balls, we can define for any open subset U and any implementation x of a real number, the type of evidence $[x] \Vdash U$ that $[x]$ lies in U as

$$[x] \Vdash U \triangleq \bigcup \{ [x] \Vdash B_\varepsilon(q) \mid B_\varepsilon(q) \subseteq U \}.$$

We can now specify the task of making the nondeterministic decision as finding a function that for all implementations x of real numbers computes

$$([x] \Vdash \cdot > 0) + ([x] \Vdash \cdot < \varepsilon),$$

and returns **true** in the case of **inl** and **false** in the case of **inr**. Because of the proof terms that such a function returns, it would be guaranteed to behave as was originally desired. The derivation of the above function would mirror the definition using regular functions: we would first use APPROX($\varepsilon/2$) to determine $\sum_{q:\mathbb{Q}} [x] \Vdash B_\varepsilon(q)$, and check to see if the result whether $q > \varepsilon/2$, and then use the WEAKEN rule to weaken to $[x] \Vdash \cdot > 0$ if $q > \varepsilon/2$ and $[x] \Vdash \cdot < \varepsilon$ if $q \leq \varepsilon/2$.

This shift succeeds in isolating specifications of real numbers from their implementations. To specify a real number, it suffices to say what balls $B_\varepsilon(q)$ it lies in. The computational rules of an implementation, e.g. APPROX and SHRINK, explicitly indicate what the results of the computation mean in terms of the observable properties of the real number that the implementation represents.

The above characterization is informal and incomplete; the remainder of this chapter will give a rigorous description of formal topology, and in particular Section 2.3.6 will characterize metric spaces. The basic idea in constructing \mathbb{R} as a formal space is to first describe a logic of the balls $B_\varepsilon(q)$ that completely describes which balls are covered by unions of other balls. A point is specified by the formal balls that it lies in, and an implementation of a real number allows one to *compute* which formal balls that it lies in, which also serves to guarantee the soundness of the specification of that real number.

2.2 Formal topology

In this section, we formally describe formal spaces, and then the inductively generated formal spaces, which will form a category **FSpc**.

Definition 2.2. A formal space A is a preorder $\mathcal{O}_B(A)$ together with a relation $\triangleleft : \mathcal{O}_B(A) \rightarrow (\mathcal{O}_B(A) \rightarrow \Omega) \rightarrow \Omega$ that satisfies the rules in Figure 2-1.

The intuitive meaning of this structure is that $\mathcal{O}_B(A)$ is a set of basic opens on A , where $a \leq_{\mathcal{O}_B(A)} b$ if the basic open a is included in the basic open b . A subset $U : \mathcal{O}_B(A) \rightarrow \Omega$ represents an arbitrary union of the basic opens in U . This means that every open of A can be represented as a subset $U : \mathcal{O}_B(A) \rightarrow \Omega$. Then $a \triangleleft U$

$$\begin{array}{c}
\text{REFL} \frac{a \in U}{a \triangleleft U} \qquad \text{TRANS} \frac{a \triangleleft U \quad U \triangleleft V}{a \triangleleft V} \qquad \leq\text{-LEFT} \frac{a \leq b \quad b \triangleleft U}{a \triangleleft U} \\
\leq\text{-RIGHT} \frac{a \triangleleft U \quad a \triangleleft V}{a \triangleleft U \downarrow V}
\end{array}$$

Figure 2-1: Rules that a formal cover relation must satisfy.

represents the proposition that the basic open a is covered by the open U (that is, the extent of a is included in the extent of U , though it need not be the case that directly $a \in U$).

This can be extended to specify the entire covering relation, that is, when an open U is covered by another open V . We define (overloading the notation \triangleleft to have either basic opens or opens on the left side)

$$U \triangleleft V \triangleq \forall a \in U. a \triangleleft V,$$

which says that U is covered by V if all of the basic opens comprising U are covered by V . With this definition, we can define the (large) set of opens on A , $\mathcal{O}(A)$, as the type $\mathcal{O}_{\mathbf{B}}(A) \rightarrow \Omega$ where two opens $U, V : \mathcal{O}_{\mathbf{B}}(A) \rightarrow \Omega$ are equivalent if they each cover each other, i.e., $U \triangleleft V$ and $V \triangleleft U$.

The intersection of two basic opens a and b need not necessarily be a basic open itself; however, it must necessarily be the union of basic opens c which are included in both a and b . Therefore, we define the operator^{2 3}

$$\begin{aligned}
&\downarrow : \mathcal{O}_{\mathbf{B}}(A) \rightarrow \mathcal{O}_{\mathbf{B}}(A) \rightarrow (\mathcal{O}_{\mathbf{B}}(A) \rightarrow \Omega) \\
&a \downarrow b \triangleq \bigcup \{c : A \mid c \leq a \text{ and } c \leq b\}.
\end{aligned}$$

The \downarrow operator can be extended to apply to opens just by taking unions, i.e.,

$$U \downarrow V \triangleq \bigcup_{a \in U} \bigcup_{b \in V} a \downarrow b.$$

Having defined \downarrow , we can return to understand the rules on a formal cover described in Figure 2-1. REFL says that a basic open is covered by an open that includes that basic open. $\leq\text{-LEFT}$ is similar, saying that if U covers a basic open b , then it covers a basic open a that is included in b . TRANS describes transitivity of the cover relation: if a is covered by U , and U is covered by V , then a is covered by V . $\leq\text{-RIGHT}$ says

² For any type A and $x : A$, the notation $\{x\}$ denotes the singleton subset, $\lambda y. x \equiv y$.

³ There are many analogies between the construction of the topos of sheaves on a site and the construction of the lattice of opens of a formal space. The definition of \downarrow operator is just a special case of the construction of the product of representable presheaves, where the presheaves are over the category $\mathcal{O}_{\mathbf{B}}(A)$ (considering a preorder as a category). We will refrain from pointing out the other various correspondences.

$$\begin{array}{ccc}
\frac{x \Vdash a \quad a \triangleleft U}{x \Vdash U} \text{ SPLIT} & \frac{}{x \Vdash \top} \text{ MEET-0} & \frac{x \Vdash a \quad x \Vdash b}{x \Vdash a \downarrow b} \text{ MEET-2}
\end{array}$$

Figure 2-2: Rules that a point of a formal space must satisfy.

that if a basic open a is covered by two opens U and V , it must be covered by their intersection.

We should emphasize that these definitions and rules are purely formal: basic opens are just “symbols” in some sense, rather than subsets as they are in classical topology. The rules should match spatial intuition but are not justified by it. The rules can be understood from a *computational* perspective as well, which is most apparent when considering the definition of a point in a formal space.

Definition 2.3. A point x of a formal space A is a subset $(x \Vdash \cdot) : \mathcal{O}_B(A) \rightarrow \Omega$ (read “ x lies in”) satisfying the rules in Figure 2-2. The points of A form a set $\text{Pt}(A)$ where two points x and y are considered equivalent if for all $a : \mathcal{O}_B(A)$, $x \Vdash a$ if and only if $y \Vdash a$.

As with the \downarrow operator, we extend $x \Vdash \cdot$ to operate on opens, rather than just basic opens, by taking a union,⁴

$$x \Vdash U \triangleq \exists a \in U. x \Vdash a.$$

The symbol $\top : \mathcal{O}(A)$ denotes the trivially true propositional function (i.e., subset that is the whole set), and denotes the open that represents the entire space.

Intuitively, $x \Vdash a$ means that the point x lies in the basic open a . Points are described by which opens they lie in, but not every collection of basic opens actually specifies a point. The SPLIT rule says that if x lies in a and if a is covered by U , then x lies in U . This can also be read computationally. If $x \Vdash a$, and we provide a cover U of a (imagine U is very fine), then x *computes* some basic open $b \in U$ such that $x \Vdash b$. We call this computation *splitting* a point with an open cover (following Palmgren [26]). MEET-0 says that a point must lie in the entire space, and MEET-2 says that if x lies in both U and V , then it must lie in their intersection. Another intuitive interpretation is that SPLIT allows one to refine information about a point, MEET-0 gives us *some* information about the point (so there is at least something to refine), and MEET-2 allows us to combine two pieces of information into one.

Definition 2.4. The one-point space $*$ is defined by taking $\mathcal{O}_B(*) \triangleq *$, where $*$ also denotes the preorder with one element tt , and defining the cover relation

$$\text{tt} \triangleleft U \triangleq \text{tt} \in U.$$

⁴The existential quantifier in this definition should be viewed as a dependent pair, where it is possible to *compute* the basic open that exists. This definition agrees with formal topology as formalized in Martin-Löf type theory but differs from Maietti and Sambin [21], who use an existential quantifier that does not guarantee a computational interpretation.

$$\frac{a \triangleleft U}{f^*(a) \triangleleft f^*(U)} \text{ SPLIT} \quad \frac{}{\top \triangleleft f^*(\top)} \text{ MEET-0} \quad \frac{}{f^*(a) \downarrow f^*(b) \triangleleft f^*(a \downarrow b)} \text{ MEET-2} .$$

Figure 2-3: Rules that a continuous map of formal spaces must satisfy.

One can confirm that this covering relation satisfies the necessary rules. Then $\mathcal{O}(\ast) \cong \Omega$, and therefore we can interpret the subset $(x \Vdash \cdot) : \mathcal{O}_B(A) \rightarrow \Omega$ corresponding to a point x of a space A also as an inverse image map $x^* : \mathcal{O}_B(A) \rightarrow \mathcal{O}(\ast)$. In $\mathcal{O}(\ast)$, the \downarrow operator corresponds to conjunction of propositions, and \triangleleft to implication. We can use this correspondence to generalize the definition of points of a space to continuous maps from one space to another:

Definition 2.5. A continuous map f from a formal space Γ to a formal space A , written $f : \Gamma \rightarrow_c A$, is a map $f^* : \mathcal{O}_B(A) \rightarrow \mathcal{O}(\Gamma)$ satisfying the rules in Figure 2-3. The continuous maps from Γ to A form a set $\Gamma \rightarrow_c A$ where two maps f and g are considered equivalent if for all $a : \mathcal{O}_B(A)$, $f^*(a) = g^*(a)$.

2.2.1 Inductively generated formal spaces

Unfortunately, in the general case, it seems impossible to form product spaces for those formal spaces that have no structure beyond their satisfaction of the laws in Figure 2-1. This motivates the definition of *inductively generated formal spaces* by Coquand et al. [7]. This will be the main category that we work with in this thesis. As a category, the inductively generated formal spaces, **FSpc**, have products and form a full subcategory of the more general class of formal spaces.

For inductively generated formal spaces, the cover relation takes a particular form: it is generated by an indexed family of axioms of the form $a \triangleleft U$ for concrete a s and U s.

Definition 2.6. A formal space A is inductively generated if for each $a : \mathcal{O}_B(A)$ there is an index type $I_a : \mathcal{U}$ indexing a family of opens $C_a : I_a \rightarrow (\mathcal{O}_B(A) \rightarrow \Omega)$ such that the covering relation is equivalent to the one inductively generated by the following constructors (which always satisfy the rules in Figure 2-1):

$$\frac{a \in U}{a \triangleleft U} \text{ REFL} \quad \frac{a \leq b \quad b \triangleleft U}{a \triangleleft U} \leq\text{-LEFT}$$

$$\frac{i : I_b \quad a \leq b \quad a \downarrow C_b(i) \triangleleft U}{a \triangleleft U} \text{ COV-AXIOM} .$$

The covering relation generated by an axiom set (I, C) is the *least* covering relation satisfying $a \triangleleft C_a(i)$ for all $a : \mathcal{O}_B(A)$ and $i : I_a$ [7]. The COV-AXIOM rule at once ensures that the required covers are present and that TRANS and \leq -RIGHT hold.

Most importantly, inductively generated formal spaces have products [7] and pullbacks [16] (whereas formal spaces in general may not). Coquand et al. and Vickers demonstrate inductive generation of all spaces used in this thesis [7, 41, 40, 43, 45]. These constructions are critical in enabling computation over these spaces.

2.2.2 Locales

Formal topology describes spaces with a particular base of opens, but often it is easier to express certain constructions instead with *locale theory*, where a space is described without reference to a particular base. Every formal space determines a locale, but (predicatively) one cannot in general construct a formal space from a locale.

There are really two key reasons (which are related) why we deal with formal topology at all, rather than exclusively using locales:

- Formal topology is more concrete and, in practice, is useful for giving concrete methods of constructing spaces (similar to the use of generators and relations for presenting locales).
- Predicatively, a general construction of product spaces and pullbacks does not appear to be possible for locales [7]. These can only be constructed for (inductively generated) formal spaces.

Once a space has been constructed using formal topology, it is sometimes convenient to shift to the language of locale theory, where no fundamental distinction is drawn between the basic opens and opens in general.

Definition 2.7. *A locale A is a distributive lattice $\mathcal{O}(A)$ that has top and bottom elements, \top and \perp , respectively, and that has small joins such that meets distribute over joins:*⁵

$$a \wedge \bigvee_{i:I} b_i = \bigvee_{i:I} a \wedge b_i.$$

Theorem 2.8. *For any formal space A , the preorder $\mathcal{O}(A)$ of opens (defined in Section 2.2) yields a locale.*

Proof sketch. We need to show $\mathcal{O}(A)$ has the requisite operations. We define

$$\begin{aligned} \top &\triangleq \lambda_{-}. \top \\ \perp &\triangleq \lambda_{-}. \perp \\ U \wedge V &\triangleq U \downarrow V \\ \bigvee_{i:I} U_i &\triangleq \bigcup_{i:I} U_i. \end{aligned}$$

⁵Having small joins means having I -indexed joins for any small type $I : \mathcal{U}$. Any lattice of opens $\mathcal{O}(A)$ where $\top \neq \perp$ is necessarily large. Since locale theory is generally impredicative, care must be taken with defining the predicative analogue presented here. Palmgren [25] offers a more careful treatment of predicativity and universes in formal topology.

One can confirm that these operations are well-defined (in fact, monotone) with respect to the preorder on $\mathcal{O}(A)$ and that they satisfy the requisite algebraic laws. \square

We call the lattice $\mathcal{O}(A)$ the *opens* of A , which describes the observable or “affirmable” properties of A [39]. If $U \leq \bigvee_{i:I} V_i$, we call the family $(V_i)_{i:I}$ an *open cover* of U .

We next define the points and continuous maps, which are closely analogous to the version for formal spaces.

Definition 2.9. A point x of a space A is a subset $x \Vdash \cdot : \mathcal{O}(A) \rightarrow \mathcal{U}$ (read “ x lies in”) such that

$$\frac{x \Vdash U \quad U \leq \bigvee_{i:I} V_i}{\exists i : I. x \Vdash V_i} \text{ SPLIT} \quad \frac{}{x \Vdash \top} \text{ MEET-0} \quad \frac{x \Vdash U \quad x \Vdash V}{x \Vdash U \wedge V} \text{ MEET-2}.$$

By the same reasoning as with formal spaces, we can rephrase these rules so that they suggest the rules for continuous maps:

$$\frac{U \leq \bigvee_{i:I} V_i}{x^*(U) \leq x^*\left(\bigvee_{i:I} V_i\right)} \text{ SPLIT} \\ \frac{}{\top \leq x^*(\top)} \text{ MEET-0} \quad \frac{}{x^*(U) \wedge x^*(V) \leq x^*(U \wedge V)} \text{ MEET-2}.$$

However, since SPLIT implies that x^* preserves small joins and is monotone, this means that the reversed inequalities in the MEET rules must hold, and so we can simplify these rules to a simpler and more algebraic definition:

Definition 2.10. A continuous map $f : A \rightarrow_c B$ between locales is a map $f^* : \mathcal{O}(B) \rightarrow \mathcal{O}(A)$ that preserves \top , binary meets, and small joins. The map f^* is known as the inverse image map.

A continuous map $f : A \rightarrow_c B$ transforms covers on B into covers on A . Locales and continuous maps form a category.

Definition 2.11. Two spaces A and B are homeomorphic, written $A \cong B$, if they are isomorphic in the category of continuous maps, that is, if there are continuous maps $f : A \rightarrow_c B$ and $g : B \rightarrow_c A$ such that $g \circ f = \text{id}_A$ as well as $f \circ g = \text{id}_B$.

Definition 2.12 (Specialization order). For two maps $f, g : A \rightarrow_c B$, define $f \leq g$ if for every $U : \mathcal{O}(B)$, $f^*(U) \leq g^*(U)$.

Whenever it is the case that $f \leq g$, g can always behave as f . For instance, if f and g are points, then given an open cover $\top \leq \bigvee_{i:I} U_i$, there is some $i^* : I$ such that $f \Vdash U_{i^*}$. If $f \leq g$, then $g \Vdash U_{i^*}$ as well.

Definition 2.13. An open $U : \mathcal{O}(A)$ is positive, written $\text{Pos}(U)$ if every open cover of U is nonempty, i.e., if for every small family $(V_i)_{i:I}$ such that $U \leq \bigvee_{i:I} V_i$, I is inhabited.

Positivity encodes the notion of being “strictly bigger than \perp ,” and assuming classical logic, $\text{Pos}(U)$ is equivalent to $\neg(U \leq \perp)$.

Proposition 2.14. If $f : \mathcal{O}(A) \rightarrow \mathcal{O}(B)$ preserves joins, then for $U : \mathcal{O}(B)$, $\text{Pos}(f(U))$ implies $\text{Pos}(U)$.

Proof. Suppose $U \leq \bigvee_{i:I} V_i$. Since f preserves joins, it is monotone, so

$$f(U) \leq f\left(\bigvee_{i:I} V_i\right) = \bigvee_{i:I} f(V_i).$$

Since $f(U)$ is positive, I is inhabited by the above cover, and thus U is positive. \square

2.2.3 The computational content of formal topology

A continuous map $f : \Gamma \rightarrow_c A$ (with a point of a space a special case, where $\Gamma \cong *$) is defined by two pieces of data: its inverse image map $f^* : \mathcal{O}(A) \rightarrow \mathcal{O}(\Gamma)$ and a formal proof that f^* preserves small joins, \top , and binary meets. It is reasonable to consider the inverse image map f^* a *specification*, as it describes the observable behavior of the output in terms of observable properties of the input. This accords with the fact that any two continuous maps with the same inverse image maps are considered equal; the formal proofs of structure preservation can be ignored when reasoning, except for the requirement that those formal proofs must exist. However, the formal proofs are necessary for computing concrete results, and so we can consider the formal proof that f^* preserves joins and finitary meets as the *implementation* of the behavior specified by f^* . It is remarkable, then, that the specification is “complete,” uniquely specifying a continuous map (by definition) if it exists.

Suppose we define a function $f : \Gamma \rightarrow_c A$. The map f can be applied to an input $x : \text{Pt}(\Gamma)$ to produce an output point $f(x) : \text{Pt}(A)$. How then does one inspect $f(x)$ to get concrete results about where the output lies within A ?

One computes concrete results in formal topology by splitting a point with an open cover. This open cover may be arbitrarily fine. By combining the SPLIT and MEET-0 rules, we observe that if we present a point y in a space A with an open cover of the whole space $\top \leq \bigvee_{i:I} U_i$, then y can *compute* an $i : I$ such that y lies in U_i . If a point lies in several opens in an open cover, it may return any index corresponding to an open it lies in. In the case of the point $f(x)$ where $f : \Gamma \rightarrow_c A$, f translates a cover of A into a cover of Γ , from which x can then compute some open from that cover of Γ that it lies in.

For instance, a point $x : \text{Pt}(\mathbb{R})$, when given a formal proof of the open cover

$$\top \leq \bigvee_{b:\mathbb{B}} \text{if } b \text{ then } (\cdot < 1) \text{ else } (\cdot > -1),$$

must either return a concrete Boolean **true** together with a proof $x \Vdash \cdot < 1$ or return **false** with a proof $x \Vdash \cdot > -1$. If a point lies in both opens, whether it returns **true** or **false** depends on both the “implementation” of the point (the formal proofs that x satisfies **SPLIT** and **MEET-0**) as well as the formal proof of the cover.

In implementing a real number, one must provide a function which, given any tolerance $\varepsilon > 0$, produces a rational approximation within that tolerance. Imagine two implementations of the real number 0, where the first always returns 0, whereas the second returns $\varepsilon/2$. Consider a formal proof of the above cover of \mathbb{R} that proceeds in the following manner: first, approximate to a tolerance of 1 with the cover $\top \leq \bigvee_{q \in \mathbb{Q}} B_1(q)$. Then, for a given $q \in \mathbb{Q}$, if the lower endpoint $q - 1$ of the resulting approximating interval $B_1(q)$ is at most 0, then prove that $B_1(q) \leq (\cdot < 1)$, and if not, prove $B_1(q) \leq (\cdot > -1)$. With this particular formal covering proof, the first implementation of 0 will return **true**, $(\cdot < 1)$, but the second will return **false**, $(\cdot > -1)$. However, if this cover is applied to the real number 2, the observation will be **false**, no matter the implementation of 2 or the particular proof of the covering.

Note that the **MEET-2** rule does not figure into this notion of computation. We can think of maps $x \Vdash \cdot : \mathcal{O}(A) \rightarrow \mathcal{O}(\ast)$ that satisfy **SPLIT** and **MEET-0** but not **MEET-2** as nondeterministic values (see Chapter 3). But there is a notion of incremental computation where **MEET-2** has a computational meaning.

Suppose we probe a point $x : \mathbf{Pt}(A)$ with a cover $\top \leq \bigvee_{i \in I} U_i$ and learn that $x \Vdash U_{i_\ast}$ for some particular $i_\ast : I$. Then we may decide to further refine our knowledge of x by probing with another open cover $U_{i_\ast} \leq \bigvee_{j \in J} V_j$, or equivalently, $U_{i_\ast} \leq \bigvee_{j \in J} U_{i_\ast} \wedge V_j$, to learn that $x \Vdash U_{i_\ast} \wedge V_{j_\ast}$ for some $j_\ast : J$. This gives a sort of “sequential composition” of splitting with covers. The **MEET-2** rule allows two independent threads of computation to be joined. For instance, we might probe x with an open cover to learn $x \Vdash U$, and independently with another open cover (of the whole space, not just U) to learn that $x \Vdash V$. By **MEET-2**, we can combine what we’ve learned to determine $x \Vdash U \wedge V$, and can continue to refine where x is with covers of $U \wedge V$. For nondeterministic values, where **MEET-2** may not hold, this isn’t satisfied, since (informally) each independent thread of refinement may get a different nondeterministic realization.

2.3 Spaces

2.3.1 Discrete spaces

The simplest kinds of spaces are those that just represent sets. Let **Set** denote the category of sets. The objects of **Set** are types A together with an equivalence relation $=_A : A \rightarrow A \rightarrow \Omega$ on A , and the arrows $f : A \rightarrow B$ are those equivalence-preserving functions, that is, functions $f : A \rightarrow B$ on the underlying types such that for all $a, a' : A$, if $a =_A a'$, then $f(a) =_B f(a')$.

We will show that it is possible to work with sets within the framework of spaces. Precisely, we can define a full and faithful functor **Discrete** : **Set** \rightarrow **FSpc** that exhibits **Set** as the discrete spaces, which form a full subcategory of **FSpc**. Given a set A

(with equivalence relation $=_A$), we construct a formal space whose type of basic opens is A , with the inclusion preorder given by $=_A$. We think of a basic open $a : A$ as representing the subset $(\cdot =_A a)$ of A , which is open in the discrete topology (and hence the preorder defining inclusion of basic opens is discrete). Since every element $a : A$ represents both an open and a point of the discrete space, for $a : A$ we will use the notation $\cdot = a$ to refer to the open, reserving a for the point of the space.

In $\text{Discrete}(A)$, we have $a \triangleleft U$ if and only if there is some $a' : A$ such that $a =_A a'$ and $a' \in U$. We can use this to simplify the rules that a point x of a discrete space or a continuous map $f : \Gamma \rightarrow_c A$ between discrete spaces must satisfy:

$$\begin{array}{c}
\frac{x \Vdash a \quad a =_A b}{x \Vdash b} \text{ SPLIT} \qquad \frac{}{\exists a : A. x \Vdash a} \text{ MEET-0} \qquad \frac{x \Vdash a \quad x \Vdash b}{a =_A b} \text{ MEET-2} \\
\\
\frac{a =_A b}{f^*(a) =_{\Gamma \rightarrow \Omega} f^*(b)} \text{ SPLIT} \qquad \frac{}{\exists a : A. \gamma \in f^*(a)} \text{ MEET-0} \\
\\
\frac{\gamma \in f^*(a) \quad \gamma' \in f^*(b) \quad \gamma =_{\Gamma} \gamma'}{a =_A b} \text{ MEET-2} .
\end{array}$$

We observe that f^* identifies a relation between Γ and A . We read $\gamma \in f^*(a)$ as “ γ is in the preimage of a under f ” (which is equivalent to saying that f maps γ to a). The SPLIT rule says that this relation respects equality on A . MEET-0 says that the relation is total (from Γ to A), and MEET-2 says that the relation respects equality on Γ and also that each input maps to at most one output (up to $=_A$). Accordingly, f^* is a functional relation from the set Γ to the set A , so it is in bijective correspondence with $\Gamma \rightarrow_{\text{Set}} A$, the collection of functions from the set Γ to the set A , meaning that the functor Discrete is full and faithful: continuous maps between discrete spaces are just functions on their underlying sets.

Recall the functor $\text{Pt} : \mathbf{FSpc} \rightarrow \mathbf{Set}$ which takes a space A to its (large) set of points $\text{Pt}(A)$, where two points are considered equal if they lie in the same basic opens. This is right adjoint to Discrete , i.e., $\text{Discrete} \dashv \text{Pt}$, giving a correspondence for a set A and a space B

$$\frac{A \rightarrow_{\text{Set}} \text{Pt}(B)}{\text{Discrete}(A) \rightarrow_c B} .$$

2.3.2 Open subspaces

Given a space A and an open $U : \mathcal{O}(A)$, we can form the open subspace $\{A \mid U\}$ of A by making $\mathcal{O}(\{A \mid U\})$ a quotient of $\mathcal{O}(A)$, identifying opens $P, Q : \mathcal{O}(A)$ in $\{A \mid U\}$ when $P \wedge U = Q \wedge U$. The quotient still defines a space, since the operation $\cdot \wedge U : \mathcal{O}(A) \rightarrow \mathcal{O}(A)$ preserves binary meets and small joins.

2.3.3 The Sierpinski space

The Sierpinski space Σ is fundamental in topology, defining the space of possible “truth values.” Just as the subsets of a set S are in correspondence with functions $S \rightarrow \Omega$, the opens of a space A are in correspondence with the continuous maps $A \rightarrow_c \Sigma$ [39],

$$\mathcal{O}(A) \cong (A \rightarrow_c \Sigma).$$

The basic opens of Σ are $\mathcal{O}_B(\Sigma) \triangleq \mathbb{B}_\leq$, where \mathbb{B}_\leq is \mathbb{B} with the “truth order,” i.e., where **false** is strictly less than **true**. The Sierpinski space has no covering axioms. We have thus defined Σ .

In particular, we have

$$(* \rightarrow_c \Sigma) \cong \mathcal{O}(*) \cong \Omega,$$

so the points of Σ are just the propositions. Given any point x of Σ , the corresponding proposition is $x \Vdash \mathbf{false}$. In particular, there are points \top_Σ and \perp_Σ , where $\top_\Sigma \Vdash \mathbf{false}$ but not $\perp_\Sigma \Vdash \mathbf{false}$.

2.3.4 Disjoint unions (sums)

We have not come across any characterization of sum spaces in **FSpc**, so we describe them here.

From a family of spaces $(A_i)_{i:I}$ parameterized over some index type I^6 , we can form their disjoint union space $\sum_{i:I} A_i$, which is the coproduct of the A_i s in **FSpc**. Intuitively, $\sum_{i:I} A_i$ pastes all the A_i s together, where points from different spaces are not considered near each other. When $I \equiv \mathbb{B}$ (as types), this specializes to binary sums, which we denote with $+$. For instance, we have the homeomorphism in **FSpc**

$$\mathbb{B} \cong * + *,$$

and more generally, for any type I (considered as a set with \equiv as its equivalence relation),

$$\text{Discrete}(I) \cong \sum_{i:I} *.$$

The preorder of basic opens of $\sum_{i:I} A_i$ is the coproduct of the basic opens of the constituent spaces,

$$\mathcal{O}_B\left(\sum_{i:I} A_i\right) \triangleq \sum_{i:I} \mathcal{O}_B(A_i).$$

The preorder relation for the coproduct preorder $\sum_{i:I} \mathcal{O}_B(A_i)$ is generated as the

⁶ In this section, we require the index type I to satisfy uniqueness of identity proofs (UIP), which states that

$$\forall a, b : I. \forall p, q : a \equiv b. p \equiv q.$$

inductive type with the single constructor

$$\frac{a \leq_{A_i} b}{(i, a) \leq_{\sum_{j:I} A_j} (i, b)}.$$

The axioms for $\sum_{i:I} A_i$ are then just a sort of “coproduct” of the axioms of the constituent spaces. For a basic open $(j, a) : \mathcal{O}_B(\sum_{i:I} A_i)$, for each axiom $a \triangleleft_{A_j} U$, we add the axiom

$$(j, a) \triangleleft \text{InDisjunct}_j(U),$$

where $\text{InDisjunct}_j : \mathcal{O}(A_j) \rightarrow \mathcal{O}(\sum_{i:I} A_i)$ is inductively generated by the constructor

$$\frac{a \in U}{(j, a) \in \text{InDisjunct}_j(U)}.$$

Theorem 2.15.

$$A_j \cong \left\{ \sum_{i:I} A_i \mid \text{InDisjunct}_j(\top) \right\}$$

Proof sketch. Define $f : A_j \rightarrow_c \{\sum_{i:I} A_i \mid \text{InDisjunct}_j(\top)\}$ to have its inverse image map $f^* : \mathcal{O}(\{\sum_{i:I} A_i \mid \text{InDisjunct}_j(\top)\}) \rightarrow \mathcal{O}(A_j)$ as the inductive family generated by the constructor

$$\frac{a =_{A_j} a'}{a' \in f^*(j, a)}.$$

Note that with this definition, $f^*(i, b) = \perp$ whenever $i \neq j$.

Define $g : \{\sum_{i:I} A_i \mid \text{InDisjunct}_j(\top)\} \rightarrow_c A_j$ by

$$g^*(a) \triangleq (j, a).$$

It should be clear that f^* and g^* are inverses of each other, so we must just confirm that f^* and g^* indeed define continuous maps, which is mostly a calculational matter. For f , SPLIT is straightforward. We only note that proving that g satisfies MEET-0 reduces to

$$\top \triangleleft \text{InDisjunct}_j(\top),$$

which is exactly the additional covering rule given by the open subspace of f ’s output. \square

Once open embeddings have been defined (Section 3.4.2), this theorem will establish the open embedding $\text{inj}_j : A_j \hookrightarrow \sum_{i:I} A_i$.

2.3.5 Products

We next introduce notation and relevant properties for the construction of product spaces in formal topology⁷.

For a family of spaces $(A_i)_{i:I}$ parameterized over some index type I ⁸, we denote their product in **FSpc** by $\prod_{i:I} A_i$. The key idea characterizing the structure of the product space is that we can make an observation on the product space by choosing a component and making an observation on that component. Since we can only make finitely many observations, any open will make non-trivial observations on only finitely many components.

More precisely, the sub-basic opens of $\prod_{i:I} A_i$ are also given by the coproduct preorder $\sum_{i:I} \mathcal{O}_B(A_i)$. However, we will use the notation $[i \mapsto a]$ for $i : I$ and $a : \mathcal{O}_B(A_i)$ to refer to such a sub-basic open, because it represents the idea that the i th component lies in a . Similarly, we can define the open $[i \mapsto U] : \mathcal{O}(\prod_{i:I} A_i)$ for an arbitrary open $U : \mathcal{O}(A_i)$ as inductively generated by the constructor

$$\frac{a \in U}{[i \mapsto a] \in [i \mapsto U]} .$$

The universal property of product spaces says that for any space Γ and maps $f_i : \Gamma \rightarrow_c A_i$, we can construct a continuous map $g : \Gamma \rightarrow_c \prod_{i:I} A_i$. The inverse image map acts according to⁹

$$g^* : \mathcal{O}_{\text{SB}}\left(\prod_{i:I} A_i\right) \rightarrow \mathcal{O}(\Gamma)$$

$$g^*([i_k \mapsto a_k]) \triangleq f_{i_k}^*(a_k).$$

If we consider the index type $I \equiv \mathbb{B}$, we get the binary products, which we denote $A \times B$ for spaces A and B . We notice that every basic open of $A \times B$ is equivalent to one of the form

$$[\text{true} \mapsto U] \wedge [\text{false} \mapsto V]$$

for $U : \mathcal{O}(A)$ and $V : \mathcal{O}(B)$, and so we may use the notation $U \times V : \mathcal{O}_B(A \times B)$ to represent a basic open of $A \times B$ (also known as an *open rectangle*). Accordingly, every open of $A \times B$ can be represented as a union of open rectangles.

Given $f : \Gamma \rightarrow_c A$ and $g : \Gamma \rightarrow_c B$, we denote by $\langle f, g \rangle : \Gamma \rightarrow_c A \times B$ the “pair” given by the universal property of products.

⁷Vickers [42] provides a full definition and characterization.

⁸ Again we require that I satisfy UIP.

⁹It suffices to only define g^* on its sub-basic opens ($\mathcal{O}_B(\prod_{i:I} A_i)$) since g^* will be required to preserve finitary meets, and the basic opens are just finitary meets of sub-basic opens.

2.3.6 Metric spaces, including \mathbb{R}

It is possible to extend a set with a metric defined on it (such as \mathbb{Q}) to a metrically complete (i.e., Cauchy complete) formal space (such as \mathbb{R}) [41]. This section describes this construction.

Suppose we are given a *metric set*, a set X with a distance metric relation $d : \mathbb{Q}^+ \times X \times X \rightarrow \Omega$, where $d(\varepsilon, x, y)$ indicates the proposition that the distance between x and y is at most ε . The predicate d must in fact define the closed-ball relation for a metric, meaning it must satisfy the following rules:¹⁰

$$\begin{array}{c} \frac{}{d(\varepsilon, x, x)} \text{ REFL} \qquad \frac{d(\varepsilon, x, y)}{d(\varepsilon, y, x)} \text{ SYM} \qquad \frac{d(\delta, x, y) \quad d(\varepsilon, y, z)}{d(\delta + \varepsilon, x, z)} \text{ TRIANGLE} \\[10pt] \frac{\forall \delta : \mathbb{Q}^+. d(\varepsilon + \delta, x, y)}{d(\varepsilon, x, y)} \text{ CLOSED} . \end{array}$$

The basic opens of the metric completion of X (if $X \equiv \mathbb{Q}$, this would be \mathbb{R}) will be the formal balls, the type $\text{Ball}(X)$ defined by the single constructor

$$\frac{\varepsilon : \mathbb{Q}^+ \quad x : X}{B_\varepsilon(x) : \text{Ball}(X)} .$$

We define a relation $<$ on balls that indicates when one ball contains another with “room to spare” all around, and then a relation \leq on balls indicating containment of one formal ball within another:

$$\begin{aligned} B_\delta(x) < B_\varepsilon(y) &\triangleq \exists \gamma : \mathbb{Q}^+. d(\gamma, x, y) \text{ and } \gamma + \delta < \varepsilon \\ B_\delta(x) \leq B_\varepsilon(y) &\triangleq \forall \gamma : \mathbb{Q}^+. B_\delta(x) < B_{\varepsilon+\gamma}(y). \end{aligned}$$

We then define the metric completion space $\mathcal{M}(X)$ of X as having the basic opens $\text{Ball}(X)$, and the following axioms:

$$\frac{\varepsilon : \mathbb{Q}^+}{\top \triangleleft \{B_\varepsilon(q) \mid q : \mathbb{Q}\}} \text{ APPROX} \qquad \frac{}{B_q(\varepsilon) \triangleleft \{B_{\varepsilon'}(q') \mid B_{\varepsilon'}(q') < B_\varepsilon(q)\}} \text{ SHRINK}$$

Definition 2.16. A function $f : X \rightarrow Y$ on metric sets X and Y is k -Lipschitz (for $k : \mathbb{Q}^+$) if for all $x, y : X$ and $\varepsilon : \mathbb{Q}^+$, if $d(\varepsilon, x, y)$ then $d(k\varepsilon, f(x), f(y))$.

We proved in Coq that f can be extended to a continuous map $g : \mathcal{M}(X) \rightarrow_c$

¹⁰ This definition of a closed-ball relation is due to O'Connor [24]. We use it so that our Coq library is compatible with the Coq Repository at Nijmegen (CoRN), which has many constructive results regarding metric spaces.

$\mathcal{M}(Y)$ defined by the inverse image map¹¹

$$g^* : \text{Ball}(Y) \rightarrow \mathcal{O}(\mathcal{M}(X))$$

$$g^*(b) \triangleq \{B_\delta(x) \mid B_{k\delta}(f(x)) < b\}.$$

The fact that f is k -Lipschitz implies that for balls $a, a' : \text{Ball}(X)$, if $a \leq a'$ and $a' \in g^*(b)$, then $a \in g^*(b)$.

Theorem 2.17. *g^* preserves joins, \top , and binary meets.*

Proof sketch. MEET-0 For every formal ball $B_\delta(x)$ of X , there is a ball of Y , $B_{\delta+k\delta}(f(x))$, such that $B_\delta(x) \leq g^*(B_{\delta+k\delta}(f(x)))$. Thus $\top \leq g^*(\top)$.

MEET-2 We are given a ball of Y , $B_{k\delta}(f(x))$, that lies strictly within two other balls of Y , $B_{k\delta}(f(x)) < B_\varepsilon(y)$ and $B_{k\delta}(f(x)) < B_{\varepsilon'}(y')$. We must prove in X that

$$B_\delta(x) \triangleleft g^*(B_\varepsilon(y) \downarrow B_{\varepsilon'}(y')).$$

By the properties of $<$, we can shrink each of the larger balls by just a bit, $\gamma : \mathbb{Q}^+$ while maintaining strict containment, so we get $B_{k\delta}(f(x)) < B_{\varepsilon-\gamma}(y)$ and $B_{k\delta}(f(x)) < B_{\varepsilon'-\gamma}(y')$. Then for some even smaller tolerance $\gamma' < \gamma$, we use the covering axiom $\text{APPROX}(\gamma'/k)$ to derive

$$B_\delta(x) \triangleleft \{B_{\gamma'/k}(x') \mid x' : X\} \downarrow B_\delta(x).$$

This reduces our problem to showing that any ball of X with radius at most γ'/k that is contained within $B_\delta(x)$ is covered by $g^*(B_\varepsilon(y) \downarrow B_{\varepsilon'}(y'))$, which in fact follows by reflexivity, which we will now show. We must only consider the “worst case,” where we have a ball $B_{\gamma'/k}(z)$ such that $B_{\gamma'/k}(z) \leq B_\delta(x)$. We will prove $B_{\gamma'/k}(z) \in g^*(B_\varepsilon(y) \downarrow B_{\varepsilon'}(y'))$. We do so in two steps, showing that $B_{\gamma'/k}(z) \in g^*(B_\gamma(f(z)))$ and that $B_\gamma(f(z)) \in B_\varepsilon(y) \downarrow B_{\varepsilon'}(y')$. The former is almost immediate. The latter is surprisingly intricate, depending on the fact that f is k -Lipschitz as well as using the triangle inequality.

SPLIT We first prove that g^* preserves unary covers, and then proceed by induction on covering axioms. Preserving unary covers reduces to proving that, given balls $b \leq b'$ in Y that $g^*(b) \leq g^*(b')$, which follows from the fact that if $a < b$ and $b \leq b'$, then $a < b'$.

We now proceed by induction on the covering axioms.

In the case of $\text{APPROX}(\varepsilon)$, given balls $B_\delta(x) : \text{Ball}(X)$ and $B_\gamma(y) : \text{Ball}(Y)$ such that $B_\delta(x) \in g^*(B_\gamma(y))$ show that

$$B_\delta(x) \triangleleft g^*(B_\gamma(y) \downarrow \{B_\varepsilon(y') \mid y' : Y\}).$$

¹¹ This definition of the inverse image map is based on a theorem of Vickers (Theorem 4.9 and Remark 4.10 in [41]) extending 1-Lipschitz functions to their metric completions. Vickers omits the proof that the inverse image map defines a continuous map, saying it is “routine to check.”

To do so, we can find some $\alpha : \mathbb{Q}^+$ such that $\alpha < \varepsilon/k$, and cover $B_\delta(x)$ with balls of this radius,

$$B_\delta(x) \triangleleft B_\delta(x) \downarrow \{B_\alpha(x') \mid x' : X\}.$$

Then

$$B_\delta(x) \downarrow \{B_\alpha(x') \mid x' : X\} \triangleleft g^*(B_\gamma(y) \downarrow \{B_\varepsilon(y') \mid y' : Y\})$$

follows by reflexivity.

In the case of SHRINK, the covering in fact follows from reflexivity.

□

The extension of Lipschitz functions from metric sets to their metric completions can be used to define addition $(+ : \mathbb{R} \times \mathbb{R} \rightarrow_c \mathbb{R})$, for instance, in conjunction with facts relating products of spaces and of metric sets, such as [41]

$$\mathcal{M}(X \times Y) \cong \mathcal{M}(X) \times \mathcal{M}(Y).$$

Since $\times : \mathbb{R} \times \mathbb{R} \rightarrow_c \mathbb{R}$ is not Lipschitz, it cannot be defined using the above construction. However, since it is locally Lipschitz, it can be defined as a “gluing” of many Lipschitz maps defined on open subspaces using overlapping pattern matching, as described in Section 4.1.4.

In \mathbb{R} , open intervals and open “rays” (e.g., $(0 < \cdot)$) are evidently open, as they can be described as unions of open balls. For instance, we define

$$\begin{aligned} (0 < \cdot) &: \mathcal{O}(\mathbb{R}) \\ (0 < \cdot) &\triangleq \{B_\varepsilon(x) \mid 0 \leq x - \varepsilon\}. \end{aligned}$$

In $\mathbb{R} \times \mathbb{R}$, the relations $(<), (\neq) : \mathcal{O}(\mathbb{R} \times \mathbb{R})$ are open. We can define $<$ by

$$(<) \triangleq \{(B_\varepsilon(x), B_\delta(y)) \mid x + \varepsilon \leq y - \delta\}$$

and then \neq in terms of that.

2.4 As a programming language

Since **FSpc** is a cartesian monoidal category (i.e., it has well-behaved products), it admits a restricted λ -calculus syntax that compiles to continuous maps, which this thesis uses freely.

Following Escardó [9], we describe the syntax as an inductive family $\vdash : \text{list}(\mathbf{FSpc}) \rightarrow \mathcal{U}$ with constructors

$$\frac{(x : X) \in \Gamma}{\Gamma \vdash x : X} \text{VAR}$$

$$\frac{f : A_1 \times \cdots \times A_n \rightarrow_c B \quad \Gamma \vdash x_1 : A_1 \quad \cdots \quad \Gamma \vdash x_n : A_n}{\Gamma \vdash f(x_1, \dots, x_n) : B} \text{APP},$$

where in the VAR rule $(x : X) \in \Gamma$ denotes a witness that X is a member of the list Γ . Let $\text{Prod} : \text{list}(\mathbf{FSpc}) \rightarrow \mathbf{FSpc}$ compute the product of a list of types.

Theorem 2.18. *Given any expression $\Gamma \vdash e : A$, we can construct a term $\llbracket e \rrbracket : \text{Prod}(\Gamma) \rightarrow_c A$, and conversely, given a function $f : \text{Prod}(\Gamma) \rightarrow_c B$, there is an expression $\Gamma \vdash e : B$.*

Proof. First, we construct a continuous map from syntax by induction on its structure. In the VAR case, we compute from the witness $(x : X) \in \Gamma$ a projection $\llbracket x \rrbracket : \text{Prod}(\Gamma) \rightarrow_c X$. In the APP case, we are given maps $\llbracket x_i \rrbracket : \Gamma \rightarrow_c A_i$ for $i \in 1, \dots, n$. We use these maps to build a map $x : \Gamma \rightarrow_c A_1 \times \cdots \times A_n$ by the universal property for products, i.e., $x \triangleq \langle \llbracket x_1 \rrbracket, \dots, \llbracket x_n \rrbracket \rangle$. Then $f \circ x : \Gamma \rightarrow_c B$.

Conversely, given $f : \text{Prod}(\Gamma) \rightarrow_c B$, we can build the term

$$\Gamma \vdash f(\pi_1, \dots, \pi_n) : B,$$

where each π_i is a variable $\Gamma \vdash \pi_i : \Gamma[i]$. □

This allows us to define a continuous map by introducing variables and applying continuous maps to them. For instance, we can define a map like

$$f : \mathbb{B} \times \mathbb{R} \times \mathbb{R} \rightarrow_c \mathbb{R}$$

$$f(x, y, z) \triangleq \text{if}(x, y, -y) + z \times z$$

rather than having to manually arrange it as a “linear” composition of continuous maps. If we don’t want to give a name to a function such as the one above, we may choose to write it with an “anonymous” lambda,

$$\lambda(x, y, z). \text{if}(x, y, -y) + z \times z.$$

In later chapters, we will expand this syntax. In chapter 3, we will show that similar languages can be define for representing partial or nondeterministic functions. In chapter 4, we will add syntax for pattern matching, and in chapter 5, we will allow quantification over compact spaces (in certain cases).

2.4.1 Opens and Σ -valued maps

For any space A , there is a correspondence $\mathcal{O}(A) \cong A \rightarrow_c \Sigma$ between opens of A and Σ -valued continuous maps on A . We can use this to describe opens via Σ -valued

continuous maps; this is particularly convenient for describing opens of spaces which are finite products. We will use the notation $\{x : A \mid e\}$ where e is a Σ -valued term that may mention x , $x : A \vdash e : \Sigma$, to denote the open subspace $\{A \mid \llbracket e \rrbracket\}$. This thesis will readily conflate opens and Σ -valued continuous maps, implicitly converting between the two.

Chapter 3

Partiality and nondeterminism

The real line \mathbb{R} is connected, meaning that any continuous map $f : \mathbb{R} \rightarrow_c A$ to a discrete set A must be a constant map. In particular, every map $f : \mathbb{R} \rightarrow_c \mathbb{B}$ is constant. The practical implications of connectedness are severe: it is impossible to (continuously) make (non-trivial) discrete decisions over variables that come from connected spaces such as \mathbb{R} .

In order to make decisions, we must give something up. While it is impossible to make discrete decisions on \mathbb{R} that are total and deterministic, we *can* make decisions that are either partial (only defined on some open subspace of the input space) or nondeterministic (could potentially give different answers even when given the exact same input).

This chapter defines notions of partiality and nondeterminism related to continuous maps and characterizes the open maps and open embeddings as those continuous maps having partial and/or nondeterministic inverses. While each of these subjects has been studied in the context of constructive topology, we contribute the integrated characterization relating them, summarized by Figures 3-1 and 3-4.

Recall that points (or continuous maps) must satisfy SPLIT (preserving joins), MEET-0 (preserving \top), and MEET-2 (preserving binary meets). Each rule corresponds to a computational property. The SPLIT rule says that we can refine our knowledge of a point by “splitting” it with an open cover. The MEET-0 rule can be viewed as a rule enforcing totality: viewing \top as the predicate representing the entire space, MEET-0 says that a point must lie in the entire space. If we were not to require the MEET-0 rule, it would be legal to define a relaxed “point” x that does not satisfy *any* properties. Accordingly, we can view values x that do not necessarily satisfy MEET-0 as *partial* points, and continuous maps that fail to satisfy MEET-0 as partial continuous maps.

The MEET-2 rule can be viewed as enforcing determinism. Spatially, the rule says that if a point lies in two opens, it must lie in their intersection. Computationally, it says that it should be possible to reconcile answers given from two independent lines of “questioning” from the SPLIT rule. Relaxed “points” which are not required to satisfy either MEET-0 or MEET-2 can be viewed as both partial and nondeterministic.

We can consider categories with spaces as objects but whose maps are like continuous maps, but the inverse image maps need not necessarily satisfy either MEET-0




Figure 3-1: The lattice of categories representing potentially nondeterministic or partial maps on spaces.

or MEET-2. This gives us a lattice of categories that represent nondeterministic and partial maps, depicted in Figure 3-1¹, where each arrow denotes a faithful (“forgetful”) functor where a particular rule is no longer required for inverse image maps. What is remarkable about these forgetful functors is that they have right adjoints, such that they induce a family of (strong) monads on **FSp** for representing partiality and nondeterminism. The remainder of this chapter will characterize these monads.

To understand how partiality and nondeterminism arise by forgetting MEET-0 and MEET-2, respectively, it is instructive to observe how the categories of partial and nondeterministic values restrict to the discrete spaces. For discrete spaces, we have that $a \triangleleft U$ if and only if there is some a' such that $a = a'$ and $a' \in U$. The SPLIT rule, then, says that an inverse image map must respect the equivalence relations on the relevant sets. For discrete spaces $a \downarrow b$ is inhabited if and only if $a = b$, so MEET-2 reduces to the requirement

$$\frac{x \Vdash a \quad x \Vdash b}{a = b} .$$

Accordingly, a point x of the space given by a discrete set S is a subset $(x \Vdash \cdot)$ of S that has exactly one member: MEET-0 says it has at least one member and MEET-2 says it has at most one member. Of course, such subsets are in bijective correspondence with elements of S , so that points of the discrete space S correspond to elements of S . A continuous map $f : S \rightarrow_c T$ of discrete spaces is a relation $f^* : T \rightarrow S \rightarrow \Omega$ which is functional from S to T , in that for all $s : S$, there is some $t : T$ such that $s \in f^*(t)$ (by MEET-0), and for all $s : S$ and $t, t' : T$, if $s \in f^*(t)$ and $s \in f^*(t')$, then $t = t'$ (by MEET-2). As functional relations are in bijective correspondence with functions, so are continuous maps of discrete spaces $S \rightarrow_c T$ in correspondence with functions $S \rightarrow T$ on their underlying sets.

If we do not require MEET-2, the relaxed continuous maps are the multivalued functions from S to T ; this justifies the sense in which we consider **FSp**_{nd} to be the category of nondeterministic maps. If we instead do not require MEET-0, the relaxed continuous maps are the partial functions from S to T . If we require neither MEET-0 nor MEET-2, these relaxed continuous maps are just those (equality-respecting)

¹nd stands for nondeterministic, p for partial.

relations between S and T .




Figure 3-2: The restriction of the categories of nondeterministic and partial maps of spaces to the discrete spaces.

3.1 Partiality

Partiality allows definition of a continuous map that is only defined on an open subspace of the domain.

Definition 3.1. A partial map f from A to B , written $f : A \xrightarrow{p}_c B$, is a map $f^* : \mathcal{O}(B) \rightarrow \mathcal{O}(A)$ that preserves joins and binary meets, but not necessarily \top .

Example. Reconsider the task of approximately comparing a real number with 0. We can define a partial comparison

$$\text{cmp} : \mathbb{R} \xrightarrow{p}_c \mathbb{B}$$

by only defining a continuous map on the open subspace $\{\mathbb{R} \mid \cdot \neq 0\}$ of \mathbb{R} . We specify its observable behavior with its inverse image map

$$\begin{aligned} \text{cmp}^*(\cdot = \text{true}) &\triangleq \cdot > 0 \\ \text{cmp}^*(\cdot = \text{false}) &\triangleq \cdot < 0. \end{aligned}$$

The inverse image map cmp^* in fact defines a partial map, as cmp^* preserves joins and binary meets, but it is not total, since it fails to preserve \top .

Proof. To confirm cmp^* preserves binary meets: it suffices to check binary meets of basic opens, which is only interesting when they differ, so we confirm

$$\begin{aligned} \text{cmp}^*((\cdot = \text{true}) \wedge (\cdot = \text{false})) &= \text{cmp}^*(\perp) \\ &= \perp \\ &= (\cdot > 0) \wedge (\cdot < 0) \\ &= \text{cmp}^*(\cdot = \text{true}) \wedge \text{cmp}^*(\cdot = \text{false}). \end{aligned}$$

However, cmp^* doesn't preserve \top , since

$$\top \not\leq f^*(\top) = (\cdot < 0) \vee (\cdot > 0).$$

□

Theorem 3.2. *There is a bijective correspondence*

$$\frac{A \xrightarrow{p}_c B}{\sum_{U:\mathcal{O}(A)} (\{A \mid U\} \rightarrow_c B)}$$

between partial maps and continuous maps defined on some open subspace of the domain.

Proof. We will show that any $f : A \xrightarrow{p}_c B$ can be considered as a continuous map $f : \{A \mid f^*(\top)\} \rightarrow_c B$. That is, if we consider opens of A equivalent whenever they are equal when met with $f^*(\top)$, then f^* preserves joins and finitary meets up to this weaker equivalence. Since f^* already preserves joins and binary meets, it certainly preserves them with this weaker equivalence on A . It only remains to prove that f^* preserves \top up to this weaker equivalence, which follows the fact that

$$f^*(\top) =_{\{A \mid f^*(\top)\}} \top$$

since

$$f^*(\top) \wedge f^*(\top) =_A \top \wedge f^*(\top).$$

Now, we will show the opposite direction: Given some $U : \mathcal{O}(A)$ and $f : \{A \mid U\} \rightarrow_c B$, then we can also consider f as a partial map $f : A \xrightarrow{p}_c B$. That is, we know that f^* preserves joins and finitary meets up to equivalence on $\{A \mid U\}$, and must show that f^* preserves joins and binary meets on A . This follows from the fact that the operation $\cdot \wedge U$ of intersection with U preserves joins and binary meets. □

Corollary 3.3. *There is a bijective correspondence*

$$\frac{A \xrightarrow{p}_c *}{\mathcal{O}(A)}$$

between partial maps from A to the one-point space and opens of A .

Proof. This follows from the above theorem and the fact that $*$ is the terminal object in \mathbf{FSpc} , so there is exactly one continuous map from any space to $*$. □

Corollary 3.4. *There is a bijective correspondence*

$$\frac{* \xrightarrow{p}_c A}{\sum_{P:\Omega} (P \rightarrow \text{Pt}(A))}$$

between partial values of A and points of A “guarded” by some arbitrary proposition P .

Proof. This follows from the correspondence between opens of $*$ and propositions. \square

This tells us that to prove that a partial point is actually total could require proving an arbitrary proposition. Assuming classical logic, the proposition P is either true, in which case we indeed have a point of A , or it is false, in which case no point of A was defined.

We can extend specialization order to partial maps in the obvious way. Just as in the continuous case, $f \leq g$ means that g is “at least as defined” as f is.

3.1.1 Lifted spaces

As mentioned previously, the “forgetful” functor from \mathbf{FSpc} to \mathbf{FSpc}_p has a right adjoint, such that there is a (strong) monad $\cdot_\perp : \mathbf{FSpc} \rightarrow \mathbf{FSpc}$ giving a correspondence

$$\frac{A \xrightarrow{p}_c B}{\underline{\underline{A \rightarrow_c B_\perp}}}$$

between continuous maps and their partial counterparts. By this correspondence, we can think of the space B_\perp as the space of partial values. We will call spaces of the form B_\perp *lifted spaces*. Vickers describes them briefly in *Topology via Logic* [39].




Figure 3-3: A schematic diagram of the lifted space A_\perp .

Figure 3-3 depicts a picture that in some sense represents what A_\perp looks like: it has an open subspace $\mathbf{up}(A)$ (defined in Section 3.4.2) that looks exactly like A , and a single point $\perp : \mathbf{Pt}(A_\perp)$ outside of that open subspace, which represents an undefined or nonterminating value.

We will now describe the construction of lifted spaces and their properties. Algebraically, we can think of the construction of lifted spaces as a free construction on the lattice of opens. Suppose we want to construct from a space A its lifted space A_\perp . Then we want to have an operator $\Downarrow : \mathcal{O}(A) \rightarrow \mathcal{O}(A_\perp)$ which should preserve the structure from A that we want to keep: joins and binary meets. We will build such an operator using the structure of A as a formal space.

Since we want \Downarrow to preserve joins, it suffices to define it for basic opens. We will add in a new top element \top as well. Accordingly, the lifted space A_\perp has basic opens generated by the constructors

$$\frac{a : \mathcal{O}_B(A)}{\Downarrow a : \mathcal{O}_B(A_\perp)} \qquad \frac{}{\top : \mathcal{O}_B(A_\perp)}$$

and the preorder on the basic opens is inductively generated by the constructors

$$\frac{a \leq_A b}{\Downarrow a \leq_{A_\perp} \Downarrow b} \qquad \frac{}{a \leq_{A_\perp} \top}.$$

We extend \Downarrow to operate on opens of A so that it preserves joins: for $U : \mathcal{O}(A)$, $\Downarrow U$ is generated by the constructor

$$\frac{a \in U}{\Downarrow a \in \Downarrow U}.$$

Now we describe the covering axioms in A_\perp . We simply copy over the covering axioms from A . For each axiom $a \triangleleft U$ in A , we add the axiom $\Downarrow a \triangleleft \Downarrow U$ in A_\perp .

Proposition 3.5. *Indeed \top is the top element of A_\perp , $\top \leq \top$.*

Proof. It suffices to show that for any basic open $a : \mathcal{O}_B(A_\perp)$, $a \leq \top$, which follows directly from the definition of \leq_{A_\perp} . \square

Lifted spaces satisfy a property much stronger than compactness²:

Proposition 3.6. *If $\top \triangleleft U$, then $\top \in U$.*

Proof. It is equivalent to prove that if $\top \triangleleft U$ then $\top \in U$. We can proceed by induction on the proof of covering. The root node of the proof could not have been the use of an axiom, because there are no axioms for covering \top . Therefore, there must be some basic open $u \in U$ such that $\top \leq u$. Since \top is the largest basic open, it must be that $u \equiv \top$, so $\top \in U$. \square

This means that any covering of A_\perp is necessarily trivial. What this means is that one cannot get (nontrivial) information about a point of A_\perp by splitting the point with a cover. This makes sense, because we require refinement by splitting to occur in finite time, and we want to allow A_\perp to represent points that may be partial. To get useful information from a point in a lifted space, one should *prove* that the point actually lies in $\Downarrow \top$.

Every space A_\perp has a point \perp which has no interesting information: it lies only in the entire space. This corresponds to the undefined or nonterminating value. We

² A space is compact if whenever $\top \triangleleft U$, then there is a (Kuratowski-) finite subset K of U such that $\top \triangleleft K$.

describe the basic opens that \perp lies in by the inverse image map

$$\begin{aligned}\perp \Vdash \cdot &: \mathcal{O}_B(A) \rightarrow \mathcal{O}(1) \\ \perp \Vdash \top &\triangleq \top \\ \perp \Vdash \Downarrow a &\triangleq \perp.\end{aligned}$$

Proposition 3.7. \perp *indeed describes a point of A_\perp .*

Proof. **MEET-0** Follows from $\perp \Vdash \top$.

MEET-2 Trivially satisfied, since there are no nontrivial intersections of basic opens that \perp lies in.

SPLIT Follows from Proposition 3.6. □

The idea that \perp is the “least defined” point can be made formal, in that it is minimal in terms of specialization order.

Proposition 3.8. *For any (generalized) point $x : \Gamma \rightarrow_c A_\perp$, $\perp \leq x$ (in terms of specialization order).*

Proof. Since \perp lies only in \top , it suffices to show that $\top \leq x^*(\top)$. However, since $\top = \top$, this is equivalent to $\top \leq x^*(\top)$, which is true of every continuous map. □

Having established some intuition about what lifted spaces represent, we can confirm they fulfill their original purpose:

Theorem 3.9. *There is the bijective correspondence*

$$\frac{A \xrightarrow{p}_c B}{A \rightarrow_c B_\perp}.$$

Proof. Given $f : A \xrightarrow{p}_c B$, we define $g : A \rightarrow_c B_\perp$ by

$$\begin{aligned}g^* &: \mathcal{O}_B(B_\perp) \rightarrow \mathcal{O}(A) \\ g^*(\top) &\triangleq \top \\ g^*(\Downarrow b) &\triangleq f^*(b).\end{aligned}$$

The map g^* preserves \top since $g^*(\top) = \top$ and preserves joins and binary meets since f^* does.

Conversely, given $g : A \rightarrow_c B_\perp$, we define $f : A \xrightarrow{p}_c B$ by

$$\begin{aligned}f^* &: \mathcal{O}(B) \rightarrow \mathcal{O}(A) \\ f^*(U) &\triangleq g^*(\Downarrow U).\end{aligned}$$

The map f^* preserves joins and binary meets since g^* and \Downarrow each do. □

This correspondence is natural in A and B , giving an adjunction and thus a monad. It remains to show that \cdot_\perp in fact defines a *strong* monad. Its tensorial strength $s : A \times B_\perp \xrightarrow{p}_c A \times B$ is defined by the inverse image map

$$\begin{aligned} s^* : \mathcal{O}_B(A \times B) &\rightarrow \mathcal{O}(A \times B_\perp) \\ s^*(a \times b) &\triangleq a \times \Downarrow b. \end{aligned}$$

The inverse image map s^* preserves joins and binary meets since \Downarrow does. We claim (but do not prove) that s satisfies the strong monad laws.

3.2 Nondeterminism

Nondeterminism allows the definitions of programs whose observable behavior might depend on the *exact implementation* of their inputs (specifically, the formal proofs that their inputs preserve joins and finitary meets). Rather than viewing such behavior as breaking the abstraction provided by the equivalence relation on points (since points that lie in the same opens may be treated differently), we can instead choose to maintain this abstraction and view such behavior as fundamentally nondeterministic.

For instance, we can perform a nondeterministic approximate comparison of a real number with 0.

Example. Fix some error tolerance parameter $\varepsilon > 0$. We may define a total but nondeterministic approximate comparison with 0

$$\text{cmp} : \mathbb{R} \xrightarrow{nd}_c \mathbb{B},$$

allowing error up to ε , by specifying its observable behavior with the inverse image map

$$\begin{aligned} \text{cmp}^*(\cdot = \text{true}) &\triangleq \cdot > -\varepsilon \\ \text{cmp}^*(\cdot = \text{false}) &\triangleq \cdot < \varepsilon. \end{aligned}$$

We can confirm that cmp^* in fact defines a nondeterministic map, as it preserves joins and \top , but fails to preserve binary meets.

Proof. Since cmp 's codomain is discrete, it trivially satisfies SPLIT. First, we confirm it preserves \top :

$$\begin{aligned} \text{cmp}^*(\top) &= \text{cmp}^*(\cdot = \text{true}) \vee \text{cmp}^*(\cdot = \text{false}) \\ &= (\cdot > -\varepsilon) \vee (\cdot < \varepsilon) \\ &= \top. \end{aligned}$$

However, it fails to preserve binary meets, since

$$\text{cmp}^*((\cdot = \text{true}) \wedge (\cdot = \text{false})) = \text{cmp}^*(\perp) = \perp$$

but

$$\begin{aligned} \text{cmp}^*(\cdot = \text{true}) \wedge \text{cmp}^*(\cdot = \text{false}) &= (\cdot > -\varepsilon) \wedge (\cdot < \varepsilon) \\ &= -\varepsilon < \cdot < \varepsilon, \end{aligned}$$

which is a positive open (thus, not \perp). \square

As we might expect, we can perform nondeterministic joins of (generalized) points in \mathbf{FSpc}_{nd} .

Proposition 3.10. *For any inhabited index type I and space A , there is the nondeterministic map $\sqcup : \prod_{i:I} A \xrightarrow{nd}_c A$ defined by*

$$\begin{aligned} (\sqcup)^* : \mathcal{O}(A) &\rightarrow \mathcal{O}\left(\prod_{i:I} A\right) \\ (\sqcup)^*(a) &\triangleq \bigvee_{i:I} [i \mapsto a]. \end{aligned}$$

Proof. Let $i^* : I$ witness that I is inhabited. Then this map satisfies MEET-0 since

$$\top \leq [i^* \mapsto \top] \leq \bigvee_{i:I} [i \mapsto \top] \leq (\sqcup)^*(\top).$$

Given any $a \triangleleft_A U$, we have

$$(\sqcup)^*(a) = \bigvee_{i:I} [i \mapsto a] \leq \bigvee_{i:I} [i \mapsto U] = (\sqcup)^*(U),$$

so SPLIT is satisfied as well. \square

Computationally, we notice that $\bigsqcup_{i:I} x_i$ (for $x_i : \Gamma \xrightarrow{nd}_c A$) will behave as x_{i^*} .

We can extend specialization order to nondeterministic maps too, in the obvious way, and in this case $f \leq g$ represents that g can potentially behave as f (though it may also behave differently, as well).

3.2.1 Nondeterministic powerspaces

As mentioned previously, the “forgetful” functor from \mathbf{FSpc} to \mathbf{FSpc}_{nd} has a right adjoint, such that there is a monad $\mathcal{P}_\diamond^+ : \mathbf{FSpc} \rightarrow \mathbf{FSpc}$ giving a correspondence

$$\frac{A \xrightarrow{nd}_c B}{A \rightarrow_c \mathcal{P}_\diamond^+(B)}$$

between continuous maps and their nondeterministic counterparts [40]. By this correspondence, we can think of the space $\mathcal{P}_\diamond^+(B)$ as the space of nondeterministic values.

Spaces of the form $\mathcal{P}_\diamond^+(B)$ are known as *positive lower powerspaces* (also known as positive or inhabited Hoare powerlocales).

There is a map $\diamond : \mathcal{O}(A) \rightarrow \mathcal{O}(\mathcal{P}_\diamond^+(A))$ (read “possibly”) that can be intuitively understood in the following sense: if $U : \mathcal{O}(A)$ is interpreted as a property of points of A , and a point $s : * \rightarrow_c \mathcal{P}_\diamond^+(A)$ (where $*$ is the one-point space) is interpreted as a subspace of A , $\diamond U : \mathcal{O}(\mathcal{P}_\diamond^+(A))$ holds of s if U holds of some point in s .

Just like \downarrow for lifted spaces, the \diamond operator preserves all the structure in the lattice $\mathcal{O}(A)$ that ought to be preserved in $\mathcal{O}(\mathcal{P}_\diamond^+(A))$: that is, it preserves joins and \top . This implies a particular fact that is critical in the computational properties of the nondeterministic powerspaces:

Theorem 3.11. *For every cover $\top \leq \bigvee_{i:I} V_i$ in A there is a cover $\top \leq \bigvee_{i:I} \diamond V_i$ in $\mathcal{P}_\diamond^+(A)$.*

Proof. Suppose $\top \leq \bigvee_{i:I} V_i$ in A . Since \diamond preserves \top and joins,

$$\bigvee_{i:I} \diamond V_i = \diamond \left(\bigvee_{i:I} V_i \right) = \diamond \top = \top.$$

□

The previous theorem can be understood computationally as allowing simulation of some nondeterministic result.

Classically, the points of $\mathcal{P}_\diamond^+(B)$ are in correspondence with closed nonempty subspaces of B . Constructively, these subspaces also are *overt* [40], which is helpful for some computational tasks.

The specialization order on $\mathcal{P}_\diamond^+(A)$ corresponds to subspace inclusion of possible values.

Vickers [40] describes how to construct this powerspace, as well the powerspace for $\mathbf{FSpc}_{nd,p}$, predicatively within \mathbf{FSpc} .

Vickers [40] shows that \mathcal{P}_\diamond^+ defines a monad, but it remains to show that \mathcal{P}_\diamond^+ in fact defines a *strong* monad. Its tensorial strength $s : A \times \mathcal{P}_\diamond^+(B) \xrightarrow{nd}_c A \times B$ is defined by the inverse image map

$$\begin{aligned} s^* : \mathcal{O}_B(A \times B) &\rightarrow \mathcal{O}(A \times \mathcal{P}_\diamond^+(B)) \\ s^*(a \times b) &\triangleq a \times \diamond b. \end{aligned}$$

The inverse image map s^* preserves joins and \top since \diamond does.

3.3 Both partiality and nondeterminism

The category $\mathbf{FSpc}_{nd,p}$, where both nondeterminism and partiality are allowed, is quite similar to \mathbf{FSpc}_{nd} , where only nondeterminism is allowed.

In $\mathbf{FSpc}_{nd,p}$, inverse image maps are required only to preserve joins, and need not preserve any meets. Accordingly, $\mathbf{FSpc}_{nd,p}$ is equivalent to the category of suplattices, which has been studied in various contexts [37, 40, 15, 4].

3.3.1 Lower powerspaces

As mentioned previously, the “forgetful” functor from \mathbf{FSpc} to $\mathbf{FSpc}_{nd,p}$ has a right adjoint, such that there is a monad $\mathcal{P}_\diamond : \mathbf{FSpc} \rightarrow \mathbf{FSpc}$ giving a correspondence

$$\frac{A \xrightarrow{nd,p}_c B}{A \rightarrow_c \mathcal{P}_\diamond(B)}$$

between continuous maps and their nondeterministic and partial counterparts [40]. By this correspondence, we can think of the space $\mathcal{P}_\diamond(B)$ as the space of nondeterministic and partial values. Spaces of the form $\mathcal{P}_\diamond(B)$ are known as *lower powerspaces* (also known as Hoare powerlocales).

There is a map $\diamond : \mathcal{O}(B) \rightarrow \mathcal{O}(\mathcal{P}_\diamond(B))$ (read “possibly”) which distributes over joins but not necessarily meets. In particular,

$$\diamond U \leq \bigvee_{i \in I} \diamond V_i$$

holds in $\mathcal{O}(\mathcal{P}_\diamond(B))$ whenever $U \leq \bigvee_{i \in I} V_i$ in $\mathcal{O}(B)$.

The monad \mathcal{P}_\diamond is, like \mathcal{P}_\diamond^+ , a strong monad, and its tensorial strength is almost the same: $s : A \times \mathcal{P}_\diamond(B) \xrightarrow{nd,p}_c A \times B$ is defined by the inverse image map

$$\begin{aligned} s^* : \mathcal{O}_B(A \times B) &\rightarrow \mathcal{O}(A \times \mathcal{P}_\diamond(B)) \\ s^*(a \times b) &\triangleq a \times \diamond b. \end{aligned}$$

The inverse image map s^* preserves joins since \diamond does.

3.4 Open maps and open embeddings

In pattern matching for functional programming, one may pattern match on an inductive type by checking whether it has the form of a particular constructor applied to some argument (i.e., it is in the image of the map defined by a particular constructor). The analogue of constructors for overlapping pattern matching will be the open maps and the open embeddings.

In this section, we will observe how *open maps* can be viewed as continuous maps that have nondeterministic and partial inverses, and *open embeddings* can be viewed as those that have partial (but deterministic) inverses. This is what makes the open maps and open embeddings relevant for constructing the patterns in overlapping pattern matching. We will also prove some properties of open maps and open embeddings that ensure that they will behave well as patterns.

3.4.1 Open maps

Definition 3.12 ([35]). A continuous map $f : A \rightarrow_c B$ is an open map if the inverse image map $f^* : \mathcal{O}(B) \rightarrow \mathcal{O}(A)$ has a left adjoint $f_! : \mathcal{O}(A) \rightarrow \mathcal{O}(B)$, called the direct image map, that satisfies the Frobenius law,

$$f_!(U \wedge f^*(V)) = f_!(U) \wedge V.$$

That is, f is an open map if the *image* of any open in A is open in B ; the direct image map provides this mapping.

Example. Identity maps $\text{id} : A \rightarrow_c A$ are open maps, and the composition of open maps is an open map.

Proof. We have $\text{id}_!(U) = U$, which clearly satisfies $\text{id}_! \dashv \text{id}^*$ since they are all identity maps. It is also immediate that the Frobenius law holds for $\text{id}_!$. Given open maps $f : A \rightarrow_c B$ and $g : B \rightarrow_c C$, we claim that the direct image map is $(g \circ f)_! = g_! \circ f_!$. We first confirm the adjunction:

$$\begin{aligned} (g \circ f)_!(U) \leq V &\iff g_!(f_!(U)) \leq V \\ &\iff f_!(U) \leq g^*(V) && (g_! \dashv g^*) \\ &\iff U \leq f^*(g^*(V)) && (f_! \dashv f^*) \\ &\iff U \leq (g \circ f)^*V. \end{aligned}$$

It only remains to confirm the Frobenius law for the composition $g_! \circ f_! : \mathcal{O}(A) \rightarrow \mathcal{O}(C)$:

$$\begin{aligned} (g \circ f)_!(U \wedge (g \circ f)^*(V)) &= g_!(f_!(U \wedge f^*(g^*(V)))) \\ &= g_!(f_!(U) \wedge g^*(V)) && (\text{Frobenius law for } f) \\ &= g_!(f_!(U)) \wedge V \\ &= (g \circ f)_!(U) \wedge V. && (\text{Frobenius law for } g) \end{aligned}$$

□

Proposition 3.13. In the pullback square where p and q are open maps,

$$\begin{array}{ccc} A \times_X B & \xrightarrow{\theta} & A \\ \downarrow \varphi & \lrcorner & \downarrow p \\ B & \xrightarrow{q} & X \end{array},$$

define $f : A \times_X B \rightarrow_c X$ by $f = p \circ \theta = q \circ \varphi$. Then

$$f_!(\top) = p_!(\top) \wedge q_!(\top).$$

Proof. Note that for any $U : \mathcal{O}(B)$, $p^*(q_!(U)) = \theta_!(\varphi^*(U))$ (proof in section 5.2 of [28]). Then

$$\begin{aligned}
p_!(\top) \wedge q_!(\top) &= p_!(\top \wedge p^*(q_!(\top))) && \text{(Frobenius reciprocity)} \\
&= p_!(p^*(q_!(\top))) \\
&= p_!(\theta_!(\varphi^*(\top))) && \text{(the earlier equality)} \\
&= p_!(\theta_!(\top)) && (\varphi^* \text{ preserves } \top) \\
&= f_!(\top).
\end{aligned}$$

□

Proposition 3.14. *Parallel composition of open maps yields an open map. That is, given $f : A \rightarrow_c B$ and $g : X \rightarrow_c Y$ open, $f \otimes g : A \times X \rightarrow_c B \times Y$ is open.*

Proof. We claim that the direct image map operates on basic opens (which are open rectangles of $A \times X$) by $(f \otimes g)_!(a \times x) = f_!(a) \times g_!(x)$ (this extends to all opens by taking joins). We confirm the adjunction (using the fact that every open is a join of basic opens):

$$\begin{aligned}
(f \otimes g)_! \left(\bigvee_{i:I} a_i \times x_i \right) \leq \bigvee_{j:J} (b_j \times y_j) &\iff \bigvee_{i:I} f_!(a_i) \times g_!(x_i) \leq \bigvee_{j:J} (b_j \times y_j) \\
&\iff \forall i : I. \exists j : J. f_!(a_i) \times g_!(x_i) \leq b_j \times y_j \\
&\iff \forall i : I. \exists j : J. (f_!(a_i) \leq b_j) \text{ and } g_!(x_i) \leq y_j \\
&\iff \forall i : I. \exists j : J. (a_i \leq f^*(b_j)) \text{ and } x_i \leq g^*(y_j) \\
&\iff \forall i : I. \exists j : J. a_i \times x_i \leq f^*(b_j) \times g^*(y_j) \\
&\iff \bigvee_{i:I} a_i \times x_i \leq \bigvee_{j:J} f^*(b_j) \times g^*(y_j) \\
&\iff \bigvee_{i:I} a_i \times x_i \leq (f \otimes g)^* \left(\bigvee_{j:J} b_j \times y_j \right).
\end{aligned}$$

It suffices to confirm the Frobenius law holds on basic opens [35]:

$$\begin{aligned}
(f \otimes g)_!((a \times x) \wedge (f \otimes g)^*(b \times y)) &= (f \otimes g)_!((a \times x) \wedge (f^*(b) \times g^*(y))) \\
&= (f \otimes g)_!((a \wedge f^*(b)) \times (x \wedge g^*(y))) \\
&= f_!(a \wedge f^*(b)) \times g_!(x \wedge g^*(y)) \\
&= (f_!(a) \wedge b) \times (g_!(x) \wedge y) \\
&\quad \text{(Frobenius law for } f \text{ and } g) \\
&= (f_!(a) \times g_!(x)) \wedge (b \times y) \\
&= (f \otimes g)_!(a \times x) \wedge (b \times y).
\end{aligned}$$

□

We will now describe some facts that relate the open maps to nondeterministic maps.

Proposition 3.15. *For any open map $p : A \rightarrow_c B$, there is a (potentially) partial and nondeterministic inverse map $p^{-1} : B \xrightarrow{nd,p}_c A$ defined by*

$$\begin{aligned} p^{-1*} : \mathcal{O}(B) &\rightarrow \mathcal{O}(A) \\ p^{-1*}(U) &\triangleq p_!(U). \end{aligned}$$

Proof. We only must prove that $p_!$ preserves joins, which follows from the fact that $p_!$ is a left adjoint (to p^*). \square

The map $\{\cdot\} : A \rightarrow_c \mathcal{P}_\diamond(A)$, defined as the “return” operation of the \mathcal{P}_\diamond monad, is open, with a direct image map taking opens U of A to $\diamond U$ in $\mathcal{P}_\diamond(A)$.

3.4.2 Open embeddings

Proposition 3.16. *An open map $f : A \rightarrow_c B$ factors through its direct image $f_!(\top)$, i.e., there is an \tilde{f} such that the following diagram commutes:*

$$\begin{array}{ccc} A & \xrightarrow{\tilde{f}} & \{B \mid f_!(\top)\} \\ & \searrow f & \downarrow \iota_{[f_!(\top)]} \\ & & B \end{array}$$

Proof. This statement is equivalent to that for all $U : \mathcal{O}(B)$, $f^*(U) \leq f^*(U \wedge f_!(\top))$. This is indeed the case:

$$\begin{aligned} f^*(U \wedge f_!(\top)) &= f^*(U) \wedge f^*(f_!(\top)) && (f^* \text{ preserves meets}) \\ &\geq f^*(U) \wedge \top && (f^* \circ f_! \text{ inflationary}) \\ &= f^*(U). \end{aligned}$$

\square

Definition 3.17. *A map $f : A \rightarrow_c B$ is an open embedding (or open inclusion) if A is homeomorphic to its image under f in B , that is, if there is an open subspace $U : \mathcal{O}(B)$ and homeomorphism $\tilde{f} : A \rightarrow_c \{B \mid U\}$ such that the following diagram commutes:*

$$\begin{array}{ccc} A & \xrightleftharpoons[\tilde{f}^{-1}]{\tilde{f}} & \{B \mid U\} \\ & \searrow f & \downarrow \iota_{[U]} \\ & & B \end{array}$$

The notation $f : A \hookrightarrow B$ indicates that the continuous map $f : A \rightarrow_c B$ is an open embedding.

Theorem 3.18. *A map $f : A \rightarrow_c B$ is an open embedding if and only if it is an open map and its direct image map $f_!$ preserves binary meets.*

Proof. Suppose $f : A \rightarrow_c B$ is an open embedding that factors through $\{B \mid U\}$, and let \tilde{f} and \tilde{f}^{-1} be the maps as in the above diagram. Then its direct image map is given by

$$\begin{aligned} f_! : \mathcal{O}(A) &\rightarrow \mathcal{O}(B) \\ f_!(V) &\triangleq \tilde{f}^{-1*}(V) \wedge U. \end{aligned}$$

We now confirm that $f_! \dashv f^*$. We have

$$\begin{aligned} f_!(V) \leq W &\iff \tilde{f}^{-1*}(V) \wedge U \leq W \\ &\iff \tilde{f}^{-1*}(V) \wedge \tilde{f}^{-1*}(\tilde{f}^*(U)) \leq W \\ &\iff \tilde{f}^{-1*}(V \wedge \tilde{f}^*(U)) \leq W \\ &\iff \tilde{f}^{-1*}(V \wedge \top) \leq W \\ &\iff \tilde{f}^{-1*}(V) \leq W \\ &\iff V \leq \tilde{f}^*(W) \end{aligned}$$

Moreover, $f_!$ preserves meets:

$$\begin{aligned} f_!(V \wedge W) &= \tilde{f}^{-1*}(V \wedge W) \wedge U \\ &= \tilde{f}^{-1*}(V) \wedge \tilde{f}^{-1*}(W) \wedge U \\ &= (\tilde{f}^{-1*}(V) \wedge U) \wedge (\tilde{f}^{-1*}(W) \wedge U) \\ &= f_!(V) \wedge f_!(W). \end{aligned}$$

We also confirm $f_!$ satisfies the Frobenius law:

$$\begin{aligned} f_!(V \wedge f^*(W)) &= f_!(V) \wedge f_!(f^*(W)) && (f_! \text{ preserves meets}) \\ &= f_!(V) \wedge \tilde{f}^{-1*}(f^*(W)) \wedge U \\ &= f_!(V) \wedge \tilde{f}^{-1*}(f^*(W)) && (f_!(V) \text{ already at most } U) \\ &= f_!(V) \wedge \tilde{f}^{-1*}(\tilde{f}^*(W \wedge U)) \\ &= f_!(V) \wedge W \wedge U \\ &= f_!(V) \wedge W && (f_!(V) \text{ already at most } U) \end{aligned}$$

Thus an open embedding is an open map with a meet-preserving direct image map. We now prove the converse. Given an open map $f : A \rightarrow_c B$ with a meet-preserving direct image map $f_!$, we claim that $A \cong \{B \mid f_!(\top)\}$. By Proposition 3.16, we already have a continuous map $\tilde{f} : A \rightarrow_c \{B \mid f_!(\top)\}$. We define $g : \{B \mid f_!(\top)\} \rightarrow_c A$ by

$$\begin{aligned} g^* : \mathcal{O}(A) &\rightarrow \mathcal{O}(\{B \mid f_!(\top)\}) \\ g^*(V) &\triangleq f_!(V) \wedge f_!(\top). \end{aligned}$$

We claim g^* indeed defines a continuous map. It preserves joins and binary meets since it is the composition of $f_!$ and $\cdot \wedge f_!(\top)$, both of which preserve joins and binary meets, so it suffices to show that $g^*(\top) = \top$, which is indeed the case, as

$$g^*(\top) = f_!(\top) \wedge f_!(\top) = f_!(\top)$$

which is equivalent to \top in $\{B \mid f_!(\top)\}$. \square

Example. Identity maps $\text{id} : A \hookrightarrow A$ are open embeddings, and the composition of open embeddings is an open embedding.

Proof. We have $\text{id}_!(U) = U$, which preserves binary meets. Given open maps $f : A \rightarrow_c B$ and $g : B \rightarrow_c C$, the composition $g_! \circ f_! : \mathcal{O}(A) \rightarrow \mathcal{O}(C)$ preserves binary meets since $f_!$ and $g_!$ both do. \square

Proposition 3.19. *Parallel composition of open embeddings yields an open embedding. That is, given $f : A \hookrightarrow B$ and $g : X \hookrightarrow Y$ open embeddings, $f \otimes g : A \times X \hookrightarrow B \times Y$ is an open embedding.*

Proof. It suffices to confirm that the direct image preserves binary meets, and it suffices to check this by only checking the basic opens, which in this case are open rectangles:

$$\begin{aligned} (f \otimes g)_!((a_1 \times x_1) \wedge (a_2 \times x_2)) &= (f \otimes g)_!((a_1 \wedge a_2) \times (x_1 \wedge x_2)) \\ &= f_!(a_1 \wedge a_2) \times g_!(x_1 \wedge x_2) \\ &= (f_!(a_1) \wedge f_!(a_2)) \times (g_!(x_1) \wedge g_!(x_2)) \\ &\quad (f_! \text{ and } g_! \text{ preserve binary meets}) \\ &= (f_!(a_1) \times g_!(x_1)) \wedge (f_!(a_2) \times g_!(x_2)) \\ &= (f \otimes g)_!(a_1 \times x_1) \wedge (f \otimes g)_!(a_2 \times x_2) \end{aligned}$$

\square

Proposition 3.20. *The pullback of an open embedding is an open embedding.*

Proof. Since (up to homeomorphism) an open embedding is just the inclusion of an open subspace, it suffices to prove that the pullback of the inclusion of an open subspace is an open embedding. Given the open inclusion $\iota[U] : \{X \mid U\} \hookrightarrow X$, for any map $f : A \rightarrow_c X$ we have the pullback square

$$\begin{array}{ccc} \{A \mid f^*(U)\} & \xrightarrow{\iota[f^*(U)]} & A \\ \downarrow f_{!U} \lrcorner & & \downarrow f \\ \{X \mid U\} & \xleftarrow{\iota[U]} & X \end{array},$$

which expresses the fact that the preimage of opens are open. It is possible to confirm that the diagram commutes and that $\{A \mid f^*(U)\}$ satisfies the universal property of pullbacks. \square

The open embeddings are closely related to the partial maps.

Proposition 3.21. *Given an open embedding $f : A \hookrightarrow B$, the “inverse map” $f^{-1} : B \xrightarrow{nd,p} A$ that any open map has is in fact deterministic, i.e., $f^{-1} : B \xrightarrow{p} A$.*

Proof. This follows directly from the fact that $f_!$ preserves joins and binary meets. \square

We will now observe that the map $\mathbf{up} : A \rightarrow_c A_\perp$ that is the “return” operation of the lifting monad \cdot_\perp is an open embedding. This allows us to view A as an open subspace of A_\perp . We can define \mathbf{up} via the inverse and direct image maps

$$\begin{aligned}\mathbf{up}^* &: \mathcal{O}_B(A_\perp) \rightarrow \mathcal{O}(A) \\ \mathbf{up}^*(\top) &\triangleq \top \\ \mathbf{up}^*(\Downarrow a) &\triangleq a\end{aligned}$$

$$\begin{aligned}\mathbf{up}_! &: \mathcal{O}(A) \rightarrow \mathcal{O}(A_\perp) \\ \mathbf{up}_!(U) &\triangleq \Downarrow U.\end{aligned}$$

Proposition 3.22. *This indeed defines an open embedding $\mathbf{up} : A \hookrightarrow A_\perp$.*

Proof. We first confirm \mathbf{up} defines a continuous map:

MEET-0

$$\mathbf{up}^*(\top) \geq \mathbf{up}^*(\top) = \top.$$

MEET-2 Follows from the facts

$$\begin{aligned}\top \Downarrow \Downarrow a &= \Downarrow a \\ \Downarrow a \Downarrow \Downarrow b &= \Downarrow(a \downarrow b).\end{aligned}$$

SPLIT It suffices to check every axiom of A_\perp : every axiom is of the form $\Downarrow a \triangleleft_{A_\perp} \Downarrow U$, where $a \triangleleft_A U$. So we must confirm

$$\mathbf{up}^*(\Downarrow a) \triangleleft \mathbf{up}^*(\Downarrow U),$$

which is equivalent to $a \triangleleft U$, which we know by assumption.

We will now confirm that $\mathbf{up}_! = \Downarrow$ preserves joins and binary meets. We observe that it preserves binary meets since

$$\Downarrow a \downarrow \Downarrow b = \Downarrow(a \downarrow b).$$

It preserves joins since if $a \triangleleft_A U$ then $\Downarrow a \triangleleft_{A_\perp} \Downarrow U$.

Finally, it is straightforward to confirm that $\mathbf{up}_! \dashv \mathbf{up}^*$ and that the Frobenius law holds. \square

3.5 Summary

The open maps and the open embeddings form part of lattice of categories of continuous maps whose *inverses* are potentially partial and nondeterministic; the inverse of an open map is partial and nondeterministic, whereas the inverse of an open embedding is partial, but deterministic. This will make the open maps and open embeddings relevant for pattern matching to produce either continuous maps or nondeterministic maps on spaces. Figure 3-4 depicts this lattice alongside the original one.




Figure 3-4: A summary of various categories of morphisms on formal spaces and functors relating them. “Forgetful” functors have tails and are unnamed. Functors named “ -1 ” denote inverse functions determined by direct image maps. All paths in the diagram commute except for those involving the (dashed) forgetful functor $o \rightarrow c$.

Usually, when pattern matching, we do not expect that a single pattern covers the entire input space, and accordingly we expect that the “inverse” map of a particular pattern should be partial, rather than total. This explains why we focused on the open maps and open embeddings, whose inverses are necessarily partial. The greyed categories in Figure 3-4 describe those categories of maps whose inverses are in fact total. The inverses of surjective open maps are total but nondeterministic, and the inverses of homeomorphisms are just regular continuous maps (total and deterministic).

3.6 Variants of the programming language

We can extend the programming language for **FSpc** described in Section 2.4 to programming languages for nondeterministic or partial variants. This works since each of these variants describe strong monads M that act on **FSpc**.

For any strong monad M , we can define a syntax \vdash_M for programming in that particular language:

$$\frac{(x : X) \in \Gamma}{\Gamma \vdash_M x : X} \text{VAR}$$

$$\frac{f : A_1 \times \cdots \times A_n \rightarrow_M B \quad \Gamma \vdash_M x_1 : A_1 \quad \cdots \quad \Gamma \vdash_M x_n : A_n}{\Gamma \vdash_M f(x_1, \dots, x_n) : B} \text{APP},$$

Theorem 3.23. *Given any expression $\Gamma \vdash_M e : A$, we can construct a term $\llbracket e \rrbracket : \text{Prod}(\Gamma) \rightarrow_M A$ in the Kleisli category for M .*

Proof. We proceed by induction on its structure of the syntax. In the VAR case, we compute from the witness $(x : X) \in \Gamma$ a projection $\llbracket x \rrbracket : \text{Prod}(\Gamma) \rightarrow_c X$, which can also be considered as map in the Kleisli category for M . In the APP case, we are given maps $\llbracket x_i \rrbracket : \Gamma \rightarrow_M A_i$ for $i \in 1, \dots, n$. We use these maps to build a map $x : \Gamma \rightarrow_c M(A_1) \times \cdots \times M(A_n)$ by the universal property for products, i.e., $x \triangleq \langle \llbracket x_1 \rrbracket, \dots, \llbracket x_n \rrbracket \rangle$.

The task then is to define a map $g : M(A_1 \times \cdots \times A_n) \rightarrow_M A_1 \times \cdots \times A_n$ such that we can produce the composition $f \circ g \circ x : \Gamma \rightarrow_M A_1 \times \cdots \times A_n$, which has the right type. This can be accomplished via repeated use of the strength operation $A \times M(B) \rightarrow_M A \times B$ (in conjunction with other monad operations). \square

Chapter 4

Overlapping pattern matching

This chapter introduces a programming construct that ought to be useful for programming with spaces: *overlapping pattern matches*.

Reconsider the approximate root-finding task introduced in Section 1.1. The code

$$\begin{aligned} \text{roots}_f &: * \xrightarrow{nd}_c \{ * \mid \forall x \in K. f(x) \neq 0 \} + \{ x : K \mid |f(x)| < \varepsilon \} \\ \text{roots}_f &\triangleq \text{case}(\text{tt}) \begin{cases} \iota[\exists x \in K. |f(x)| < \varepsilon](y) & \Rightarrow \text{inr}(\text{simulate}(y)) \\ \iota[\forall x \in K. f(x) \neq 0](n) & \Rightarrow \text{inl}(n) \end{cases} \end{aligned}$$

accomplishes this task with an overlapping pattern match. The program is nondeterministic (indicated by the flavor of the arrow \xrightarrow{nd}_c), returning either an $\text{inr}(x)$ that is “almost a root” or a proof indicating that there are no roots in K . The notations and auxiliary definitions in this program will be explained over the course of this chapter and the following one, and we will return to this example to give a more thorough analysis in section 5.4.

Intuitively, overlapping patterns behave much like traditional pattern matches as in functional programming, but an input that matches several patterns may nondeterministically follow any matching branch.

The syntax of an overlapping pattern match determines a unique map that is potentially partial and nondeterministic, but there are simple conditions that ensure that a map is total and deterministic:

1. *Totality*: together, the cases cover the entire input space.
2. *Determinism*: patterns are open embeddings, and “maxima” exist for each pair of branches on their overlap.

Formal (constructive) proofs of these properties contribute to the computational behavior of the pattern match.

In the roots_f example, totality is satisfied essentially because either of the two earlier bulleted statements must hold (Section 5.4 provides further detail).

Contributions

- We demonstrate (Section 4.1), by informal example, the utility of overlapping patterns for a variety of purposes, including
 - (nondeterministic) approximate computation and decision-making (4.1.2)
 - manipulating partial data (4.1.3)
 - “sheafification” of constructions (4.1.4).
- We describe constructions of a few variations on overlapping pattern matches. We first characterize what families of patterns may be used if no properties are assumed about the code for the branches (4.2); these constructions generalize the conventional notion of pattern matching in functional programming. We then describe a construction that generalizes the sheaf gluing property of representable sheaves, where overlapping branches need not agree exactly (4.3). We then discuss some additional properties of overlapping patterns (4.4).

4.1 Example uses of overlapping pattern matching

This section presents an informal introduction to overlapping pattern matching, providing examples that demonstrate its usefulness in a variety of situations.

4.1.1 A familiar example

Here we will fix some notation and define concepts that will be used to explain the more exotic examples later in this section, by starting with a familiar example: we can define functions on coproducts (sums) via pattern matching, just as is possible in any functional programming language:

$$\begin{aligned} \text{forget_sign} &: \{\mathbb{R} \mid \cdot < 0\} + \{\mathbb{R} \mid \cdot > 0\} \rightarrow_c \mathbb{R} \\ \text{forget_sign}(x) &\triangleq \text{case}(x) \begin{cases} \text{inl}(\ell) & \Rightarrow \iota[\cdot < 0](\ell) \\ \text{inr}(r) & \Rightarrow \iota[\cdot > 0](r) \end{cases}. \end{aligned}$$

For general spaces A and B , $\text{inl} : A \hookrightarrow A + B$ and $\text{inr} : B \hookrightarrow A + B$ are open embeddings (theorem 2.15); every pattern must be an open embedding when defining deterministic functions. To confirm that `forget_sign` defines a continuous map, we must check totality and determinism conditions. Since `inl` and `inr` together cover the input space, and since their images are disjoint, the program must be both total and deterministic.

4.1.2 Approximate computation and decision-making

Connected spaces, such as \mathbb{R} , cannot be partitioned into disjoint open subspaces. Accordingly, it is impossible to make a nontrivial discrete decision over these spaces;

for instance, any function $f : \mathbb{R} \rightarrow_c \mathbb{B}$ must be a constant function. Instead, it makes sense to allow decisions to be nondeterministic, which can be accomplished by producing output in the space $\mathcal{P}_\diamond^+(\mathbb{B})$ of nondeterministic Boolean values, rather than \mathbb{B} .

For any error tolerance $\varepsilon > 0$, we can approximately compare a real number with 0 using the function

$$\begin{aligned} \text{positive?}_\varepsilon &: \mathbb{R} \xrightarrow{nd}_c \mathbb{B} \\ \text{positive?}_\varepsilon(x) &\triangleq \text{case}(x) \left\{ \begin{array}{ll} \iota[\cdot > -\varepsilon] (_) \Rightarrow \text{true} \\ \iota[\cdot < \varepsilon] (_) \Rightarrow \text{false} \end{array} \right. . \end{aligned}$$

To confirm this definition is well-defined, we confirm that the two patterns cover \mathbb{R} . If an input lies in $-\varepsilon < \cdot < \varepsilon$, then the resulting behavior is the union of the behaviors of both branches: $\text{true} \sqcup \text{false} : * \xrightarrow{nd}_c \mathbb{B}$; that is, either decision might be made. Note that determining whether an input satisfies $\cdot > -\varepsilon$ is undecidable (likewise for $\cdot < \varepsilon$).

We can nondeterministically break a real number into an integer and a fractional part in $\{\mathbb{R} \mid -\varepsilon < \cdot < 1\}$ with

$$\begin{aligned} \text{int_frac}_\varepsilon &: \mathbb{R} \xrightarrow{nd}_c \mathbb{Z} \times \{\mathbb{R} \mid -\varepsilon < \cdot < 1\} \\ \text{int_frac}_\varepsilon(x) &\triangleq \text{case}(x) \left\{ \text{nplus}(n, y) \Rightarrow (n, y) , \right. \end{aligned}$$

where the open map $\text{nplus} : \mathbb{Z} \times \{\mathbb{R} \mid -\varepsilon < \cdot < 1\} \rightarrow_o \mathbb{R}$ adds its arguments. It is possible for a real number to have several preimages under nplus ; one can think of the pattern match as nondeterministically choosing all possible preimages.

Section 5 describes binary covers, giving many further examples of approximate decision-making using overlapping patterns.

4.1.3 Manipulating partial data

We can write, for instance, a partial Boolean **or** operation,

$$\begin{aligned} \text{or} &: \mathbb{B}_\perp \times \mathbb{B}_\perp \xrightarrow{p}_c \mathbb{B} \\ \text{or}(p) &\triangleq \text{case}(p) \left\{ \begin{array}{ll} \text{up}(\text{true}), _ & \Rightarrow \text{true} \\ _, \text{up}(\text{true}) & \Rightarrow \text{true} \\ \text{up}(\text{false}), \text{up}(\text{false}) & \Rightarrow \text{false} \end{array} \right. . \end{aligned}$$

This example demonstrates that it is possible to pattern match on products as well as to nest patterns. For instance, $\text{up}(\text{true})$ corresponds to the nested pattern

$$* \xrightarrow{\text{true}} \mathbb{B} \xrightarrow{\text{up}} \mathbb{B}_\perp .$$

Since \mathbb{B} is discrete, its points are open embeddings.

Unlike a short-circuiting “or” operation in most programming languages, this **or** is

strict in neither argument, since

$$\text{or}(\text{up}(\text{true}), \perp) = \text{or}(\perp, \text{up}(\text{true})) = \text{true},$$

and **true** is a total value.

We must check that the map is deterministic, by ensuring pairs of branches are compatible on their overlaps. The first and second cases do overlap, but they agree exactly, and other case pairs are disjoint from each other. So **or** is well-defined.

As we can see, when constructing a partial map with a pattern match, the patterns need not cover the entire input space.

Other operations on lifted spaces are naturally expressed with overlapping patterns. For instance, any $f : A \rightarrow_c B$ can be lifted to operate on lifted spaces with the overlapping pattern

$$\begin{aligned} f_\perp : A_\perp &\xrightarrow{p}_c B \\ f_\perp(x) &\triangleq \text{case}(x) \left\{ \text{up}(z) \Rightarrow f(z), \right. \end{aligned}$$

which forces its argument to compute the result, but if the input is \perp , then so is the result. Since there is only one branch that does not return \perp , the gluing condition is satisfied. Lifting of spaces forms a monad, whose join operation can be expressed with the pattern match

$$\begin{aligned} \text{join} : (A_\perp)_\perp &\xrightarrow{p}_c A \\ \text{join}(x) &\triangleq \text{case}(x) \left\{ \text{up}(\text{up}(y)) \Rightarrow y. \right. \end{aligned}$$

Comparison with Haskell pattern matching

Pattern matching definitions of the form $A_{1\perp} \times \dots \times A_{n\perp} \xrightarrow{p}_c B$ may resemble pattern matches in Haskell. The translated definitions sometimes, but not always, represent valid Haskell functions. For instance, the partial Boolean **or** operation corresponds to the following prospective Haskell definition

```
or :: Bool -> Bool -> Bool
or True _ = True
or _ True = True
or False False = False
```

which has overlapping cases, thus making it invalid Haskell code (Haskell gives a warning and ignores the second case). The problem is that both of the first two cases apply if both arguments are **True**, and in Haskell there is no way to guarantee that the results are compatible in either case. Our pattern matching construction allows overlapping patterns, since it requires proof that overlapping patterns indeed are reconcilable on their overlap.

4.1.4 “Sheafification” of constructions

Overlapping patterns naturally allow “sheafification” (see Section 4.3) of constructions, yielding a very general tool that aids in defining continuous maps. For instance, suppose we want to define multiplication on \mathbb{R} , where \mathbb{R} is defined as a metric completion of the set \mathbb{Q} (as in [41]). Using a straightforward construction to extend a Lipschitz function over metric sets to their metric completions, it is possible to define a family of Lipschitz functions that compute multiplication where one argument is bounded,

$$\text{scale}_L : \{\mathbb{R} \mid -L < \cdot < L\} \times \mathbb{R} \rightarrow_c \mathbb{R}$$

for $L : \mathbb{Q}^+$. Multiplication on \mathbb{R} is not Lipschitz and so cannot be defined directly with the Lipschitz construction. However, multiplication is *locally* Lipschitz: the family of bounded multiplication functions defined above cover the entire input domain $\mathbb{R} \times \mathbb{R}$, and so it is possible to define multiplication with the overlapping pattern

$$\begin{aligned} \times : \mathbb{R} \times \mathbb{R} &\rightarrow_c \mathbb{R} \\ x \times y &\triangleq \text{case}(x) \\ &\left\{ [L : \mathbb{Q}^+] \quad \iota[-L < \cdot < L](x') \Rightarrow \text{scale}_L(x', y). \right. \end{aligned}$$

The covering condition is clearly satisfied, and the gluing condition is satisfied because any two branches agree exactly on their region of overlap.

Just as the Lipschitz extension construction can be “sheafified” to allow extension of locally Lipschitz functions, pattern matching naturally enables “sheafification” of constructions more generally.

4.2 Pattern families

This section characterizes those families of patterns which may be used to match on a scrutinee that comes from a space A , if nothing is to be assumed about the branches. The idea is that we compose a function $f : A \xrightarrow{?}_c B$ by factoring through a disjoint sum over collection of spaces representing the possible patterns and branches, $\sum_{i:I} U_i$, i.e., a composition

$$A \xrightarrow{\text{inv}} \sum_{i:I} U_i \xrightarrow{e} B$$

of a “pattern matching” part inv followed by the “branch execution” part e . The collection of branches $(e_i : U_i \xrightarrow{?}_c B)_{i:I}$ exactly correspond to the branch execution function e , but the pattern matching part inv is more interesting; this section will address those families of patterns which may yield valid functions of this sort.

Semantically, we think of a single pattern as representing a space U together with a map $p : U \rightarrow_c A$ that represents the possibility that the scrutinee can be represented as a point in the image of p . For a single pattern $p : U \rightarrow_c A$ to be somehow implementable, it must have a well-behaved inverse $p^{-1} : A \xrightarrow{nd,p}_c U$ that is partial and also may be nondeterministic. That means that p should at the very least

be an open map. If we are building a program that is deterministic, then p^{-1} should be deterministic, which implies p must be an open embedding. If we are building a program that is total, then we do not need each p^{-1} to be total, but we do need the collection of them to cover A .

In general, we have an entire family of patterns $p_i : U_i \rightarrow_o A$ for $i : I$ where I is some index type. We can use this to construct the nondeterministic inverse

$$\begin{aligned} \text{inv} : A &\xrightarrow{nd,p}_c \sum_{i:I} U_i \\ \text{inv}(x) &\triangleq \bigsqcup_{i:I} \text{inj}_i(p_i^{-1}(x)). \end{aligned}$$

Therefore, for any index type I , any collection of opens maps $p_i : U_i \rightarrow_o A$ is a collection of patterns for defining a nondeterministic and partial map using an overlapping pattern match. Given an arbitrary collection of branches $e_i : U_i \xrightarrow{nd,p}_c B$, or equivalently, $e : \sum_{i:I} U_i \xrightarrow{nd,p}_c B$, the overlapping pattern match is just the composition $e \circ \text{inv} : A \xrightarrow{nd,p}_c B$. As expected, the overlapping pattern match is a nondeterministic union of its branches.

For those categories/languages that require totality or determinism, we'd like to characterize the families of patterns which are suitable in those cases. These will be subcollections of the collection of pattern families in the nondeterministic and partial case.

We will observe that they form a lattice of Grothendieck pretopologies. This structure is useful: it tells us that there are certain techniques that we can always use to form pattern families, and that pattern families will have important structural properties. For instance, the transitivity axiom corresponds to the ability to flatten nested pattern matches into a single one.

Definition 4.1 ([19]). *A Grothendieck pretopology (or basis for a Grothendieck topology) is an assignment to each space A of a collection of families $(U_i \rightarrow_c A)_{i:I}$ of continuous maps, called covering families such that*

1. homeomorphisms cover – every family consisting of a single isomorphism $\{V \xrightarrow{\cong}_c A\}$ is a covering family;
2. stability axiom – the collection of covering families is stable under pullback: if $(U_i \rightarrow_c A)_{i:I}$ is a covering family and $f : V \rightarrow_c A$ is any continuous map, then the family of pullbacks $(f^*U_i \rightarrow_c V)_{i:I}$ is a covering family;
3. transitivity axiom – if $(U_i \rightarrow_c A)_{i \in I}$ is a covering family and for each i also $(U_{i,j} \rightarrow_c U_i)_{j : J_i}$ is a covering family, then also the family of composites $(U_{i,j} \rightarrow_c U_i \rightarrow_c A)_{i:I, j:J_i}$ is a covering family.

Proposition 4.2 (Product axiom). *In any Grothendieck pretopology, given covering families $(p_i : U_i \rightarrow_c A)_{i:I}$ and $(q_j : V_j \rightarrow_c B)_{j:J}$, there is a product covering family $(p_i \otimes q_j : U_i \times V_j \rightarrow_c A \times B)_{(i,j):I \times J}$.*

Proof. First, we use pullback stability along $\text{fst} : A \times B \rightarrow_c A$ to produce the covering family $(\text{fst}^* p_i : U_i \times B \rightarrow_c A \times B)_{i:I}$. Then, for each $i : I$, we use pullback stability along $\text{snd} : U_i \times B \rightarrow_c B$ to produce the covering family $(\text{snd}^* q_j : U_i \times V_j \rightarrow_c U_i \times B)_{j:J}$. Finally, by the transitivity axiom, we get the covering family

$$(\text{fst}^* p_i \circ \text{snd}^* q_j : U_i \times V_j \rightarrow_c A \times B)_{i:I, j:J}.$$

Then one can confirm that

$$\text{fst}^* p_i \circ \text{snd}^* q_j = p_i \otimes q_j.$$

□

Proposition 4.3. *The collection of pattern families for $\mathbf{FSpc}_{nd,p}$, that is, open maps $p_i : U_i \rightarrow_o A$, form a Grothendieck pretopology $\mathcal{J}_{nd,p}$.*

Proof. 1. Homeomorphisms cover, since any homeomorphism is (an open embedding and thus) an open map.

2. Stability follows from the fact that the pullback of an open map against any continuous map is open (see Proposition C3.1.11 (i) in [15]).

3. Transitivity follows from the fact that open maps are closed under composition. □

4.2.1 Totality

If we want to ensure that an overlapping pattern match is total, only requiring that the branches themselves $e : \sum_{i:I} U_i \xrightarrow{nd}_c B$ are total, then it suffices to require that $\text{inv}^*(\top) = \top$, or equivalently,

$$\top \leq \bigvee_{i:I} p_{i!}(\top),$$

meaning that the patterns cover the whole input space.

Proposition 4.4. *The collection of pattern families for \mathbf{FSpc}_{nd} , that is, open maps $p_i : U_i \rightarrow_o A$ satisfying $\top \leq \bigvee_{i:I} p_{i!}(\top)$, form a Grothendieck pretopology \mathcal{J}_{nd} .*

Proof. 1. An isomorphism $p : V \rightarrow_o A$ satisfies $p_!(\top) = \top$ and so it alone is a pattern family on A .

2. Given a pattern family $p_i : U_i \rightarrow_o A$ and a continuous map $f : V \rightarrow_c A$, the pullback object $f^* U_i$ is homeomorphic to $\{V \mid f^*(p_{i!}(\top))\}$. Using this definition of the pullback object, we confirm

$$\bigvee_{i:I} (f^* p_i)_!(\top) = \bigvee_{i:I} f^*(p_{i!}(\top)) = f^* \left(\bigvee_{i:I} p_{i!}(\top) \right) = f^*(\top) = \top.$$

3. The family of composites indeed covers, since

$$\begin{aligned}
\bigvee_{i:I, j:J_i} (p_i \circ p_{i,j})!(\top) &= \bigvee_{i:I, j:J_i} p_{i!}(p_{i,j}!(\top)) \\
&= \bigvee_{i:I} p_{i!} \left(\bigvee_{j:J_i} p_{i,j}!(\top) \right) \quad (\text{direct images preserve joins}) \\
&= \bigvee_{i:I} p_{i!}(\top) \\
&= \top.
\end{aligned}$$

□

4.2.2 Determinism

Suppose we want to ensure an overlapping pattern match is deterministic, while allowing any arbitrary family of branches $e : \sum_{i:I} U_i \xrightarrow{p}_c B$ (where each branch e_i is indeed deterministic).

In this case, we will recover the familiar condition required of pattern matching in functional programming: disjointness of patterns (i.e., patterns are not allowed to overlap).

Definition 4.5. *A family of open maps $p_i : U_i \rightarrow_o A$ is pairwise disjoint if whenever $p_{i!}(\top) \wedge p_{j!}(\top)$ is positive in A , then $i \equiv j$.*

Theorem 4.6. *The map inv^* preserves binary meets if and only if both of the following conditions hold:*

1. *Each pattern $p_i : U_i \rightarrow_o A$ is in fact an open embedding.*
2. *The patterns are pairwise disjoint.*

Proof. First, we prove the two conditions hold if inv^* preserves binary meets.

1. It suffices to show that each $p_{i!}$ preserves binary meets, which follows from the calculation

$$p_{i!}(a \wedge b) = \text{inv}^*((i, a) \wedge (i, b)) = \text{inv}^*(i, a) \wedge \text{inv}^*(i, b) = p_{i!}(a) \wedge p_{i!}(b).$$

2. We must show that if $p_{i!}(\top) \wedge p_{j!}(\top)$ is positive in A , then $i \equiv j$.

$$\begin{aligned}
p_{i!}(\top) \wedge p_{j!}(\top) &= \text{inv}^*(i, \top) \wedge \text{inv}^*(j, \top) \\
&= \text{inv}^*((i, \top) \wedge (j, \top)).
\end{aligned}$$

Since inv^* preserves joins, by Proposition 2.14 if $p_{i!}(\top) \wedge p_{j!}(\top)$ is positive, so is $(i, \top) \wedge (j, \top)$, which implies that $i \equiv j$.

Now we prove the converse: if a family $p_i : U_i \hookrightarrow A$ of open embeddings is pairwise disjoint, its inv map is in fact deterministic. It suffices to prove that inv^* preserves binary meets of basic opens. For any open $a : \mathcal{O}(A)$ and proposition Q , let $\chi_Q(a) \triangleq \bigvee_{q:Q} a$. If Q is true, then $\chi_Q(a) = a$ and if Q is false, then $\chi_Q(a) = \perp$.

$$\begin{aligned}
\text{inv}^*((k, a) \wedge (\ell, b)) &= \bigvee_{i:I} p_{i!}(\chi_{i \equiv k}(a) \wedge \chi_{i \equiv \ell}(b)) \\
&= \bigvee_{i:I} p_{i!}(\chi_{i \equiv k}(a)) \wedge p_{i!}(\chi_{i \equiv \ell}(b)) && (p_{i!} \text{ preserves meets}) \\
&= \bigvee_{i:I, j:I} p_{i!}(\chi_{i \equiv k}(a)) \wedge p_{j!}(\chi_{j \equiv \ell}(b)) && (p_i \text{ s pairwise disjoint}) \\
&= \left(\bigvee_{i:I} p_{i!}(\chi_{i \equiv k}(a)) \right) \wedge \left(\bigvee_{j:I} p_{j!}(\chi_{j \equiv \ell}(b)) \right) \\
&= \text{inv}^*(k, a) \wedge \text{inv}^*(\ell, b).
\end{aligned}$$

□

Proposition 4.7. *The collection of pattern families for \mathbf{FSpc}_p , that is, open embeddings $p_i : U_i \hookrightarrow A$ that are pairwise disjoint, form a Grothendieck pretopology \mathcal{J}_p .*

Proof. 1. Any cover with a single pattern is trivially pairwise disjoint.

2. By Proposition 3.20, the pullback of an open embedding is an open embedding. Thus it remains to confirm that pullback preserves disjointness. From the computation

$$(f^*p_i)!(\top) \wedge (f^*p_j)!(\top) = f^*(p_i!(\top)) \wedge f^*(p_j!(\top)) = f^*(p_i!(\top) \wedge p_j!(\top)),$$

and since f^* preserves joins, by Proposition 2.14 if $(f^*p_i)!(\top) \wedge (f^*p_j)!(\top)$ is positive, then so is $p_i!(\top) \wedge p_j!(\top)$, which implies that $i \equiv j$. Therefore, the pullback of the cover is still pairwise disjoint.

3. Since the composition of open embeddings is an open embedding, it only remains to confirm that transitivity preserves disjointness. Suppose we consider two composites $p_i \circ p_{i,j}$ and $p_{i'} \circ p_{i',j'}$ such that

$$(p_i \circ p_{i,j})!(\top) \wedge (p_{i'} \circ p_{i',j'})!(\top) = p_{i!}(p_{i,j!}(\top)) \wedge p_{i'!}(p_{i',j'!}(\top))$$

is positive. This implies that the larger open

$$p_{i!}(\top) \wedge p_{i'!}(\top)$$

is also positive, and therefore, $i \equiv i'$. We will now prove that $j \equiv j'$. Since

$i \equiv i'$, by equality induction we can consider them both i , so that we know that

$$\begin{aligned} (p_i \circ p_{i,j})!(\top) \wedge (p_i \circ p_{i,j'})!(\top) &= p_{i!}(p_{i,j!}(\top)) \wedge p_{i!}(p_{i,j'!}(\top)) \\ &= p_{i!}(p_{i,j!}(\top) \wedge p_{i,j'!}(\top)) \quad (p_{i!} \text{ preserves meets}) \end{aligned}$$

is positive. Note that any direct image map $f_!$ preserves joins, and accordingly, if $f_!(U)$ is positive then U is positive. Therefore,

$$p_{i,j!}(\top) \wedge p_{i,j'!}(\top)$$

is positive, and by pairwise disjointness of the covering family, we know $j \equiv j'$. Therefore, $(i, j) \equiv (i', j')$. □

4.2.3 Totality and determinism

To create a pattern match which is both total and deterministic, we simply combine the separate conditions for totality and determinism.

Proposition 4.8. *The collection of pattern families for \mathbf{FSpc} , that is, open embeddings $p_i : U_i \hookrightarrow A$ that are pairwise disjoint and cover A , form a Grothendieck pretopology $\mathcal{J} = \mathcal{J}_{nd} \cap \mathcal{J}_p$ on \mathbf{FSpc} .*

Proof. It is straightforward from the definition of a Grothendieck pretopology that they are closed under arbitrary intersection. □

Definition 4.9. *An open U is clopen (for “closed” and “open”) if it has a Boolean complement, i.e., there is another open V such that $U \vee V = \top$ and $U \wedge V = \perp$.*

Proposition 4.10. *For a pattern family $(p_i : U_i \hookrightarrow A)_{i:I}$ on \mathbf{FSpc} , if the index type I has decidable equality, then for each $i : I$, $p_{i!}(\top)$ is clopen in A .*

Proof. Fix some $i : I$. We claim that the Boolean complement of $p_{i!}(\top)$ is $\bigvee_{j:I|j \neq i} p_{j!}(\top)$. Since I has decidable equality, their join is \top :

$$p_{i!}(\top) \vee \bigvee_{j:I|j \neq i} p_{j!}(\top) = \bigvee_{i:I} p_{i!}(\top) = \top.$$

Pairwise disjointness implies that their meet is \perp :

$$p_{i!}(\top) \wedge \bigvee_{j:I|j \neq i} p_{j!}(\top) = \bigvee_{j:I|j \neq i} p_{i!}(\top) \wedge p_{j!}(\top) = \bigvee_{j:I|j \neq i} \perp = \perp.$$

□

This means that if the index type I has decidable equality, the predicate corresponding to any given pattern is decidable. This recovers the usual understanding of pattern matching in functional programming, where patterns are disjoint and correspond to decidable predicates.

4.2.4 Syntax of patterns

We will define the admissible syntax of patterns, where $p : A \dashv \Gamma$ intuitively means that the syntax p provides a “context” of pattern matching variables Γ by pattern matching on a space A . For instance, we could have the pattern

$$(\text{up}(x), y) : A_{\perp} \times B \dashv x : A, y : B$$

on a space $A_{\perp} \times B$ that provides variables $x : A$ and $y : B$ to be used in the branch corresponding to that pattern.

$$\begin{array}{c} \frac{f : * \rightarrow_{?} A}{f : A \dashv_{?} \cdot} \text{CONSTANT} \qquad \frac{}{v : A \dashv_{?} v : A} \text{VAR} \qquad \frac{}{_ : A \dashv_{?} \cdot} \text{WILDCARD} \\[10pt] \frac{p : U \dashv_{?} \Gamma \quad f : U \rightarrow_{?} A}{f(p) : A \dashv_{?} \Gamma} \text{COMPOSE} \qquad \frac{p : A \dashv_{?} \Gamma \quad q : B \dashv_{?} \Delta}{p, q : A \times B \dashv_{?} \Gamma, \Delta} \text{PRODUCT} \end{array}$$

Figure 4-1: The syntax for patterns. The symbol $?$ can either be o , for open maps, or oe for open embeddings.

Figure 4-1 characterizes the syntax of patterns. We can either create patterns of open maps (when building a nondeterministic function) or open embeddings (when building a necessarily deterministic continuous map).

Theorem 4.11. *For $?$ instantiated as either o (for open maps) or oe (for open embeddings), given any pattern derivation $p : A \dashv_{?} \Gamma$, there is a map $\llbracket p \rrbracket : \text{Prod}(\Gamma) \times \Delta \rightarrow_{?} A$ for some space Δ which collects the “discarded variables” from the wildcards.*

Proof. By induction on the derivation of the pattern:

CONSTANT By assumption, we have a constant map of the right kind.

VAR We use $\text{id} : A \rightarrow_{oe} A$ (which is also an open map).

WILDCARD We *also* use $\text{id} : A \rightarrow_{oe} A$, but since we require $\Gamma \cong *$, we set the “garbage space” Δ to be A .

COMPOSE By induction, we have a map $\llbracket p \rrbracket : \Gamma \times \Delta \rightarrow_{?} U$ for some Δ . Then we use the composition $(f \otimes \text{id}_{\Delta}) \circ \llbracket p \rrbracket : \Gamma \times \Delta \rightarrow_{?} A$, threading through the “garbage” space Δ .

PRODUCT By induction, we have maps $\llbracket p \rrbracket : \Gamma_1 \times \Delta_1 \rightarrow_{?} A$ and $\llbracket q \rrbracket : \Gamma_2 \times \Delta_2 \rightarrow_{?} B$. Regardless of whether $?$ = o or $?$ = oe , their parallel composition $\llbracket p \rrbracket \otimes \llbracket q \rrbracket : \Gamma_1 \times \Gamma_2 \rightarrow_{?} A \times B$ is also a map of the right kind, which we can use, together with some homeomorphisms to rearrange the “garbage” spaces Δ_1 and Δ_2 and to pass them through.

□

Proposition 4.12. *If $p : A \dashv_{oe} \Gamma$, then $p : A \dashv_o \Gamma$ as well.*

Proof. Straightforward, following from the fact that every open embedding is an open map. □

Note that the definition of patterns we give here does *not* admit some structural rules that may be expected, such as

Strengthening Given $p : A \dashv \Gamma$, we cannot necessarily derive $p : A \dashv \Delta$ where Δ is a sub-list of Γ .

Exchange Given $p : A \dashv \Gamma, \Delta$, we cannot necessarily derive $p : A \dashv \Delta, \Gamma$.

These structural properties aren't necessary since patterns are *providing* a context to be *used* by the expression language, which *does* have the corresponding rules weakening and exchange. Regardless, the WILDCARD rule still gives a sort of “explicit” strengthening, allowing the user to discard some variables.

We can now define a general (mostly) syntactic rule for “compiling” the sorts of pattern matches described in this section.

Theorem 4.13. *We can compile the syntax¹*

$$\frac{\Gamma \vdash_{\mathcal{J}} s : A \quad \prod_{i:I} p_i : A_i \dashv_{\mathcal{J}} A \quad \prod_{i:I} \Gamma, A_i \vdash_{\mathcal{J}} e_i : B \quad (p_i)_{i:I} \in \mathcal{J}_{\mathcal{J}}}{\Gamma \vdash_{\mathcal{J}} \left(\text{case}(s) \left\{ [i : I] \quad p_i \Rightarrow e_i \right\} : B \right)} \text{CASE-}\mathcal{J}$$

, where \mathcal{J} and \mathcal{J} are instantiated with any of the combinations in this table:

\mathcal{J}	\mathcal{J}
	oe
p	oe
nd	o
nd, p	o

Proof. The syntactic constructions give us maps $\llbracket s \rrbracket : \Gamma \xrightarrow{\mathcal{J}}_c A$, $\llbracket p_i \rrbracket : A_i \times \Delta_i \rightarrow_i A$ (for some spaces Δ_i representing discarded variables in the pattern p_i), and $e_i : \Gamma \times A_i \xrightarrow{\mathcal{J}}_c B$. The condition $(p_i)_{i:I} \in \mathcal{J}_{\mathcal{J}}$ is interpreted to mean that these “compiled” maps appropriately cover A , so that we get an appropriately behaved map $\text{inv} : A \xrightarrow{\mathcal{J}}_c \sum_{i:I} A_i \times \Delta_i$.

We must produce a map $f : \Gamma \xrightarrow{\mathcal{J}}_c B$. We can do so by defining²

$$f : \Gamma \xrightarrow{\mathcal{J}}_c B$$

$$f(\gamma) \triangleq \text{let } \langle i, x \rangle \triangleq \text{inv}(s(\gamma)) \text{ in } e_i(\gamma, \text{fst}(x)).$$

¹ As mnemonics, I stands for *index* type, s stands for *scrutinee* of a case expression, p stands for *pattern*, and e for *expression*.

² While the “let-in” syntax for using the universal property (itself a sort of pattern matching) of sums has not been formally described, hopefully it is clear how it can be implemented via categorical semantics.

□

Note that the condition $(p_i)_{i:I} \in \mathcal{J}_?$ that the patterns lie in the appropriate Grothendieck topology is trivial when $? = (nd, p)$.

4.3 Pattern matching with “gluing” conditions

One limitation of the pattern matching constructions described in the previous section is that, when determinism is required, patterns are required to be disjoint (and if totality is also required, then under mild conditions each pattern must be clopen as well).

But we should be able to produce deterministic functions even when patterns overlap if we require some notion of compatibility of overlapping branches. The previous example of the partial Boolean `or` operation is such a case where the patterns overlap, but the function is still deterministic, because the outputs are compatible whenever there is overlap.

There is an analogous concept of a sort of “overlapping pattern matching” in sheaf theory, which effectively says that if branches are *identical* when they overlap, one can glue together the branches to form a single function. This is nice, but in fact we do not need to require that maps be identical when they overlap; we can get away with something more permissive. This section will characterize this more permissive notion of overlapping pattern matching that demands some sort of “compatibility” of overlapping branches.

First, let’s briefly examine the similar concept from sheaf theory. Though we described four Grothendieck pretopologies in the previous section, we have yet to define perhaps the most popular Grothendieck pretopology treated in sheaf theory, which is the *open cover topology*.

Definition 4.14. *The open cover topology states that a collection of maps $(p_i : A_i \hookrightarrow A)$ covers the space A if each p_i is an open embedding and together their images cover A , i.e.,*

$$\top \leq \bigvee_{i:I} p_{i!}(\top).$$

Proposition 4.15. *The open cover topology is a Grothendieck pretopology.*

Proof. A minor variation of the proof of Proposition 4.8, which proved that if one also demands pairwise disjointness of covers, then one has a Grothendieck pretopology. □

The open cover topology is subcanonical, which implies that it is possible to define a continuous map $A \rightarrow_c B$ by splitting up A into a potentially overlapping collection of open spaces $(U_i)_{i:I}$ with an a cover $(p_i : U_i \hookrightarrow A)_{i:I}$ in the open cover topology and defining continuous maps $f_i : U_i \rightarrow_c B$, so long as the f_i s agree on their overlap (which we will define more precisely soon).

Accordingly, we will restrict our attention to deterministic maps in this section, as the constructions in the previous section were sufficient in the nondeterministic case.

First, we'll consider the construction of a partial map.

To make the analysis easier, before handling the general syntactic case, we'll initially restrict to a simpler scenario, where we directly scrutinize the input, and then do not use the original input in any of the branches:

$$f : A \xrightarrow{p}_c B$$

$$f(x) \triangleq \text{case}(x) \left\{ [i : I] \quad p_i(x_i) \quad \Rightarrow e_i(x_i) \right\},$$

where i ranges over the index type $I : \mathcal{U}$, each $p_i : A_i \hookrightarrow A$ is an open embedding, each $e_i : A_i \xrightarrow{p}_c B$ is an expression representing a deterministic but partial map, and each $x_i : A_i$ is a pattern variable binding. Note that each p_i must be an open embedding rather than just an open map, or else a single pattern on its own could introduce nondeterminism.

We can certainly define an implementation of the **case** expression that is potentially nondeterministic

$$f : A \xrightarrow{nd,p}_c B$$

$$f \triangleq \bigsqcup_{i:I} e_i \circ p_i^{-1}.$$

We will now attempt to devise a condition on the *branches* (rather than any condition on the patterns) to ensure that f is actually deterministic, i.e., that f^* preserves binary meets. First, we define a notion of compatibility of continuous maps that is weaker than equality but strong enough to ensure that the map f is deterministic when overlapping branches are compatible.

Definition 4.16. *Two continuous maps $f, g : A \rightarrow_c B$ have a maximum if the map*

$$(f \sqcup g)^* : \mathcal{O}(B) \rightarrow \mathcal{O}(A)$$

$$(f \sqcup g)^*(U) \triangleq f^*(U) \vee g^*(U)$$

preserves binary meets, which implies that it defines a continuous map $f \sqcup g : A \rightarrow_c B$.

Trivially, any map has a maximum with itself, $f \sqcup f = f$, so if two maps are equal, then they have a maximum.

Remark 4.17. *If $f, g : A \rightarrow_c B$ have a maximum, then for all $U, V : \mathcal{O}(B)$,*

$$f^*(U) \wedge g^*(V) \leq f^*(V) \vee g^*(U).$$

Assuming classical logic, this implies that for every (global) point x of A , either $f(x) \leq g(x)$ or $g(x) \leq f(x)$ (in terms of specialization order).

Proof. Since

$$(f \sqcup g)^*(U \wedge V) = (f \sqcup g)^*(U) \wedge (f \sqcup g)^*(V),$$

we have

$$\begin{aligned}
f^*(U \wedge V) \vee g^*(U \wedge V) &= (f^*(U) \vee g^*(U)) \wedge (f^*(V) \vee g^*(V)) \\
&= (f^*(U) \wedge f^*(V)) \vee (f^*(U) \wedge g^*(V)) \\
&\quad \vee (g^*(U) \wedge f^*(V)) \vee (g^*(U) \wedge g^*(V)) \\
&= f^*(U \wedge V) \vee g^*(U \wedge V) \\
&\quad \vee (f^*(U) \wedge g^*(V)) \vee (g^*(U) \wedge f^*(V))
\end{aligned}$$

and thus

$$(f^*(U) \wedge g^*(V)) \vee (g^*(U) \wedge f^*(V)) \leq f^*(U \wedge V) \vee g^*(U \wedge V),$$

which implies that

$$\begin{aligned}
f^*(U) \wedge g^*(V) &\leq f^*(U \wedge V) \vee g^*(U \wedge V) \\
&\leq f^*(V) \vee g^*(U).
\end{aligned}$$

We now prove that assuming classical logic, for any point x of A , either $f(x) \leq g(x)$ or $g(x) \leq f(x)$. Suppose that $f(x) \not\leq g(x)$, which means there is some $U : \mathcal{O}(B)$ such that $f(x) \Vdash U$ (i.e., $x^*(f^*(U)) = \top$) but not $g(x) \Vdash U$ (i.e., $x^*(g^*(U)) = \perp$). Using the law we derived, we know (for all V)

$$x^*(f^*(U)) \wedge x^*(g^*(V)) \leq x^*(f^*(V)) \vee x^*(g^*(U)),$$

and thus

$$x^*(g^*(V)) \leq x^*(f^*(V)),$$

which means that $g(x) \leq f(x)$. □

Example. The maps $t_1, t_2 : \mathbb{B} \rightarrow_c \Sigma$ defined by

$$\begin{array}{ll}
t_1(\text{true}) \triangleq \top_\Sigma & t_2(\text{true}) \triangleq \perp_\Sigma \\
t_1(\text{false}) \triangleq \perp_\Sigma & t_2(\text{false}) \triangleq \top_\Sigma
\end{array}$$

have the maximum $f : \mathbb{B} \rightarrow_c \Sigma$ defined by $f(x) \triangleq \top_\Sigma$.

Theorem 4.18 (Determinism). *If branches are “compatible” when they overlap: for each $i, j : I$, there is a map $e_{ij} : A_i \times_A A_j \rightarrow_c B$ that is the maximum of $e_i \circ \theta_{ij}$ and $e_j \circ \varphi_{ij}$ in the pullback diagram*

$$\begin{array}{ccccc}
B & & & & \\
\swarrow e_i & & \xleftarrow{\theta_{ij}} & & A_i \\
& \nearrow e_{ij} & A_i \times_A A_j & \xrightarrow{\theta_{ij}} & A_i \\
& & \downarrow \varphi_{ij} \lrcorner & & \downarrow p_i \\
& & A_j & \xrightarrow{p_j} & A
\end{array}$$

$\nearrow e_j$

then the overlapping pattern match f is in fact deterministic, i.e., $f : A \xrightarrow{p}_c B$.

Proof. We must show f^* preserves binary meets. Since f^* preserves joins, it is monotone, and thus

$$f^*(U \wedge V) \leq f^*(U) \wedge f^*(V),$$

so it remains only to prove the reverse inequality.

Define $p_{ij} : A_i \times_A A_j \hookrightarrow A$ by $p_{ij} = p_j \circ \varphi_{ij} = p_i \circ \theta_{ij}$. By Proposition 3.13, we have

$$p_{ij!}(\top) = p_{i!}(\top) \wedge p_{j!}(\top).$$

We first note the equivalence

$$f^*(U) = \bigvee_{i:I} p_{i!}(e_i^*(U)) = \bigvee_{i,j:I} p_{ij!}(e_{ij}^*(U)).$$

That the middle expression is at most the right follows from the fact that $e_{ii} = e_i$, $p_{ii} = p_i$. The reverse inequality follows from the following inequality:

$$\begin{aligned} p_{ij!}(e_{ij}^*(U)) &= p_{ij!}((e_i \circ \theta_{ij})^*(U) \vee (e_j \circ \varphi_{ij})^*(U)) && (e_{ij} \text{ “compatible”}) \\ &= p_{ij!}((e_i \circ \theta_{ij})^*(U)) \vee p_{ij!}(e_j \circ \varphi_{ij})^*(U)) && (p_{ij!} \text{ preserves joins}) \\ &= p_{i!}(\theta_{ij!}(\theta_{ij}^*(e_i^*(U)))) \vee p_{j!}(\varphi_{ij!}(\varphi_{ij}^*(e_j^*(U)))) && (\text{pullback equations}) \\ &= p_{i!}(\theta_{ij!}(\top) \wedge e_i^*(U)) \vee p_{j!}(\varphi_{ij!}(\top) \wedge e_j^*(U)) && (\text{Frobenius reciprocity}) \\ &\leq p_{i!}(e_i^*(U)) \vee p_{j!}(e_j^*(U)). \end{aligned}$$

This inequality has a sort of partial converse:

$$\begin{aligned} p_{ij!}(\top) \wedge p_{i!}(e_i^*(U)) &= p_{i!}(\theta_{ij!}(\top)) \wedge p_{i!}(e_i^*(U)) \\ &= p_{i!}(\theta_{ij!}(\top) \wedge e_i^*(U)) && (p_{i!} \text{ preserves meets}) \\ &\leq p_{i!}(\theta_{ij!}(\top) \wedge e_i^*(U)) \vee p_{j!}(\varphi_{ij!}(\top) \wedge e_j^*(U)) \\ &= p_{ij!}(e_{ij}^*(U)). \end{aligned}$$

Note that the direct image map $p_{i!}$ (as well as the other direct image maps) preserves meets since p_i is an open embedding.

Putting these arguments together, we can confirm that f^* preserves meets:

$$\begin{aligned}
f^*(U) \wedge f^*(V) &= \left(\bigvee_{i:I} p_{i!}(e_i^*(U)) \right) \wedge \left(\bigvee_{j:I} p_{j!}(e_j^*(V)) \right) \\
&= \bigvee_{i,j:I} p_{i!}(e_i^*(U)) \wedge p_{j!}(e_j^*(V)) \\
&= \bigvee_{i,j:I} p_{i!}(\top) \wedge p_{i!}(e_i^*(U)) \wedge p_{j!}(\top) \wedge p_{j!}(e_j^*(V)) \\
&= \bigvee_{i,j:I} p_{ij!}(\top) \wedge p_{i!}(e_i^*(U)) \wedge p_{ij!}(\top) \wedge p_{j!}(e_j^*(V)) \\
&\leq \bigvee_{i,j:I} p_{ij!}(e_{ij}^*(U)) \wedge p_{ij!}(e_{ij}^*(V)) && \text{(previous argument)} \\
&= \bigvee_{i,j:I} p_{ij!}(e_{ij}^*(U \wedge V)) && (p_{ij!} \text{ preserves meets}) \\
&= f^*(U \wedge V)
\end{aligned}$$

□

It can be informative to consider the case where every open embedding is just an inclusion of an open sublocale (which is, up to isomorphism, just as general). Then the pullback just represents the intersection, noting the homeomorphism

$$\{A \mid p_{i!}(\top) \wedge p_{j!}(\top)\} \cong A_i \times_A A_j.$$

If we additionally want totality, it is easy to specify the conditions which guarantee that, which are effectively the same as in the previous section:

Proposition 4.19 (Totality). *If each $e_i : A_i \rightarrow_c B$ is total, then the overlapping pattern match f is in fact total, i.e., $f : A \xrightarrow{nd}_c B$, if and only if the patterns together cover A , that is,*

$$\top \leq \bigvee_{i:I} p_{i!}(\top),$$

Proof. Given that each e_i is total, we calculate that

$$f^*(\top) = \bigvee_{i:I} p_{i!}(e_i^*(\top)) = \bigvee_{i:I} p_{i!}(\top)$$

And thus $f^*(\top) = \top$ if and only if the patterns together cover A . □

Note that we defined a nondeterministic map in the above definition. If we also add the previous condition on branches for determinism, we will get a total and deterministic (i.e., continuous) map.

More general syntax

We now return to the more general syntax, where the definition is of the form

$$f : \Gamma \rightarrow_c B$$

$$f(\gamma) \triangleq \text{case}(s(\gamma)) \left\{ [i : I] \quad p_i(x_i) \quad \Rightarrow e_i(\gamma, x_i) \quad , \right.$$

we had

$$f = \bigsqcup_{i:I} e_i \circ \langle \text{id}, p_i^{-1} \circ s \rangle.$$

We can put in a more convenient form,

$$\begin{aligned} f &= \bigsqcup_{i:I} e_i \circ \langle \text{id}, p_i^{-1} \circ s \rangle \\ &= \bigsqcup_{i:I} e_i \circ (\text{id} \otimes p_i^{-1}) \circ \langle \text{id}, s \rangle \\ &= \left(\bigsqcup_{i:I} e_i \circ (\text{id} \otimes p_i^{-1}) \right) \circ \langle \text{id}, s \rangle \\ &= \left(\bigsqcup_{i:I} e_i \circ (\text{id} \otimes p_i)^{-1} \right) \circ \langle \text{id}, s \rangle, \end{aligned}$$

where it becomes apparent that the left parenthesized term takes on the form of a “simpler” pattern match as we studied earlier,

$$\tilde{f} : \Gamma \times A \rightarrow_c B$$

$$\tilde{f}(z) \triangleq \text{case}(z) \left\{ [i : I] \quad \gamma, p_i(x_i) \quad \Rightarrow e_i(\gamma, x_i) \quad , \right.$$

such that

$$f = \tilde{f} \circ \langle \text{id}, s \rangle.$$

We can use this decomposition to give sufficient conditions for determinism in the general case.

Theorem 4.20 (Determinism). *If branches are “compatible” when they overlap: for each $i, j : I$, there is a map $e_{ij} : \Gamma \times (A_i \times_A A_j) \rightarrow_c B$ that is the maximum of $e_i \circ (\text{id} \otimes \theta_{ij})$ and $e_j \circ (\text{id} \otimes \varphi_{ij})$ in the pullback diagram*

$$\begin{array}{ccccc} B & & & & \\ & \nwarrow e_i & & & \\ & & \Gamma \times (A_i \times_A A_j) & \xrightarrow{\text{id} \otimes \theta_{ij}} & \Gamma \times A_i \\ & \nearrow e_{ij} & \downarrow \text{id} \otimes \varphi_{ij} \lrcorner & & \downarrow \text{id} \otimes p_i \\ & & \Gamma \times A_j & \xrightarrow{\text{id} \otimes p_j} & \Gamma \times A \\ & \nwarrow e_j & & & \end{array}$$

then the overlapping pattern match f is in fact deterministic, i.e., $f : \Gamma \xrightarrow{p}_c B$:

$$\frac{\prod_{i:I} p_i : A_i \dashv_{oe} A \quad \prod_{i:I} \Gamma, A_i \vdash_p e_i : B \quad \text{pairwise branch compat.}}{\Gamma \vdash_p \left(\text{case}(s) \left\{ [i : I] \quad p_i \Rightarrow e_i \right\} : B \right)} \text{ CASE-P}$$

Proof. Since $f = \tilde{f} \circ \langle \text{id}, s \rangle$ and since s is deterministic, it suffices to confirm that \tilde{f} is deterministic.

We must only confirm that our diagram above is actually equivalent to the one that we would get for \tilde{f} from Proposition 4.18,

It suffices to confirm that $\Gamma \times (A_i \times_A A_j)$ is indeed the pullback of $\text{id} \otimes p_i$ and $\text{id} \otimes p_j$. This follows from the fact that the pullback of “parallel” maps such as this is the product of the pullbacks, and that the pullback of id_Γ with itself is Γ . \square

To define a continuous map that is total and deterministic, we simply combine the requirements:

$$\frac{\prod_{i:I} \Gamma, A_i \vdash e_i : B \quad \top \leq \bigvee_{i:I} p_i(\top) \quad \text{pairwise branch compat.}}{\Gamma \vdash \left(\text{case}(s) \left\{ [i : I] \quad p_i \Rightarrow e_i \right\} : B \right)} \text{ CASE}$$

4.4 Properties of overlapping pattern matching

4.4.1 Partial patterns and lifted spaces

It’s possible to see the conditions for totality of overlapping patterns in another way. There is a natural correspondence between a partial pattern match

$$f : A \xrightarrow{p}_c B$$

$$f(x) \triangleq \text{case}(x) \left\{ [i : I] \quad p_i(x_i) \Rightarrow e_i(x_i) \right.$$

where each $e_i : A_i \xrightarrow{p}_c B$ and a total pattern match that instead outputs to a lifted space,

$$f : A \rightarrow_c B_\perp$$

$$f(x) \triangleq \text{case}(x) \begin{cases} [i : I] & p_i(x_i) & \Rightarrow e_i(x_i) \\ - & & \Rightarrow \perp_B \end{cases}$$

where here we consider each $e_i : A_i \rightarrow_c B_\perp$ as a total map to the lifted space.

We call the additional case a catch-all case. The catch-all case already suffices to cover the whole space, of course, trivially showing that the totality condition holds for such a transformed pattern match. Since \perp_B has maxima with all points of B , the additional determinism conditions induced by adding the catch-all case are satisfied automatically.

It should be apparent that these two pattern match definitions are equivalent.

4.4.2 Related to Grothendieck pretopologies

As the various pattern families we have considered thus far all form Grothendieck pretopologies, all pattern matches have certain properties. For instance, the facts that homeomorphisms alone are covers and covers can be composed each correspond to constructions involving overlapping patterns.

Any homeomorphism is a cover by itself (“reflexivity”), corresponding to a pattern that may change the “view” of data to an isomorphic space (without breaking it apart). For instance, given the homeomorphism $\text{exp} : \mathbb{R} \hookrightarrow \{\mathbb{R} \mid \cdot > 0\}$, we can define multiplication on $\{\mathbb{R} \mid \cdot > 0\}$ using addition on \mathbb{R} with the pattern match

$$\text{mult} : \{\mathbb{R} \mid \cdot > 0\} \times \{\mathbb{R} \mid \cdot > 0\} \rightarrow_c \{\mathbb{R} \mid \cdot > 0\}$$

$$\text{mult}(p) \triangleq \text{case}(p) \begin{cases} \text{exp}(x), \text{exp}(y) & \Rightarrow \text{exp}(x + y) \end{cases}.$$

There is also a notion of composition of covers (“transitivity”), which corresponds to the potential to flatten nested patterns. For instance, the definition

$$f : A_\perp + B_\perp \xrightarrow{p}_c A + B$$

$$f(x) \triangleq \text{case}(x) \begin{cases} \text{inl}(\ell) & \Rightarrow \text{case}(\ell) \begin{cases} \text{up}(\ell') & \Rightarrow \text{inl}(\ell') \end{cases} \\ \text{inr}(r) & \Rightarrow \text{case}(r) \begin{cases} \text{up}(r') & \Rightarrow \text{inr}(r') \end{cases} \end{cases}$$

can be flattened into

$$f(x) \triangleq \text{case}(x) \begin{cases} \text{inl}(\text{up}(\ell)) & \Rightarrow \text{inl}(\ell) \\ \text{inr}(\text{up}(r)) & \Rightarrow \text{inr}(r) \end{cases}.$$

Pulling back covers

Since covering families of open embeddings form a Grothendieck pretopology, given a cover $(p_i : B_i \hookrightarrow B)_{i:I}$ of B and a map $f : A \rightarrow_c B$, the family

$$f^*p_i : \{A \mid f^*(p_i!(\top))\} \hookrightarrow A \quad (i : I)$$

of open embeddings pulled back along f covers B , where we have the pullback diagram

$$\begin{array}{ccc} \{A \mid f^*(p_i!(\top))\} & \xrightarrow{f^*p_i} & A \\ \downarrow & \lrcorner & \downarrow f \\ B_i & \xrightarrow{p_i} & B \end{array}$$

By using pulled-back covers, it is possible to pattern match on an input $x : A$ by doing a case analysis on $f(x) : B$, such that in each branch it is known that x lies in a particular open subspace of A (rather than only knowing that $f(x)$ lies in a particular open subspace of B). The introductory example, as explained in section 5.4, uses a pulled-back cover in this way.

4.4.3 Evaluation of patterns

The purpose of this section is to explain how overlapping pattern matches behave under evaluation (i.e., substitution or composition). Intuitively, when we evaluate an overlapping pattern match on a value x , we collect those patterns that x falls into, and the result is the nondeterministic union of the branches corresponding to that pattern. We will now make this notion formal.

A map $f : A \xrightarrow{?}_c B$ (for $? \in \{c, nd, p, (nd, p)\}$) defined by the pattern match

$$f(x) \triangleq \mathbf{case}(x) \left\{ [i : I] \quad p_i(x_i) \quad \Rightarrow \quad e_i(x_i) \right.$$

where each $p_i : A_i \rightarrow_o A$ and $e_i : A_i \xrightarrow{?}_c B$, when evaluated on a (generalized) point $x : \Gamma \xrightarrow{?}_c A$, can be described by

$$(f \circ x)(\gamma) = \mathbf{case}(\gamma) \left\{ [i : I] \quad f^*p_i(\gamma_i) \quad \Rightarrow \quad e_i(x_i(\gamma_i)), \right.$$

where each f^*p_i and x_i come from the pullback diagram

$$\begin{array}{ccc} \{\Gamma \mid f^*(p_i!(\top))\} & \xrightarrow{f^*p_i} & \Gamma \\ \downarrow x_i & \lrcorner & \downarrow x \\ A_i & \xrightarrow{p_i} & A \end{array}$$

In particular, suppose for each $i : I$, the open

$$f^*(p_{i!}(\top)) : \mathcal{O}(\Gamma)$$

is either \top or \perp . If x is a global point, meaning $\Gamma \cong *$, then classically this is the case. Let I_\top and I_\perp represent the indices where the opens are \top and \perp , respectively, and for $i : I_\top$, let $x_{A_i} : \Gamma \rightarrow_c \{A \mid p_{i!}(\top)\}$ be x with a restricted range. Then for any $U : \mathcal{O}(B)$,

$$\begin{aligned} (f \circ x)^*(U) &= \bigvee_{i:I} (e_i \circ x_i)^*(U) \\ &= \bigvee_{i:I_\top} (e_i \circ x_i)^*(U) \\ &= \bigvee_{i:I_\top} (e_i \circ p_i^{-1} \circ x_{A_i})^*(U) \end{aligned}$$

These arguments confirm the intuitive explanation about evaluation of overlapping patterns: collect the possible branches, apply them to the point, and the behavior will be the nondeterministic union of those results.

For instance, in section 4.1.3, we found that the partial Boolean `or` satisfies

$$\text{or}(\text{up}(\text{true}), \perp) = \text{true}.$$

We confirm this by “pulling back” the pattern match defining `or` on this input to

$$\text{or}(\text{up}(\text{true}), \perp) = \text{case}(\text{tt}) \left\{ \begin{array}{ll} _ & \Rightarrow \text{true} \\ \text{abort}(x) & \Rightarrow \text{abort}(x) \\ \text{abort}(x) & \Rightarrow \text{abort}(x) \\ _ & \Rightarrow \perp \end{array} \right.,$$

where for any space A , `abort` : $\emptyset \hookrightarrow A$ is the trivial embedding of the empty space \emptyset . Patterns that match the scrutinee reduce to the embedding `id` : $* \hookrightarrow *$ (the wildcard pattern), and patterns that it misses entirely reduce to `abort`. We can then collect only the possible branches to conclude that

$$\text{or}(\text{up}(\text{true}), \perp) = \perp \sqcup \text{true} = \text{true}.$$

Chapter 5

Binary covers: nondeterministic decision procedures

In this section, we develop a theory of *binary covers*, which are pairs of opens that cover a given space. These binary covers are useful for computing approximate decisions on spaces, where computing exact decisions might not be possible. For instance, \mathbb{R} has no nontrivial decidable predicates since it is connected, but it admits a wealth of binary covers.

The basic idea is to generalize the notions of exhaustive reasoning from general functional programming, where it is possible to compute whether or not decidable propositions hold. There is a calculus of decidable predicates: they are closed under conjunction, disjunction, and negation, and universal and existential quantification over finite¹ sets yields decidable predicates as well. We will generalize this calculus to binary covers, with finite sets generalized to *compact/overt* spaces.

We can relax the notion of decidable predicates to binary covers. Whereas a decidable proposition P is one where either P or $\neg P$ holds, a binary cover is a pair of propositions P and Q such that either P or Q holds. In a binary cover, it is possible that P and Q *both* hold, which is of course never the case for P and $\neg P$ in decidable propositions. So every decidable proposition can be interpreted as a binary cover, but not the other way around.

Translating to spatial language, a decidable predicate on a space A is an open P of A such that there is a disjoint open $\neg P$ satisfying $\top \leq P \vee \neg P$ (equivalently, P is *clopen*). A binary cover is a pair of opens P and Q that cover the space, i.e., $\top \leq P \vee Q$. For example, the `positive ε` program in Section 4.1.2 covers \mathbb{R} with the pair of opens $\cdot > -\varepsilon$ and $\cdot < \varepsilon$. More generally, any overlapping pattern match with two cases specifies a binary cover of the input space.

5.1 Binary covers as nondeterministic decisions

Definition 5.1. A binary cover of a space A is an ordered pair (P, Q) of opens of A satisfying $\top \leq P \vee Q$.

¹ By “finite,” we always mean Kuratowski-finite [14].

Proposition 5.2. *Binary covers of A are in bijective correspondence with nondeterministic Boolean-valued maps $A \xrightarrow{nd}_c \mathbb{B}$.*

Proof. Suppose we have a binary cover (P, Q) of A . Then define $f^* : \mathcal{O}(\mathbb{B}) \rightarrow \mathcal{O}(A)$ by $f^*(\cdot = \text{true}) \triangleq P$ and $f^*(\cdot = \text{false}) \triangleq Q$. To confirm f^* defines a nondeterministic map $f : A \xrightarrow{nd}_c \mathbb{B}$, we must confirm f^* preserves joins and \top . Since \mathbb{B} , as a discrete space, has no covering axioms, f^* trivially preserves joins. And f^* also preserves \top , since

$$f^*(\top) = f^*((\cdot = \text{true}) \vee (\cdot = \text{false})) = P \vee Q = \top.$$

Conversely, given a nondeterministic map $f : A \xrightarrow{nd}_c B$, we get a binary cover of A given by

$$(f^*(\cdot = \text{true}), f^*(\cdot = \text{false})),$$

which we confirm is covering since

$$f^*(\cdot = \text{true}) \vee f^*(\cdot = \text{false}) = f^*((\cdot = \text{true}) \vee (\cdot = \text{false})) = f^*(\top) = \top.$$

□

This establishes two alternative perspectives on approximate decision-making, one spatial and one more algorithmic in flavor. The algorithmic one makes clear that we should be able to use Boolean operations to combine binary covers into new ones, just as we can perform Boolean operations on the decidable predicates. This will yield a language for approximate decision procedures.

The nondeterminism means that it won't behave quite like the deterministic decision procedures. We can characterize this algebraically. In **Set**, \mathbb{B} forms a Boolean algebra. Since the functor **Discrete** : **Set** \rightarrow **FSpc** is full, faithful, and preserves binary products and the terminal object, this lifts to an internal Boolean algebra (within **FSpc**) on \mathbb{B} the space. We can again lift these operations by the functor **NonDet** : **FSpc** \rightarrow **FSpc_{nd}**; for instance, the lifted version of the Boolean “and” operation ($\&\& : \mathbb{B} \times \mathbb{B} \xrightarrow{nd}_c \mathbb{B}$) is potentially **true** if its argument potentially takes on values whose conjunction is equal to **true**.

Proposition 5.3. *However, the operations on \mathbb{B} in **FSpc_{nd}** no longer form a Boolean algebra.*

Proof. In particular, there is the nondeterministic value **both** : $* \xrightarrow{nd}_c \mathbb{B}$ satisfying

$$\text{both}^*(\cdot = \text{true}) = \text{both}^*(\cdot = \text{false}) = \top,$$

such that

$$\text{both} \parallel !\text{both} = \text{both} \neq \text{true},$$

whereas in a Boolean algebra there is the identity $x \parallel !x = \text{true}$. □

Though \mathbb{B} doesn't form a Boolean algebra in **FSpc_{nd}**, it comes close!

Proposition 5.4. *The space \mathbb{B} forms a quasi-Boolean algebra (or de Morgan algebra) in \mathbf{FSpc}_{nd} , meaning that \mathbb{B} with $\&\&$, $\|$, **true**, and **false** forms a bounded distributive lattice, and $!$ is a de Morgan involution, in that it satisfies $!!x = x$ and $!(x \&\& y) = !x \| !y$.*

Proof. It is instructive to observe how the operations act on generalized points $\Gamma \xrightarrow{nd}_c \mathbb{B}$; we will use their equivalent representation as binary covers of Γ . Observe that

$$\begin{aligned} (P_1, Q_1) \&\& (P_2, Q_2) &= (P_1 \wedge P_2, Q_1 \vee Q_2) \\ (P_1, Q_1) \| (P_2, Q_2) &= (P_1 \vee P_2, Q_1 \wedge Q_2) \\ !(P, Q) &= (Q, P) \\ \mathbf{true} &= (\top, \perp) \\ \mathbf{false} &= (\perp, \top). \end{aligned}$$

We can use this to confirm the various laws, for instance, that **true** is the identity for $\&\&$,

$$\mathbf{true} \&\& (P, Q) = (\top \wedge P, \perp \vee Q) = (P, Q).$$

We similarly can confirm **false** is the identity for $\|$, and it is easy to observe that $\&\&$ and $\|$ are commutative and associative. Absorption of $\&\&$ and $\|$ follows from the similar absorption properties of opens, and likewise for their distributivity properties.

It remains to confirm that $!$ is a deMorgan involution. We have

$$!!(P, Q) = !(Q, P) = (P, Q)$$

and

$$\begin{aligned} !((P_1, Q_1) \&\& (P_2, Q_2)) &= !(P_1 \wedge P_2, Q_1 \vee Q_2) = (Q_1 \vee Q_2, P_1 \wedge P_2) \\ &= (Q_1, P_1) \| (Q_2, P_2) = !(P_1, Q_2) \| !(P_2, Q_2). \end{aligned}$$

□

That these nondeterministic Boolean operations form a quasi-Boolean algebra is not very interesting on its own. What makes truly binary covers powerful is that they admit quantification over compact/overt spaces, which is what we will now show.

5.2 Quantification over compact/overt spaces

When working with sets, if a predicate P on a set A is decidable and if A is finite, then $\forall a : A. P(a)$ and $\exists a : A. P(a)$ are decidable as well. The spatial analogue of the finite sets are the compact/overt spaces.

5.2.1 On compact/overt spaces

Definition 5.5. *A space A is compact if for every space Γ , the functor $- \times \top_A : \mathcal{O}(\Gamma) \rightarrow \mathcal{O}(\Gamma \times A)$ has a right adjoint $\forall_A : \mathcal{O}(\Gamma \times A) \rightarrow \mathcal{O}(\Gamma)$. That is, for each*

open $U : \mathcal{O}(\Gamma \times A)$, there is an open $\forall_A U : \mathcal{O}(\Gamma)$ such that for every $V : \mathcal{O}(\Gamma)$,

$$V \leq_\Gamma \forall_A U \quad \Longleftrightarrow \quad V \times \top_A \leq_{\Gamma \times A} U.$$

Similarly, a space A is *overt* if for every space Γ , $- \times \top_A$ has a left adjoint $\exists_A : \mathcal{O}(\Gamma \times A) \rightarrow \mathcal{O}(\Gamma)$. That is, for every open $U : \mathcal{O}(\Gamma \times A)$, there is an open $\exists_A U : \mathcal{O}(\Gamma)$ such that for every $V : \mathcal{O}(\Gamma)$,

$$\exists_A U \leq_\Gamma V \quad \Longleftrightarrow \quad U \leq_{\Gamma \times A} V \times \top_A.$$

These conditions are the definitions of universal and existential quantification in terms of adjoints, viewing Γ as some context and opens as truth values in a context. This definition of compactness is equivalent to the more common one, that every open cover has finite subcover.

Definition 5.6. A compact/overt space is a space A that is both compact and overt, as well as satisfying an additional property, that says for all Γ and all $P, Q : \mathcal{O}(\Gamma \times A)$, we have

$$\begin{aligned} \forall_A(P \vee Q) &\leq_\Gamma \forall_A P \vee \exists_A Q \\ \forall_A P \wedge \exists_A Q &\leq_\Gamma \exists_A(P \wedge Q). \end{aligned}$$

Assuming classical logic, all spaces are overt, so it is acceptable intuition (spatially) to conflate the compact/overt spaces with the compact ones.

These properties allow us to define syntax for quantification of Σ -valued continuous maps on compact/overt spaces:

$$\frac{\Gamma, x : A \vdash e : \Sigma \quad A \text{ compact}}{\Gamma \vdash \forall x \in A. e : \Sigma} \quad \forall\text{-}\Sigma\text{-FORM} \qquad \frac{\Gamma, x : A \vdash e : \Sigma \quad A \text{ overt}}{\Gamma \vdash_{nd} \exists x \in A. e : \Sigma} \quad \exists\text{-}\Sigma\text{-FORM}$$

These properties allow us to quantify binary covers over compact/overt spaces, too. That is, we can add some syntax

$$\frac{\Gamma, x : A \vdash_{nd} e : \mathbb{B} \quad A \text{ compact/overt}}{\Gamma \vdash_{nd} \forall x \in A. e : \mathbb{B}} \quad \forall\text{-}\mathbb{B}\text{-FORM}$$

$$\frac{\Gamma, x : A \vdash_{nd} e : \mathbb{B} \quad A \text{ compact/overt}}{\Gamma \vdash_{nd} \exists x \in A. e : \mathbb{B}} \quad \exists\text{-}\mathbb{B}\text{-FORM}$$

that behaves as we'd expect (for a quasi-Boolean algebra, at least). We compile this syntax by defining quantification functionals of the type $(\Gamma \times A \xrightarrow{nd}_c \mathbb{B}) \rightarrow (\Gamma \xrightarrow{nd}_c \mathbb{B})$.

For a compact/overt space A , we define a universal quantification functional

$$\begin{aligned} \forall_A &: (\Gamma \times A \xrightarrow{nd}_c \mathbb{B}) \rightarrow (\Gamma \xrightarrow{nd}_c \mathbb{B}) \\ \forall_A(P, Q) &\triangleq (\forall_A P, \exists_A Q), \end{aligned}$$

and confirm that this pair of opens is indeed covering with the derivation

$$\begin{aligned}
\top_\Gamma &\leq \forall_A(\top_{\Gamma \times A}) && (\forall_A \text{ adjointness}) \\
&\leq \forall_A(P \vee Q) && ((P, Q) \text{ cover } \Gamma \times A, \forall_A \text{ monotone}) \\
&\leq \forall_A P \vee \exists_A Q. && (A \text{ compact/overt})
\end{aligned}$$

We can similarly define the existential quantification functional

$$\begin{aligned}
\exists_A : (\Gamma \times A \xrightarrow{nd}_c \mathbb{B}) &\rightarrow (\Gamma \xrightarrow{nd}_c \mathbb{B}) \\
\exists_A(P, Q) &\triangleq (\exists_A P, \forall_A Q),
\end{aligned}$$

which indeed defines a binary cover for the same reason as \forall_A . By inspection, the two quantifiers are related by the law

$$\neg(\forall_A f) = \exists_A(\neg f).$$

While intuitively it seems that these operations indeed compute some nondeterministic quantification over compact spaces, we can make this statement precise by showing that these quantification functionals are adjoints to a weakening functional. To do this, we must first define an order on binary covers that relates their degrees of truth.

The quasi-Boolean algebra structure determines a preorder that can be placed on binary covers of any space. For binary covers (P_1, Q_1) and (P_2, Q_2) of a space Γ , we say $(P_1, Q_1) \leq (P_2, Q_2)$ if and only if both $P_1 \leq P_2$ and $Q_2 \leq Q_1$. We will call this *truth order*. (Note that truth order is different from the specialization order.)

For any spaces Γ and A , the weakening operation,

$$(- \circ \text{fst}) : (\Gamma \xrightarrow{nd}_c \mathbb{B}) \rightarrow (\Gamma \times A \xrightarrow{nd}_c \mathbb{B}),$$

which acts according to

$$(- \circ \text{fst})(U, V) = (U \times \top, V \times \top),$$

is monotone with respect to truth order. Now we can show that the quantification operators on binary covers deserve their names.

Theorem 5.7. *The existential and universal quantification functionals are left and right adjoints to weakening, respectively, with respect to truth order, i.e.,*

$$\exists_A \dashv (- \circ \text{fst}) \dashv \forall_A.$$

Proof. First we prove that \exists_A is left adjoint to weakening: given (P, Q) a binary cover

of $\Gamma \times A$ and (U, V) a binary cover of Γ , we must show

$$\begin{aligned} (P, Q) &\leq (- \circ \text{fst})(U, V) \quad \text{if and only if} \quad \exists_A(P, Q) \leq (U, V) \\ (P, Q) &\leq (U \times \top, V \times \top) \quad \text{if and only if} \quad (\exists_A P, \forall_A Q) \leq (U, V) \\ P &\leq U \times \top \text{ and } V \times \top \leq Q \quad \text{if and only if} \quad \exists_A P \leq U \text{ and } V \leq \forall_A Q. \end{aligned}$$

In the final form, we see that this follows from adjoint properties of the \exists_A and \forall_A operations that act on opens of $\Gamma \times A$.

The proof that \forall_A is a right adjoint is a mirror image of the proof regarding \exists_A . \square

5.2.2 On compact/overt subspaces

Sometimes, the space that we might want to quantify over could depend on some continuous variables in the context. For instance, we may want to quantify a binary cover $f : \mathbb{R} \times \mathbb{R} \xrightarrow{nd}_c \mathbb{B}$ over the unit simplex in $\mathbb{R} \times \mathbb{R}$ (i.e., the triangle bounded by $(0, 0)$, $(1, 0)$, and $(0, 1)$). We will describe a formalism whereby it will be possible to write this as

$$\forall x \in [0, 1]. \forall y \in [0, 1 - x]. f(x, y).$$

We handle this situation by considering certain spaces whose points represent compact/overt subspaces of some space.

First, we note the connection between the overt spaces and \mathcal{P}_\diamond , which represents (some) overt subspaces of a space A as points of the powerspace $\mathcal{P}_\diamond(A)$.

Proposition 5.8 (Theorem 32 of [43]). *Every point of $\mathcal{P}_\diamond(A)$ corresponds to an overt subspace of A .²*

Similarly, for each space A there is a powerspace $\mathcal{P}_\square(A)$ whose points correspond with compact subspaces of A . Vickers describes the construction in detail [45, 40], but we summarize the salient characteristics. There is a “necessity” modality $\square : \mathcal{O}(A) \rightarrow \mathcal{O}(\mathcal{P}_\square(A))$ that distributes over meets and directed joins (analogous to the “possibility” modality $\diamond : \mathcal{O}(A) \rightarrow \mathcal{O}(\mathcal{P}_\diamond(A))$ for the lower powerspace). Continuous maps $\Gamma \rightarrow_c \mathcal{P}_\square(A)$ are in bijective correspondence with inverse image maps $\mathcal{O}(A) \rightarrow \mathcal{O}(\Gamma)$ that preserve meets and directed joins. Like \mathcal{P}_\diamond , \mathcal{P}_\square is a strong monad.

The powerspace analogue of the compact/overt spaces is called the *Vietoris powerspace* $\mathcal{P}_{\square\diamond}$. Points of $\mathcal{P}_{\square\diamond}(A)$ correspond to compact/overt subspaces of A . The space $\mathcal{P}_{\square\diamond}$ has both a “possibility” modality $\diamond : \mathcal{O}(A) \rightarrow \mathcal{O}(\mathcal{P}_{\square\diamond}(A))$ a “necessity” modality $\square : \mathcal{O}(A) \rightarrow \mathcal{O}(\mathcal{P}_{\square\diamond}(A))$ that interact exactly as with the compact/overt spaces; that is, for any opens $P, Q : \mathcal{O}(A)$, the following laws hold:

$$\square(P \vee Q) \leq \square P \vee \diamond Q \tag{5.1}$$

$$\square P \wedge \diamond Q \leq \diamond(P \wedge Q). \tag{5.2}$$

² Specifically, the points of $\mathcal{P}_\diamond(A)$ are in bijective correspondence with the *weakly closed* (defined in [43]) overt subspaces of A .

Like for the lower powerspaces, there is the “positive” subspace $\mathcal{P}_{\square\Diamond}^+(A)$ of $\mathcal{P}_{\square\Diamond}(A)$ that additionally satisfies $\top \leq \Diamond\top$ and $\Box\perp \leq \perp$.

We can add some additional syntax to make it easier to describe opens (via the correspondence between Σ -valued maps $\Gamma \rightarrow_c \Sigma$ and opens $\mathcal{O}(\Gamma)$) with these modalities:

$$\frac{\Gamma \vdash s : \mathcal{P}_{\Diamond}(A) \quad \Gamma, x : A \vdash e : \Sigma}{\Gamma \vdash \exists x \in s. e : \Sigma} \exists\text{-}\Sigma\text{-SUB-FORM}$$

$$\frac{\Gamma \vdash s : \mathcal{P}_{\Box}(A) \quad \Gamma, x : A \vdash e : \Sigma}{\Gamma \vdash \forall x \in s. e : \Sigma} \forall\text{-}\Sigma\text{-SUB-FORM}.$$

These are compiled using the correspondence [37]

$$\Sigma \cong \mathcal{P}_{\Diamond}(\ast) \cong \mathcal{P}_{\Box}(\ast).$$

Replacing Σ with $\mathcal{P}_{\Diamond}(\ast)$ in $\exists\text{-}\Sigma\text{-SUB-FORM}$ and with $\mathcal{P}_{\Box}(\ast)$ in $\forall\text{-}\Sigma\text{-SUB-FORM}$, these rules just correspond to the bind operation of a strong monad M , which composes maps $\Gamma \rightarrow M(A)$ and $\Gamma \times A \rightarrow M(B)$ to produce $\Gamma \rightarrow M(B)$.

These properties allow us to quantify binary covers over compact/overt subspaces as well, implementing the syntax

$$\frac{\Gamma \vdash s : \mathcal{P}_{\square\Diamond}(A) \quad \Gamma, x : A \vdash_{nd} e : \mathbb{B}}{\Gamma \vdash_{nd} \forall x \in s. e : \mathbb{B}} \forall\text{-}\mathbb{B}\text{-SUB-FORM}$$

$$\frac{\Gamma \vdash s : \mathcal{P}_{\square\Diamond}(A) \quad \Gamma, x : A \vdash_{nd} e : \mathbb{B}}{\Gamma \vdash_{nd} \exists x \in s. e : \mathbb{B}} \exists\text{-}\mathbb{B}\text{-SUB-FORM}.$$

Given any binary cover $f : \Gamma \times A \xrightarrow{nd}_c \mathbb{B}$ of A , we define its universal quantification

$$\forall_{[\cdot]} f : \Gamma \times \mathcal{P}_{\square\Diamond}(A) \xrightarrow{nd}_c \mathbb{B}$$

$$\forall_{[\cdot]} f \triangleq \text{case}(p) \begin{cases} \iota[\lambda(\gamma, s). \forall x \in s. f(\gamma, x) = \text{true}] (_) & \Rightarrow \text{true} \\ \iota[\lambda(\gamma, s). \exists x \in s. f(\gamma, x) = \text{false}] (_) & \Rightarrow \text{false} \end{cases}.$$

We can confirm this pattern match is covering, i.e.,

$$\top \leq (\lambda(\gamma, s). \forall x \in s. f(\gamma, x) = \text{true}) \vee (\lambda(\gamma, s). \exists x \in s. f(\gamma, x) = \text{false})$$

with the derivation³

$$\begin{aligned}
\top &\leq \lambda(\gamma, s). \forall x \in s. \top && (\Box \text{ preserves } \top) \\
&\leq \lambda(\gamma, s). \forall x \in s. f(x) = \text{true} \vee f(x) = \text{false} \\
&\leq \lambda(\gamma, s). (\forall x \in s. f(x) = \text{true}) \vee \exists x \in s. f(x) = \text{false}. && (\text{law 5.1})
\end{aligned}$$

We can define $\exists_{[\cdot]} f : \mathcal{P}_{\Box\Diamond}(A) \xrightarrow{nd}_c \mathbb{B}$ by composing $\forall_{[\cdot]} f$ with Boolean negation.

Compact/overt spaces form a convenient class of spaces over which exhaustive reasoning is possible. The continuous image of a compact/overt space is compact/overt (similar to the fact that the image of a finite set under any map is finite). Like finite subsets, compact/overt subspaces are closed under finitary union, witnessed by a construction

$$\cup : \mathcal{P}_{\Box\Diamond}(A) \times \mathcal{P}_{\Box\Diamond}(A) \rightarrow_c \mathcal{P}_{\Box\Diamond}(A),$$

but not necessarily intersection. Naturally, a finite set viewed as a discrete space is compact/overt.

An alternate understanding of binary covers

An alternative interpretation of binary covers helps explain why the two opens in a binary cover get “opposite” treatments in the definitions of conjunction and disjunction. Rather than thinking of a binary cover of A as two open sets $P, Q : \mathcal{O}(A)$ such that (informally) $A \subseteq P \cup Q$, we can instead think of it as an open set P and a *closed* set \overline{Q} which is the “set-theoretic” complement of Q , since the complement of an open set is closed. Then the fact that (P, Q) is a binary cover means that $\overline{Q} \subseteq P$. Then conjunction computes intersections of the open and closed subspaces, while disjunction computes unions. Since we are simply “encoding” closed sets with their open complements, computing the “union” of closed subspaces just corresponds to taking an intersection of their open representatives. Hence, conjunction computes intersections of the second components.

This elicits the view of binary covers as “approximate” predicates, sandwiching a closed subspace inside an open one, with wiggle room for points in between. Any points in \overline{Q} (and thus also P) will definitely compute to **true**, while any points outside of P (and thus also outside \overline{Q}) will definitely compute to **false**, while in-between points, which are in P but not \overline{Q} , are allowed to compute either way.

³ The algebraic manipulations in this derivation are justified by the corresponding ones on $\mathcal{P}_{\Box\Diamond}(A)$. For instance $\top \leq \Box \top$ implies

$$\lambda(\gamma, s). \top \leq \lambda(\gamma, s). \forall x \in s. \top.$$

5.3 Binary covers on \mathbb{R}

Here, we will describe some compact/overt subspaces and binary covers of \mathbb{R} and use these to solve some approximate decision problems on \mathbb{R} with the calculus of binary covers.

There is a construction

$$\text{HeineBorel} : \{(a, b) : \mathbb{R} \times \mathbb{R} \mid a < b\} \rightarrow_c \mathcal{P}_{\square\Diamond}^+(\mathbb{R})$$

that demonstrates for endpoints (a, b) that the closed interval from a to b is compact/overt and inhabited (a constructive analogue of the Heine-Borel theorem) [46].

We can define binary covers that give approximate order comparisons on \mathbb{R} by defining, for each $\varepsilon > 0$,

$$\begin{aligned} [\cdot < \cdot]_\varepsilon &: \mathbb{R} \times \mathbb{R} \xrightarrow{nd}_c \mathbb{B} \\ [x < y]_\varepsilon &\triangleq \text{case}(x, y) \begin{cases} \iota[\lambda(x, y). x < y] (_) & \Rightarrow \text{true} \\ \iota[\lambda(x, y). x > y - \varepsilon] (_) & \Rightarrow \text{false}. \end{cases} \end{aligned}$$

We can use this to approximately query whether an arbitrary continuous map $f : \mathbb{R} \rightarrow_c \mathbb{R}$ has no roots on a compact/overt subspace of \mathbb{R} (such as a closed interval), with the binary cover

$$\begin{aligned} \text{no_roots}_\varepsilon &: \mathcal{P}_{\square\Diamond}(\mathbb{R}) \xrightarrow{nd}_c \mathbb{B} \\ \text{no_roots}_\varepsilon(s) &\triangleq \forall x \in s. [0 < |f(x)|]_\varepsilon. \end{aligned}$$

If we observe **true** of the result on input s , then indeed f has no roots on s , while if we observe **false** of the result, then there is some $x \in s$ such that $|f(x)| < \varepsilon$.

Perhaps we have two maps $f, g : \mathbb{R} \times \mathbb{R} \rightarrow_c \mathbb{R}$ and wish to confirm, for some compact/overt space $s : \mathcal{P}_{\square\Diamond}(\mathbb{R})$ that there is some $x \in [0, 1]$ such that for every $y \in s$, $f(x, y)$ is positive while $g(x, y)$ is negative. Then

$$\begin{aligned} \text{apart}_\varepsilon &: \mathcal{P}_{\square\Diamond}(\mathbb{R}) \xrightarrow{nd}_c \mathbb{B} \\ \text{apart}_\varepsilon(s) &\triangleq \exists x \in [0, 1]. \forall y \in s. [0 < f(x, y)]_\varepsilon \wedge [g(x, y) < 0]_\varepsilon. \end{aligned}$$

approximately decides this, where $[0, 1]$ is used as shorthand for $\text{HeineBorel}(0, 1)$. Like the previous query, this one has only false negatives (and its negation has only false positives). If we observe **true**, then it is certainly true, while if we observe **false**, then for every $x \in [0, 1]$ there is a $y \in [0, 1]$ such that either $f(x, y) < \varepsilon$ or $g(x, y) > -\varepsilon$.

5.4 Approximate root-finding

Recall the approximate root-finding program from the introduction,

$$\begin{aligned} \text{roots}_f &: * \xrightarrow{nd}_c \{ * \mid \forall x \in K. f(x) \neq 0 \} + \{ x : K \mid |f(x)| < \varepsilon \} \\ \text{roots}_f &\triangleq \text{case}(\text{tt}) \begin{cases} \iota[\exists x \in K. |f(x)| < \varepsilon] (y) & \Rightarrow \text{inr}(\text{simulate}(y)) \\ \iota[\forall x \in K. f(x) \neq 0] (n) & \Rightarrow \text{inl}(n) \end{cases}. \end{aligned}$$

Now that we have explained compact/overt spaces and binary covers, we are prepared to understand this code.

We begin by confirming that the two cases cover the entire input space. We can define a binary cover

$$\begin{aligned} [\cdot \neq 0]_\varepsilon &: \mathbb{R} \xrightarrow{nd}_c \mathbb{B} \\ [x \neq 0]_\varepsilon &\triangleq \text{case}(x) \begin{cases} \iota[\cdot \neq 0] (_) & \Rightarrow \text{true} \\ \iota[\lambda x. |x| < \varepsilon] (_) & \Rightarrow \text{false} \end{cases} \end{aligned}$$

that approximately determines whether a real number is nonzero.

Then the binary cover

$$\begin{aligned} \text{no_roots}'_\varepsilon &: * \xrightarrow{nd}_c \mathbb{B} \\ \text{no_roots}'_\varepsilon &\triangleq \forall x \in K. [x \neq 0]_\varepsilon \end{aligned}$$

is defined by the same pair of opens as in the definition of roots_f , and hence those opens cover $*$.

It remains to define simulate , which in this instance is a map

$$\text{simulate} : \{ * \mid \exists x \in K. |f(x)| < \varepsilon \} \xrightarrow{nd}_c \{ x : K \mid |f(x)| < \varepsilon \},$$

but in general, for any overt space A and open $U : \mathcal{O}(A)$, defines a map

$$\text{simulate} : \{ * \mid \exists x \in A. U(x) \} \xrightarrow{nd}_c \{ x : A \mid U(x) \}$$

that, given the existence of some values that satisfy a property U of A , can non-deterministically simulate those values. It is defined by the inverse image map

$$\begin{aligned} \text{simulate}^* &: \mathcal{O}(\{A \mid U\}) \rightarrow \mathcal{O}(\{ * \mid \exists_A U \}) \\ \text{simulate}^*(V) &\triangleq \exists_A(V \wedge U). \end{aligned}$$

Proposition 5.9. *The inverse image map simulate^* preserves joins and \top .*

Proof. First, we confirm that the output lies in the open subspace U of A , that is, that simulate^* is gives equivalent results on inputs V and W satisfying $V \wedge U = W \wedge U$. This is straightforward because simulate^* immediately applies $\cdot \wedge U$ to its argument.

The map preserves joins since it is the composition $\exists_A \circ (\cdot \wedge U)$, both of which preserve joins. It preserves \top (or equivalently, U), since

$$\text{simulate}^*(\top) = \exists_A(U),$$

which is equal to \top in $\{* \mid \exists_A U\}$. □

The approximate root-finding program **roots_f** accomplishes the relatively complicated task of approximate root-finding over a very general class of functions with a very short definition that works by composing some constructs from this chapter and the previous one. Most of the real computational work is accomplished by the formal proof that K is compact/overt. This gives evidence that the property of being compact/overt corresponds to a very general and composable computational interface for exhaustive search.⁴

⁴Escardó [10] discusses the computational significance of compactness.

Chapter 6

Related work

6.1 Alternative theories of constructive topology

Synthetic topology gives an alternative computable interpretation of topology that differs subtly from formal topology and locale theory [9, 20]. Synthetic topology gives an alternative intuition about overlapping patterns. In this understanding, data types from a (functional) programming language serve as spaces, and open sets are semidecidable predicates that take values of some type as arguments; for points in the open set, the semidecider must halt, whereas for points in the closed complement, the semidecider should not halt.

In this case, an open cover $U \leq \bigvee_{i:I} V_i$ means that whenever U halts, then some semidecider V_i also halts.

This interpretation is limited to (classically) spatial locales (locales that also make sense as conventional topological spaces) and covers indexed by recursively enumerable sets. But it gives an alternative computational interpretation of overlapping pattern matching: to run a pattern match, run the semidecidable for each case in the pattern match *concurrently*. Since the cases must cover the entire space, the semidecider for one case will eventually halt and at that point proceed into the branch corresponding to that case. However, implementing overlapping patterns naively with this computational interpretation does not seem particularly efficient.

Taylor’s *abstract stone duality* [36] is a theory/language for locally compact topological spaces. The additional dependency of local compactness means that all spaces are exponentiable, which is convenient as the formulation makes use of such exponentials.

6.2 Related programming formalisms involving continuity, partiality, or nondeterminism

Marcial-Romero and Escardó [22] define a language with real number computation which has a foundational family of functions $\text{rtest}_{a,b}$ which map real numbers to nondeterministic Boolean values. This language also admits nontermination, so it

denotes its types as Hoare powerlocales (which corresponds to \mathcal{P}_\diamond in this thesis). Establishing totality requires reasoning within their operational model. Our work differs in a few ways: we describe a general principle for constructing nondeterministic maps, and we are able to use the positive Hoare powerlocale, which guarantees totality (at the cost of requiring formal proofs of covering).

Escardó’s defines a language “Real PCF” with a denotational semantics in terms of cpo’s, in which there is an operation known as a “parallel conditional,” which corresponds to the internal “or” operation on the Sierpinski space $\vee_\Sigma : \Sigma \times \Sigma \rightarrow_c \Sigma$ in our formalism. Parallel conditionals applied to construct total deterministic functions on the real numbers, which differs from most of our examples, whose computations are total but nondeterministic.

Similarly, Tsuiki’s work on computation with Gray-code based real numbers [38] is based on “indeterministic” computation, where potentially nonterminating computations must be interleaved, and those that terminate must agree in their answers.

6.3 Overlapping pattern matching

We are unaware of any other notion of pattern matching that permits patterns where determining membership is undecidable (without jeopardizing totality).

Coquand [6] gives a topologically motivated explanation of pattern matching for dependently typed functional programming, describing patterns as (disjoint) partitions of a space.

Müller [23] describes a system for exact real arithmetic that has a datatype of “lazy Booleans” analogous to our \mathbb{B}_\perp , as well as an n -ary `choose` operation on lazy Booleans that would be analogous in our language to

$$\begin{aligned} \text{choose}_n &: \prod_{i:\text{Fin}(n)} \mathbb{B}_\perp \xrightarrow{nd,p}_c \text{Fin}(n) \\ \text{choose}_n(b) &\triangleq \text{case}(b) \left\{ [i : \text{Fin}(n)] \quad \iota[[i \mapsto \cdot = \text{true}]](_) \Rightarrow i \right. \end{aligned}$$

Dijkstra [8] introduces “guarded commands,” a language construct for imperative programming languages, where a branch can be chosen nondeterministically from a list of statements each guarded by a Boolean expression.

6.4 Binary covers

Some of the modified Boolean spaces give topological interpretations of well-known fuzzy logics: \mathbb{B}_\perp corresponds to Kleene’s three-valued logic [17], $\mathcal{P}_\diamond^+(\mathbb{B})$ to Priest’s logic of paradox [29], and $\mathcal{P}_\diamond(\mathbb{B})$ to Dunn/Belnap’s four-valued logic [13].

dReal is a tool that allows computation of approximate truth values over the real numbers [12], allowing order comparisons and bounded quantifiers. The calculus of binary covers presented here, when restricted to \mathbb{R} , provides similar computational

abilities but with a different foundational framework. In a sense, it shows how it is possible to generalize the theory behind dReal to spaces other than \mathbb{R} .

6.5 Implementation

Using the theory of abstract stone duality, Bauer and Taylor have developed a computer implementation, *Marshall*, that supports exact real arithmetic computation based on the Dedekind real numbers [3].

Some aspects of formal topology have been formalized within the Matita theorem prover (another theorem prover implementing Martin-Löf type theory) [1, 5]. Their library is not as concrete as ours; as far as we are aware, they do not construct any spaces (in an empty context), nor define any continuous maps other than identity maps and compositions of continuous maps, focusing instead on the relation between various notions of spaces and maps.

Chapter 7

Discussion and conclusion

The programs described in this thesis conceivably support reasonably efficient implementations, thanks to their use of formal topology and (predicative) locale theory. Formal proofs that inverse image maps preserve meets and joins not only confirm that they indeed define continuous maps but also provide computational content.

Overlapping pattern matches seem useful for a variety of purposes but in particular highlight the importance of nondeterminism in constructing programs that manipulate spaces. The approximate root-finding procedure in the introduction and the case study of binary covers demonstrate useful nondeterministic programs that can be constructed with overlapping patterns.

7.1 Coq implementation

We developed a library for formal topology in the predicative fragment of Coq (i.e., without use of the `Prop` universe) [32]. The library includes the following constructions:

- the one-point space $*$ and the unique maps $A \rightarrow_c *$
- discrete spaces and a construction that creates a continuous map $f : A \rightarrow_c B$ from a discrete space A to any space B , given for each $a : A$ a point $f_a : * \rightarrow_c B$
- arbitrary type-indexed sum spaces, their injections, and their universal property
- arbitrary type-indexed product spaces, their projections, and their universal property
- metric “completion” spaces of metric sets, and a construction for lifting Lipschitz continuous functions from metric sets to metric completions
- lifted spaces (e.g., A_\perp), with implementations of $\perp : * \rightarrow_c A_\perp$ and $\text{up} : A \rightarrow_c A_\perp$ as continuous maps
- definitions regarding subspaces, and some particular facts about open and closed subspaces

- definition of a predicative analogue of locales, and facts relating them to formal spaces
- a simplified overlapping pattern matching construction, where patterns must correspond to opens of the space of the scrutinee

7.2 Programming with a theorem prover

A continuous map $f : A \rightarrow_c B$ is defined by an inverse image map $f^* : \mathcal{O}(B) \rightarrow \mathcal{O}(A)$ that preserves joins, \top , and binary meets. We can view f^* as a complete *specification* of the behavior of the behavior of f , and it is the formal (and constructive) proofs that f^* preserves the relevant structure that serve as the *implementation* of the computational behavior. It is not enough only to know informally that f^* preserves the relevant structure. At the same time, different formal proofs still give equivalent continuous maps, so (disregarding matters such as computational efficiency) it does not matter which formal proof is given.

We believe that theorem provers such as Coq are well-suited for programming in this manner. One can define in Gallina the inverse image map f^* , and then use the theorem-proving interface to prove that f^* preserves the relevant structure. For instance, in our Coq development, we made heavy use of “proof-relevant” typeclass-based rewriting¹. Without tools such as this, defining the formal proofs may be potentially tedious. Since the formal proofs are irrelevant to reasoning about programs (beyond the fact that they exist), their definitions can be left opaque (until it is time to run/extract programs).

It is possible to avoid the need for proofs by restricting to the language defined in Section 2.4, which basically only allows composition of existing maps. There are also no proof obligations in the pattern matching construction for partial and non-deterministic maps, but there are proof obligations for constructing either total or deterministic maps. The proof obligations needed for pattern families is always to show that they belong to a particular Grothendieck pretopology. Given the compositional structure of Grothendieck pretopologies, it should be convenient to use a theorem prover to derive pattern families. This overlapping case splitting would be similar to using the `destruct` tactic in Coq to split a term in an inductive type into possible cases. Viewing pattern families in this way, the obligation that they belong to a particular Grothendieck pretopology is more like a part of the program structure, rather than a side condition that must be verified independently of the program syntax.

7.3 Future work

We are broadly interested in understanding how constructive topology can inform the design of programming languages for continuous systems.

¹That is, rewriting of terms in `Type` rather than `Prop`. This feature was added to Coq in version 8.5.

Dependent types

For verification of systems, it is desirable to have a full dependently typed language for expressing programming tasks. For instance, reconsider the problem of an autonomous car that must decide whether or not to stop at a traffic light that has just turned yellow. We can model the problem as such: we have spaces `state_before`, `state_after` and `decision` that describe the possible states of the car when the light turns yellow and when it turns red, respectively, and the decisions that can be made when the light turns yellow; a dynamics function `dynamics : state_before × decision →c state_after`; and an open `safe : \mathcal{O} (state_after)` that indicates if a given state is outside of the intersection, and thus safe. We want to ensure that for every possible car state $s : \text{state_before}$ there is a decision $d : \text{decision}$ such that `safe(dynamics(s, d))`, but the decision must be allowed to be made nondeterministically. With dependent types, we could imagine encoding this task as

$$\prod_{s:\text{state_before}} \mathcal{P}_{\diamond}^+ \left(\sum_{d:\text{decision}} \text{safe}(\text{dynamics}(s, d)) \right).$$

It ought to be possible to accomplish this by taking the gros topos of sheaves over the site generated by the open cover topology [11]. We do not yet understand how to extend the partiality and nondeterminism monads to the entire topos, but believe that approaches from synthetic topology might apply.

Probabilistic programming

Constructive topology provides a good setting for phrasing probability theory in a constructive manner, and so it should inform the theory and semantics of probabilistic programming. From any formal space A one can construct a space $\mathcal{R}(A)$ of probability distributions on A [44, 47].

Simpson [33] describes “random locales,” which are formal spaces which are not spatial, meaning that they have no analogue in classical topology. These random locales characterize the properties that occur with probability 1. For instance, Simpson describes the sublocale **Ran** of $\prod_{:\mathbb{N}} \mathbb{B}$ of random sequences of Boolean values. On this sublocale, the pattern match

$$\begin{aligned} \text{geometric} &: \mathbf{Ran} \rightarrow_c \mathbb{N} \\ \text{geometric}(x) &\triangleq \text{case}(x) \left\{ [n : \mathbb{N}] \quad \iota[[n \mapsto \cdot = \text{true}] \wedge \bigwedge_{k < n} [k \mapsto \cdot = \text{false}]] (_) \Rightarrow n \right. \end{aligned}$$

which finds the first occurrence of a `true` in the sequence defines a valid continuous map (it would not if the codomain were $\prod_{:\mathbb{N}} \mathbb{B}$). Although **Ran** has no points $* \rightarrow_c \mathbf{Ran}$, so this `geometric` cannot be run on any deterministic inputs, **Ran** *does* have nondeterministic points $* \xrightarrow{nd}_c \mathbf{Ran}$, so `geometric` can be run on nondeterministic values.

Bibliography

- [1] Andrea Asperti, Maria Emilia Maietti, Claudio Sacerdoti Coen, Giovanni Sambin, and Silvio Valentini. Formalization of formal topology by means of the interactive theorem prover Matita. In *International Conference on Intelligent Computer Mathematics*, pages 278–280. Springer, 2011.
- [2] Andrej Bauer. Realizability as the connection between computable and constructive mathematics. In *Lecture notes for a tutorial at a satellite seminar of CCA 2004*, 2005.
- [3] Andrej Bauer. Efficient computation with Dedekind reals. In *Fifth International Conference on Computability and Complexity in Analysis, pages i–vi, Hagen, Germany*, 2008.
- [4] Francesco Ciraulo, Maria Emilia Maietti, and Giovanni Sambin. Convergence in formal topology: a unifying notion. *Journal of Logic and Analysis*, 5, 2013.
- [5] Claudio Sacerdoti Coen and Enrico Tassi. Formalising overlap algebras in Matita. *Mathematical Structures in Computer Science*, 21(04):763–793, 2011.
- [6] Thierry Coquand. Pattern matching with dependent types. In *Proceedings of the Workshop on Types for Proofs and Program*, pages 66–79, 1992.
- [7] Thierry Coquand, Giovanni Sambin, Jan Smith, and Silvio Valentini. Inductively generated formal topologies. *Annals of Pure and Applied Logic*, 124(1):71–106, 2003.
- [8] Edsger W. Dijkstra. Nondeterminacy and formal derivation of programs. *Comm. ACM*, 18(8):453–457, 1975.
- [9] Martín Escardó. Synthetic topology: of data types and classical spaces. *Electronic Notes in Theoretical Computer Science*, 87:21–156, 2004.
- [10] Martín Escardó. Infinite sets that admit fast exhaustive search. In *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*, pages 443–452. IEEE, 2007.
- [11] Michael P Fourman. Continuous truth I: Non-constructive objects. *Studies in Logic and the Foundations of Mathematics*, 112:161–180, 1984.

- [12] Sicun Gao, Jeremy Avigad, and Edmund M Clarke. Delta-decidability over the reals. In *Logic in Computer Science (LICS), 2012 27th Annual IEEE Symposium on*, pages 305–314. IEEE, 2012.
- [13] Siegfried Gottwald. Many-valued logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2015 edition, 2015.
- [14] Peter T Johnstone. *Topos theory*. Academic Press, 1977.
- [15] Peter T Johnstone. *Sketches of an elephant: A topos theory compendium*, volume 2. Oxford University Press, 2002.
- [16] Tatsuji Kawai. Localic completion of uniform spaces, 2017.
- [17] Stephen Cole Kleene. *Introduction to metamathematics*. North-Holland, 1952.
- [18] Leslie Lamport. Buridan’s principle. *Foundations of Physics*, 42/8:1056–1066, October 1984.
- [19] S.M. Lane and I. Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Hochschultext / Universitext. Springer-Verlag, 1992.
- [20] Davorin Lešnik. Synthetic topology and constructive metric spaces. *Diss. University of Ljubljana*, 2010.
- [21] Maria Emilia Maietti and Giovanni Sambin. Why topology in the minimalist foundation must be pointfree. *Logic and Logical Philosophy*, 22(2):167–199, 2013.
- [22] J Raymundo Marcial-Romero and Martín H Escardó. Semantics of a sequential language for exact real-number computation. *Theoretical Computer Science*, 379(1-2):120–141, 2007.
- [23] Norbert Th Müller. Enhancing imperative exact real arithmetic with functions and logic. In *Computability and Complexity in Analysis, Ljubljana*, 2009.
- [24] Russell O’Connor. Certified exact transcendental real number computation in Coq. In *Theorem Proving in Higher Order Logics*, pages 246–261. Springer, 2008.
- [25] Erik Palmgren. Predicativity problems in point-free topology. In *Proceedings of the Annual European Summer Meeting of the Association for Symbolic Logic, held in Helsinki, Finland*, pages 221–231, 2003.
- [26] Erik Palmgren. A constructive and functorial embedding of locally compact metric spaces into locales. *Topology and its Applications*, 154(9):1854–1880, 2007.
- [27] Pavel Pančekha, Alex Sanchez-Stern, James R Wilcox, and Zachary Tatlock. Automatically improving accuracy for floating point expressions. *ACM SIGPLAN Notices*, 50(6):1–11, 2015.

- [28] Maria Cristina Pedicchio and Walter Tholen. *Categorical foundations: special topics in order, topology, algebra, and sheaf theory*, volume 97. Cambridge University Press, 2004.
- [29] Graham Priest. The logic of paradox. *Journal of Philosophical logic*, 8(1):219–241, 1979.
- [30] Giovanni Sambin and Silvio Valentini. Building up a toolbox for Martin-Löf intuitionistic type theory: subset theory. In Giovanni Sambin and Jan Smith, editors, *Twenty-five years of Constructive Type Theory*, pages 221–244. Oxford Logic Guides, 1998.
- [31] Robert Sedgewick and Kevin Wayne. Introduction to programming in Java: An interdisciplinary approach.
- [32] Benjamin Sherman, Luke Sciarappa, and Clément Pit-Claudel. Coq library for formal topology, May 2017. Available at <https://doi.org/10.5281/zenodo.581231>.
- [33] Alex Simpson. Measure, randomness and sublocales. *Annals of Pure and Applied Logic*, 163(11):1642–1659, 2012.
- [34] Alex K Simpson. Lazy functional algorithms for exact real functionals. In *International Symposium on Mathematical Foundations of Computer Science*, pages 456–464. Springer, 1998.
- [35] Bas Spitters. Locatedness and overt sublocales. *Annals of Pure and Applied Logic*, 162(1):36–54, 2010.
- [36] Paul Taylor. A lambda calculus for real analysis. *Journal of Logic and Analysis*, 2, 2010.
- [37] Christopher F Townsend. On the parallel between the suplattice and preframe approaches to locale theory. *Annals of Pure and Applied Logic*, 137(1-3):391–412, 2006.
- [38] Hideki Tsuiki. Real number computation through gray code embedding. *Theor. Comput. Sci.*, 284(2):467–485, July 2002.
- [39] Steven Vickers. *Topology via logic*. Cambridge University Press, 1989.
- [40] Steven Vickers. The double powerlocale and exponentiation: a case study in geometric logic. *Theory and Applications of Categories*, 12(13):372–422, 2004.
- [41] Steven Vickers. Localic completion of generalized metric spaces I. *Theory and Applications of Categories*, 14(15):328–356, 2005.
- [42] Steven Vickers. Some constructive roads to Tychonoff. *From Sets and Types to Topology and Analysis: Towards Practicable Foundations for Constructive Mathematics*, 48:223, 2005.

- [43] Steven Vickers. Sublocales in formal topology. *The Journal of Symbolic Logic*, 72(02):463–482, 2007.
- [44] Steven Vickers. A localic theory of lower and upper integrals. *Mathematical Logic Quarterly*, 54(1):109–123, 2008.
- [45] Steven Vickers. The connected Vietoris powerlocale. *Topology and its Applications*, 156(11):1886–1910, 2009.
- [46] Steven Vickers. Localic completion of generalized metric spaces II: Powerlocales. *Journal of Logic and Analysis*, 1, 2009.
- [47] Steven Vickers. A monad of valuation locales. *Preprint at <http://www.cs.bham.ac.uk/~sjv/Riesz.pdf>*, 2011.
- [48] Frank Waaldijk. On the foundations of constructive mathematics—especially in relation to the theory of continuous functions. *Foundations of Science*, 10(3):249–324, 2005.
- [49] K. Weihrauch. *A simple introduction to computable analysis*. Informatik-Berichte. Fernuniv., Fachbereich Informatik, 1995.

List of symbols

- \cdot_{\perp} The lifting monad on spaces. 44, 53, 54, 59, 60, 67, 76
- $A \xrightarrow{nd,p}_c B$ The set of nondeterministic and partial maps from a space A to a space B . 47, 50, 53, 61, 62, 70, 92
- $A \cong B$ Objects A and B in a category are isomorphic (usually, either two sets are isomorphic or two spaces are homeomorphic). 23, 25, 26, 29, 30, 34, 35, 51, 62, 67, 73, 78, 85, 104
- $A \xrightarrow{nd}_c B$ The set of nondeterministic maps from a space A to a space B . 44–46, 57, 59, 63, 73, 80–88, 97, 103
- $A \xrightarrow{p}_c B$ The set of partial maps from a space A to a space B . 39–41, 43, 44, 53, 59, 60, 64, 70, 72, 75, 76
- $A \rightarrow_c B$ The set of continuous maps from a space A to a space B . 11, 12, 15, 23, 25, 26, 28–32, 34, 35, 37, 38, 40, 41, 43, 45–53, 55, 58–63, 69–71, 73, 74, 76–78, 84–87, 92, 95–97
- $a \downarrow b$ The meet (intersection) of two basic opens a and b of a space A (which is an open of A). 21–24, 33, 34, 38, 53
- $a \equiv b$ Intensional equality of a and b (which must have the same type). 16, 21, 29–32, 42, 64–66
- $a \triangleleft U$ The basic open a of a space A is covered by the open U of A . 20–23, 28, 30, 32–34, 38, 42, 45, 53
- $f \leq g$ The truth order implication f implies g for binary covers $f, g : \Gamma \xrightarrow{nd}_c \mathbb{B}$. 83, 84
- $f \otimes g$ The parallel composition $f \otimes g : A \times X \rightarrow_c B \times Y$ of $f : A \rightarrow_c B$ and $g : X \rightarrow_c Y$. 49, 52, 62, 63, 67, 74, 75
- f^* The inverse image map $f^* : \mathcal{O}(B) \rightarrow \mathcal{O}(A)$ of a map of spaces $f : A \rightarrow_{\text{?}} B$. 25
- $f_!$ The direct image map $f_! : \mathcal{O}(A) \rightarrow \mathcal{O}(B)$ of an open map $f : A \rightarrow_o B$. 48
- $p \Rightarrow e$ A branch in a pattern match, where p is a pattern such that, if matched, the expression e results. 57–61, 68, 70, 74–78, 85, 87, 88, 92, 97

- $x \Vdash U$ A point x in a space A lies in an open U of A . 19, 20, 22, 23, 25, 27–29, 38, 43, 71
- $\iota[U]$ The open embedding $\iota[U] : \{A \mid U\} \hookrightarrow A$ of an open subspace U of A into the entire space. 50, 52, 57–59, 61, 85, 87, 88, 92, 97
- $\diamond U$ The direct image of an open U of A in its partial and nondeterministic powerspace $\mathcal{P}_\diamond(A)$ or some subspace of it. 46, 47, 84, 85
- $\mathcal{O}(A)$ The (large) set of opens of a space A . 21–31, 33–35, 39–53, 65, 70, 71, 78, 80–82, 84–86, 88, 96, 97, 103
- $\mathcal{O}_B(A)$ The set of basic opens of a space A . 20–23, 29–31, 42–44, 46, 47, 53
- \iff If and only if. 48, 49, 51, 82
- Ω The (large) set of propositions. 16, 20–23, 27–29, 32, 38–40
- Σ The Sierpinski space, $\Sigma \cong *_\perp \cong \mathcal{P}_\diamond(*)$. 7, 29, 35, 36, 71, 82, 85
- $_$ A wildcard in a pattern, which matches anything. 59, 67, 76, 78, 85, 87, 88, 92, 97
- $*$ Either the type with one element or the space with one point. 22, 23, 26, 27, 29, 40, 41, 46, 57, 59, 67, 78, 80, 85, 88, 89, 95, 97, 104, 105
- \mathbb{B} The type/set/discrete space of Boolean values. 11, 15, 18, 26, 29, 31, 35, 37, 39, 44, 59, 71, 80–88, 92, 97, 103–105
- \mathbb{N} The type/set/discrete space of natural numbers. 17, 18, 97
- \mathbb{Q} The set/discrete space of rational numbers. 17–20, 27, 32–34, 61
- \mathbb{R} The space of real numbers. 3, 7, 8, 11–15, 17–20, 26, 27, 32, 34, 35, 37, 39, 44, 58, 59, 61, 76, 79, 84, 87, 88, 92, 93
- $\&\&$ Boolean and, $\&\& : \mathbb{B} \times \mathbb{B} \rightarrow_c \mathbb{B}$. 80, 81
- \parallel Boolean or, $\parallel : \mathbb{B} \times \mathbb{B} \rightarrow_c \mathbb{B}$. 80, 81
- \mathcal{P}_\diamond^+ The nondeterministic powerspace monad, $\mathcal{P}_\diamond^+ : \mathbf{FSpc} \rightarrow \mathbf{FSpc}$. 45–47, 54, 59, 92, 97
- \mathcal{P}_\diamond The partial & nondeterministic powerspace monad, $\mathcal{P}_\diamond : \mathbf{FSpc} \rightarrow \mathbf{FSpc}$. 47, 50, 54, 84, 85, 92, 104
- \mathcal{P}_\square The compact powerspace monad, $\mathcal{P}_\square : \mathbf{FSpc} \rightarrow \mathbf{FSpc}$. 84, 85
- $\mathcal{P}_{\square\diamond}$ The compact/overt powerspace monad, $\mathcal{P}_{\square\diamond} : \mathbf{FSpc} \rightarrow \mathbf{FSpc}$. 84–87
- \mathcal{U} The type of types, $\mathcal{U} : \mathcal{U}$. 16, 23–25, 34, 70

! Boolean negation, $! : \mathbb{B} \rightarrow_c \mathbb{B}$. 80, 81

\Downarrow The direct image of an open in its lifting, $\Downarrow : A \rightarrow A_\perp$. 41–44, 46, 53

false Boolean false, $\text{false} : * \rightarrow_c \mathbb{B}$. 11, 17–20, 27, 29, 31, 39, 44, 45, 59, 71, 80, 81, 85–88, 97

true Boolean true, $\text{true} : * \rightarrow_c \mathbb{B}$. 11, 17, 18, 20, 27, 29, 31, 39, 44, 45, 59, 60, 71, 78, 80, 81, 85–88, 92, 97

up The open embedding of a space in its lifting, $\text{up} : A \hookrightarrow A_\perp$. 53, 59, 60, 67, 76, 95

\emptyset The empty space. 78