

A Bunch of Set Theory and a Bit of Logic

CPSC 509: Programming Language Principles

Ronald Garcia*

14 January 2014

(Time Stamp: 12:57, Friday 25th March, 2022)

1 Representing the world using sets

Set theory is the “machine language” of mathematics. If you think about it, every program that you have ever run on a computer has ultimately produced instructions that get sent to a CPU and the CPU just churns through them. Those CPU instructions are really really low-level: add this 32 bit number to that 32 bit number, grab these bits from memory and move them somewhere else, check this number to see if it’s zero, and grab these other set of bits from memory and interpret them like a machine instruction. Somehow, these little primitive instructions make it possible to write programs that download pictures of kittens from the Internet and thereby make your life more fulfilling. At the end of the day, though, it’s all bits: 1’s and 0’s. Your kittens are *represented* by a bunch of bits that your graphics hardware knows how to translate into a vision of furry joy. Your tragic poetry is written in ASCII, which is really a bunch of characters, each of which is represented by a number, which in turn is represented by bits. So your computer consistently operates on *representations* of the things that you really care about.

Set theory is like that. It’s *painfully* low-level, and it doesn’t understand high-level concepts like programming languages, or numbers for that matter, but it’s very very powerful! Many mathematicians and logicians take set theory as *the* foundational tool for building representations of things. If you can’t represent your object using the machinery of set theory, then those mathematicians would likely deem them not so mathematical. This viewpoint is not universal, but it’s safe to say that as I write this it is dominant, and the thesis guiding this class is that, regardless of its supremacy and despite its warts, it’s *useful*.

2 Abstractions make this tractable

On computers and in mathematics, we keep this low-level stuff from exploding our brains by building up layers of *abstraction*. Much of the time we just ignore what’s going on at the low-levels, but instead think in terms of higher-level concepts that we really care about. As long as your abstractions are well-designed and clearly explained, you can ignore the details. If your abstractions are “leaky”, then you might write down things that make no sense to you but say something meaningful at the lower levels [Spolsky, 2002]. For example, I’m willing to believe that the letter ‘a’ comes before ‘b’, so in a sense: ‘a’ < ‘b’, and the C programming language tells me this, but why would I think that ‘)’ < ‘+’ ??? The C language tells me that this is true too, and if I didn’t know anything about the *ASCII encoding of characters as numbers*, then I wouldn’t understand what’s going on. In short, C’s abstraction of characters is leaky. If the language didn’t let you compare characters in terms of their underlying machine representations, this weird stuff wouldn’t happen.¹ Later, we’ll see some examples of leaky abstractions in set theory.

*© 2014 Ronald Garcia.

¹In fact, some errors in C code arise from assuming that ASCII is always the relevant character encoding! See for example <https://stackoverflow.com/questions/16400009/why-the-char-comparison-ifc-a-c-z-is-not-portable>

Now don't get me wrong, sometimes it's useful to know what's going on under the hood, if only so you can debug problems or implement abstractions of your own. This is true on both computers and in mathematics. In essence, the ideal is to be a *full stack mathematician* who can work efficiently at a high level, but understand the underpinnings that ensure that this high-level work makes sense.

Finally, bear in mind that we sometimes use low-level concepts in our high-level reasoning. In programming, sometimes all you want is a set of bits. So bits appear in our upper-level abstraction. But be warned: just because you are talking about bits in your high-level, it doesn't mean that they are being *represented* exactly as bits at the low level! Perhaps each of your high-level bits is being represented by *an entire* 32-bit machine word, for convenience of implementation. If you don't care about the space usage, then maybe that's not a big deal. Similarly, we still use sets in our higher-level reasoning too: sometimes you just need a set of kittens to make your day. But your high-level notion of "set of kittens" is going to be interpreted somehow in set theory, and that interpretation may not be readily recognizable after you compile it down. It will surely be a set, because sets are all there is at the bottom, but a set of what? That is to say, we can really think about everything we're doing in mathematics as operations on sets.²

Now, to keep the navel gazing to a reasonable minimum, we are not going to build *all* of the mathematical concepts we need from the ground up. In practice, we are going to choose a few primitive notions in addition to sets and work with them. In this document, however, we will do a bit of wallowing in the set-theoretic muck, partly so you have some experience looking under the hood at the low-level parts, partly so you can appreciate the high-level abstractions we use, and partly so that you can do your own trouble-shooting and discern whether your high-level reasoning makes sense. It's good to set some ground-rules for creating and manipulating sets. You may have seen some of these concepts before, but maybe not explained in this painful of detail. Enjoy!

3 Sets and Logic

Set theory, as we will use it in class, is really an embedding of notation and axioms that are meant to codify what sets we can describe and manipulate, into a logical language called *first-order logic*, which gives us the tools that we need to actually say stuff about sets and prove that the things we say "make sense". The two combine to create what we technically call a *first-order set theory*. So they are necessarily intertwined. Ultimately we want to get comfortable with being relatively informal about our discussions of sets, just like mathematicians tend to be, but we want to make sure that if we really *had* to, we could get super-pedantic and formal about it, literally using first-order logic all the way down. This is quite analogous to programming practice. Experienced programmers are capable of writing precise code that a computer can execute, but they may still communicate with one another using pseudocode, knowing full well that when push comes to shove they can sit down and implement it: pseudocode is sometimes a more efficient and insightful way to communicate the essence of a program than concrete code. The same is true in mathematics, except that not all mathematicians are equipped to fill in all the low-level details, but they somehow learn to stay within (or close to) the ball park, so that their proof sketches are enough guidance to a well-trained expert to fill in the dots (and correct any typos) Avigad [2020].

For this reason, I will intersperse the casual informal set notation with the super-pedantic statements within logic, so you can tell precisely what we mean when we write some informal stuff down: so you can fill in the dots when you need to.

4 Sets are about Stuff

At heart, set theory is a way of talking about containers of stuff. *Everything* in set theory boils down to the "element of" relation \in . If one set A is an element of another set B , then we write $A \in B$. That's really all there is to set theory, but we will build up a lot of stuff on top of that single concept (see, we don't even have 1's and 0's just \in). First, we need to know how this \in relationship works...what are the rules that

²One can argue about whether this is a good idea, just as one can argue whether the X86 instruction set is a good idea. Philosophers and logicians have proposed other foundations for mathematics, and with good reason! But from the early 20th century until now, set theory rules the day, much like the X86 ISA.

constrain it? These are captured by the *axioms* of a set theory. I'm not going to give *all* of the axioms of set theory, just the ones that we will use regularly. In fact, there are lots of set theories out there, but many have a pretty common core. We'll be using a set of axioms for what is called *Zermelo-Frankel Set Theory* or ZF for short. Technically we'll be using a variant called "Intuitionistic ZF" a.k.a. IZF. Now let's talk about some of the axioms, which tell us how sets work as well as what kinds of sets are out there, i.e. what sets *exist*.

5 When are two sets really the same set?

We start with a basic property of sets:

A set's identity is entirely determined by its elements.

Think of a set as an invisible bag of stuff...all we identify are the stuff, not the material that the bag is made of. This conception of sets is captured formally by the *Axiom of Extensionality*.³:

$$\forall S_1. \forall S_2. S_1 = S_2 \Leftrightarrow \forall S. S \in S_1 \Leftrightarrow S \in S_2.$$

In words, this reads: "for any sets S_1 and S_2 , they are the same set, i.e. identical, if and only if they contain the same elements." Remember: the only elements that arise in set theory are sets, so we're just worried about what sets are elements of some other set. For our purposes, this axiom defines equality of sets $S_1 = S_2$. First-order logics typically provide rules about equality that codify some "bare minimum" properties that must be true of any notion of equality. We'll discuss those later when we learn about how to reason in our logic. The remaining fundamentals of set equality/identity can be ascribed to this axiom.

Lets briefly discuss some of the logic notation that arises in the formal statement of the axiom. The symbol \forall means "for all", and \Leftrightarrow which means "if and only if". \Leftrightarrow is really an abbreviation:

$$A \Leftrightarrow B \equiv (A \Rightarrow B) \wedge (B \Rightarrow A).$$

Here the wedge \wedge is the symbol for "and". The arrow $A \Rightarrow B$ is just "if A then B" or "A implies B" or "A only if B" or "B if A": all of these statements are different ways of saying the exact same thing in words. To me the stranges ones are the last two, which take getting used to. But they help explain how "A if and only if B" textually breaks down into the two distinct implications "A only if B" and "A if B".

The \equiv symbol is used to introduce *notational definitions*. $A \equiv B$ can be read as "Whenever you see the syntax A , you can immediately rewrite it as B in any context." Think of it as analogous to a *macro*, like what you would write in the C Pre-processor. We use notational definitions to encode familiar concepts in terms of simpler ones (thereby keeping our logical language small and precise), often resulting in a more concise or traditional notation, but backed by a rigorous (possibly longer) logical presentation. For now don't worry too much about how precisely to read the above logical formula, we'll get into the details of how to understand logical propositions later.

But unpacking the axiom above a little more, we now get: "for any sets S_1 and S_2 , we know that: 1) if they are equal, then any element of S_1 is an element of S_2 and any element of S_2 is an element of S_1 ; and 2) if any element of S_1 is an element of S_2 and vice versa, then S_1 and S_2 are equal/identical, i.e. the same set."

Now we know how to judge that "two" sets are the same set, or more precisely that two set expressions *describe* the same set. Since equality/identity is a pretty key idea in logical systems, it's important that we lay this out good and early. Without a notion of equality/identity, we wouldn't be able to justify referring to a specific set, e.g., "the set that has no elements" as we will do next.

6 The Empty Set

So we now know how to determine if two set expressions describe the same set, *but we still don't know if there are any sets!*

³Roughly speaking, a set's *extension* is "the stuff that it is made up of."

Don't let the axiom of extensionality's "for all sets..." claim fool you: logic allows us to make statements like "All flying pigs prefer to drive Teslas." which might be a true statement, not because flying pigs don't like bicycles, but rather because we can establish that there are no flying pigs in the first place!⁴ This is one of the powers and challenges of logic: it empowers us to make deductions based on premises that might not be true, or might be true but we do not yet know.⁵ In short, we could still talk about "all sets" even if there were none, so we are on the hook to ensure that there's *at least one* set. Let's do that.

One set that's very important to us is the set with nothing in it: *the empty set*. This leads us to state the *Axiom of the Empty Set*:

There is an empty set, i.e. there exists a set with nothing in it.

Using formal notation, this statement looks like the following:

$$\exists S_1. \forall S_2. S_2 \in S_1 \Rightarrow \perp.$$

Here, the logical symbol \perp represents *falsehood*. In our logic, $A \Rightarrow \perp$ is our way of writing "not A ." You can read it as "Well if A then pigs fly!" as in, there is no way that A holds: its truth implies *the impossible* (proposition). If you have experience with logic, then this may look a bit odd compared to other presentations of negation, but there is some actual sense to this. In fact, we often abbreviate this kind of negation using another common notation:

$$\neg A \equiv A \Rightarrow \perp.$$

Many presentations of logic use the notation $\neg A$ for negation, but take it as primitive. We will find it very useful to treat "not A " as a non-primitive logical form, just like we treat "if and only if". Reasoning about negation reduces to reasoning about implication and reasoning about falsehood (yup, it's a thing!), which can make proofs that involve negation more mechanical and less mysterious, and can also clarify the distinction between reasoning about falsehood (which we must do) and proof by contradiction (which we will never have need for). These two distinct practices are often confused for one another unfortunately Bauer [2010].

Also, we often abbreviate the idea that some relation doesn't hold by crossing it out. For example,

$$A \notin B \equiv A \in B \Rightarrow \perp$$

and

$$X \neq Y \equiv X = Y \Rightarrow \perp.$$

Keep these notational definitions in mind: they will help you prove theorems about non-containment and inequality later.

Now, we can combine the axiom of the empty set with the axiom of extensionality to prove that the empty set is unique, i.e., there is at least one and at most one set that is empty.

Proposition 1 (The Empty Set is Unique).

$$(\exists S_1. \forall S_2. S_2 \in S_1 \Rightarrow \perp) \wedge \left(\forall S_1, S_2. (\forall S. S \in S_1 \Rightarrow \perp) \wedge (\forall S. S \in S_2 \Rightarrow \perp) \Rightarrow S_1 = S_2 \right).$$

Reading the above formula literally, it reads "there exists a set that has nothing in it; and furthermore should you find yourself with two sets that have nothing in them, then they are actually the same set." In other words, there's at least one empty set and at most one empty set, i.e. there's exactly one! This is why we get to call it *the empty set*, rather than *an empty set*, by the way. Thus the concept of emptiness suffices to *definitely describe* a particular set. *Definite description*—the presentation of sufficient conditions to identify a particular set—will be one of the central activities in this course, and much of the machinery we study will provide more and more sophisticated means to definitely describe certain sets. One motto of the course

⁴Barring recent advances in biotechnology of which I am unaware.

⁵For example, many results in algorithmic complexity theory are premised on assumptions like $P \neq NP$, even though we don't know, and may never know, if that is true.

is: “the word *define* is short for *definitely describe* and *definition* is short for *definite description*.” We introduce a formal notation for definite descriptions shortly.

Let’s use notational definitions to clarify the formal statement of the Axiom of The Empty Set: the textual statement is simple but the primitive formal statement is a hot mess! We can simplify the logical formula using a couple of abbreviations. First, as described above we use \notin to mean “not an element of” Second, we codify the concept of *unique existence*, a.k.a. *uniqueness*, using a notational definition. Let Φ be some property of sets. Then we define the $\exists!$ quantifier to express uniqueness:

$$\exists!S.\Phi(S) \equiv \exists S.\Phi(S) \wedge \left(\forall S_1, S_2. \Phi(S_1) \wedge \Phi(S_2) \Rightarrow S_1 = S_2 \right).$$

In the above notation, $\Phi(S)$ means that the logical formula $\Phi(S)$ may contain unquantified references to the set variable S . Then $\Phi(S_1)$ denotes the same formula as $\Phi(S)$ except replacing some unquantified references to S with S_1 . This notation is often intended to replace *all* references to S in $\Phi(S)$, but in some contexts that is not necessary. Since total replacement is the default, we’ll point out those cases where only some—including none!—references might be replaced.

This notational definition emphasizes the two distinct reasoning powers that uniqueness gives you: first that something with the property $\Phi(S)$ is out there (in case you need it), and second that any two set expressions that satisfy this property must actually describe the same set, so you can use logical equality rules to reason about that sameness. Reasoning with equality is an important and powerful tool.

Using these two notational definitions, we can succinctly rewrite Prop. 1 as:

$$\exists!S_1.\forall S_2.S_2 \notin S_1$$

In short: “There’s a unique set that has no elements.” This says *exactly the same thing* as the original proposition, since this “macro-expands” to it. Bear in mind that we have not proven this proposition yet, but we’ll see that it follows from the axiom of extensionality and the axiom of the empty set.

Now that we know that there is only one empty set, we can use a definite description to name it.

$$\iota S_1.\forall S_2.S_2 \notin S_1$$

This form of expression, $\iota S.\Phi(S)$, where Φ is some proposition that (probably) refers to S is what we call a *definite description* Stenlund [1975]. This particular expression can be read “the unique set S_1 such that no set S_2 is an element of it.”

A definite description is a way of...well...describing a set by stating some property $\Phi(S)$ of a set S that uniquely characterizes the set. As such, a definite description is a way of “naming” a set, kind of like “that guy who sits in the back of class every day snapping his gum.” Ideally such a description suffices to uniquely identify the set in question, in which case the description is called *proper*. If there are three guys snapping gum, then this description is *indefinite*, and if no guys snap gum then the description is *empty*. We should strive for proper definite descriptions, and avoid the latter two.

One nice thing about definite descriptions is that they not only identify a set, but do so by enunciating a property that the set in question (uniquely) satisfies. One built-in rule of our logic is that for any proper definite description $S \equiv \iota S_1.\Phi(S_1)$, we can immediately deduce that $\Phi(S)$ holds, i.e. $\Phi(\iota S_1.\Phi(S_1))$ is true. As you would expect, the unique set S that satisfies Φ ...satisfies Φ ! But don’t forget that sooner or later we must prove that Φ describes a unique set to establish that the definite description is proper.

For conciseness and tradition, we introduce a common name for the empty set using a notational definition:

$$\emptyset \equiv \iota S_1.\forall S_2.S_2 \notin S_1$$

It’s no coincidence that the definite description of the empty set looks *exactly like* the statement of uniqueness, but replaces $\exists!$ with ι . We’ve transitioned from establishing that “there exists a unique set S such that ...” to referring directly to “the unique set S such that...”.

Definite descriptions are not strictly necessary for logic: we can “compile” propositions that use definite descriptions to equivalent propositions that don’t. But definite descriptions more directly and precisely capture how we casually talk about mathematical objects. We are prone to make statements that directly refer to “the empty set”, and reason just based on that name. Definite descriptions, and the reasoning

principles they embody, mirror our informal speech. In addition to improving the mapping from prose to formality, definite descriptions clarify the relationship between definitions and reasoning principles: the definitions directly imply some principles, which we can immediately read from the structure of the description embedded in their formal name. For example, the fundamental reasoning principle of the empty set can be deduced directly from its definite description:

$$\forall S. S \notin \emptyset.$$

We'll find that the concepts underlying definite descriptions, and the reasoning principles that they entail, scale from our lowest-levels of set theory to our high-level strategies for defining components of a programming language semantics. One could also imagine teaching a computer to recognize and manipulate them on our behalf since they are so systematic.

6.1 Uniqueness, Shmuniqueness!

You may be annoyed that we are still on the hook to prove that the empty set is unique. Why don't our axioms just make the uniqueness of the empty set part of the axiom of the empty set and be done with it? We could have done that, but we didn't. Our axiom of the empty set does not establish uniqueness because it does not need to: the axiom of extensionality could be used to readily establish uniqueness, so why add more complexity and redundancy to one of your axioms?

Furthermore, by relegating all reasoning about set equality/identity to the axiom of extensionality, we are left with a clear and unambiguous understanding of what constitutes set equality. If we assert uniqueness in some other axiom, then expanding the notational definition for $\exists!$ shows that such an axiom also makes claims about equality, and might accidentally introduce unintended additional properties to set equality.

These principles reflect how mathematicians tend to think in my experience: keep the *number* of axioms you assert small, keep the *strength* of each axiom as weak as possible, and try to make each axiom as *independent* from one another as possible, such that each is necessary and none implies the others. Building up a foundational theory this way keeps axioms manageable and ensures that you can get your work done with as small of a *trusted base* as possible. Mathematicians have argued for yeeeeears about various axioms: which ones are self-evident, which ones are questionable, which ones are necessary for a theory to be useful, and which ones imply really strange things but we can't seem to do without them. This is not dissimilar to an aesthetic that drives some programming language designs: the desire to provide as few mechanisms as possible (e.g. "everything's an object" or "everything's a function") and to provide only "one way to do it" whatever "it" may be.

Now this approach to choosing axioms (and constructing programming languages) *does not* make your working set of properties particularly ergonomic, but if the language provides some mechanism for building your own abstractions, then you can build up higher-level human-scale tools (i.e. propositions) on top of them with confidence and then use those in your day-to-day work. By having a firm, clear foundation, you need not worry that your mathematical castles rest on sand, or that unnecessary inconsistencies may arise accidentally. But it goes to show that these systems were built to reason *about* set theory, not to do real work in it!

7 Beggars Can Be Choosers

Some axioms of our set theory allow us to identify bigger sets in terms of smaller sets. However, much of the time we desire to identify *smaller* sets by selecting desirable elements from some previously-determined bigger set. The tool that we use for this selection process is called the *Axiom Schema of Specification* (also known as the *Axiom Schema of Separation*)

Let $\Phi(X)$ be some logical predicate on sets X and let A be some set. Then there is a set S of all and only the elements $B \in A$ such that $\Phi(B)$ holds.

i.e. for every predicate $\Phi(X)$ we have the axiom:

$$\forall A. \exists S. \forall B. B \in S \Leftrightarrow B \in A \wedge \Phi(B).$$

This one is usually called an axiom *schema* because it is taken to stand for an infinite number of axioms: one for every Φ that you can come up with.

Notice how this axiom schema is phrased as an if-and-only if, so we can pick out a set whose contents *exactly* match the result of filtering, unlike earlier axioms like the axiom of pairing. This axiom is the one that we can use to winnow things down to exact sets.

We can prove that for any A and Φ that specification describes a unique set, at which point we are justified in associating a traditional variation of *set-builder notation* for it:

$$\{ B \in A \mid \Phi(B) \} \equiv \lambda S. \forall B. B \in S \Leftrightarrow B \in A \wedge \Phi(B).$$

This particular notation is often called a *set comprehension*.⁶ The notation $\{ B \in A \mid \Phi(B) \}$ describes the unique set characterized by being all and only the B 's in A for which $\Phi(B)$ holds. Its underlying definite description immediately provide the following reasoning principle:

$$C \in \{ B \in A \mid \Phi(B) \} \Leftrightarrow C \in A \wedge \Phi(C).$$

Anti-Patterns in Notation Many texts that introduce set theory introduce a quite similar notation to the above: similar but different in a fraught way:

$$\{ B \mid \Phi(B) \}$$

We will steer clear of this notation! The main reason for doing so is that every set description notation that we will introduce here can be justified by a unique existence proof based on the axioms of set theory. This one cannot. It corresponds to a proposition that was once proposed as an axiom of set theory but turns out to destroy the theory, called *unrestricted comprehension*:

$$\exists S. \forall B. B \in S \Leftrightarrow \Phi(B).$$

This proposition essentially says “every proposition determines a set.” Unfortunately, we can use it to prove \perp —falsehood!—without any extra premises, which renders the theory useless. A logic that can prove falsehood is called *inconsistent*, and in most situations this is a death sentence. In this particular case, The logician Bertrand Russell (among others, even before him) demonstrated how to prove \perp , a technique now named *Russell's paradox*. It involves setting $\Phi(B) \equiv B \notin B$, yielding the set S of all sets B which are not elements of themselves. Existence of this set can be used to prove \perp .

With careful choice of Φ , the unrestricted notation can be used to properly define sets, but a more rigorously motivated notation can obviate the need for such care. It's far less error-prone to appeal to robust description mechanisms that are aligned with the axioms of the theory. As it happens, definitions guided by the axiom schema of separation do not suffice for all uses of set theory, but this is because the axiom schema of separation alone does not suffice. IZF includes another related axiom that fills the gap, and we can introduce another definition form to go with it.

To quote Yaron Minsky, a proponent of statically-typed functional programming:⁷ “Make illegal states unrepresentable.”

8 Sets in Sets in Sets

Okay: so far we know what it means for two sets to be equal, we know that there's an empty set, and we know how to use predicates to curate an artisanal subset of bespoke sundries. Is that enough for us to introduce more friends for the empty set? Not so far! Why not?!? Well, all we have right now are the empty set and set comprehensions, and applying the latter to filter the former gives us back the empty set again. We're still stuck with just one set.

Our next axiom gets us some more sets, in a somewhat simple way. The *Axiom of Pairing*:

For any sets A and B , there is a set C that contains exactly A and B as elements.

⁶Set comprehensions inspired a language feature called *list comprehensions* that appears in Python, Haskell, and other languages.

⁷<https://blog.janestreet.com/effective-ml-revisited/>

Formally,

$$\forall A. \forall B. \exists C. \forall D. D \in C \Leftrightarrow A \vee B = C.$$

The \vee symbol represents *inclusive “or”*: a formula $A \vee B$ is true if and only if A is true, B is true, or *both A and B are true*. This use of inclusive-or, common in logic, can be confusing since in informal speech we often say “or” when we mean “either A or B but not both”.

This axiom is just a bit obnoxious because it tells us that there’s a set with A and B as members, but it does not outright say that the set is unique. There may be tons of other sets with this property! However, just like we dealt with the axiom of the empty set, we can use the axiom of extensionality to prove that any two sets with this property are identical. Thus, this set not only exists, but is unique. Uniqueness justifies our introduction of some possibly familiar notation:

$$\{A, B\} \equiv \iota S. \forall Q. Q \in S \Leftrightarrow Q = A \vee Q = B.$$

8.1 From Two Elements to One Element to Many Elements

Now that we have the empty set and justification to describe sets $\{A, B\}$ given any pre-existing sets A and B , we can apply the axiom of pairing, with \emptyset for A and \emptyset again for B to prove that there exists a set S such that $Q \in S$ if and only if $Q = \emptyset \dots$ or $Q = \emptyset$? If that sounds redundant to you, that’s because it is. Formally, we can prove that $Q = \emptyset \vee Q = \emptyset \Leftrightarrow Q = \emptyset$. So it’s safe to say that the set $\{\emptyset, \emptyset\}$ has only one element: $Q \in \{\emptyset, \emptyset\} \Leftrightarrow Q = \emptyset$. This kind of reasoning can be applied schematically to *any* pre-existing set A , and it’s a little odd to use the notation $\{A, A\}$ to describe such *singleton* sets, so we introduce better suggestive notation:

$$\{A\} \equiv \iota S. \forall Q. Q \in S \Leftrightarrow Q = A.$$

Our reasoning above implies that $\{A, A\} = \{A\}$: both expressions describe the same set!

At this point we we can construct a set with *zero* elements, sets with *one* element, sets with *two* elements. What about three, four, five, and the rest? Not yet: we need more axioms! However, we can repeat our singleton and coupleton constructions to form deeply nested sets of couples of singletons of couples, like $\{\{\emptyset\}\}$, $\{\{\{\emptyset\}\}\}$, \emptyset and so on.

9 Come Together, Right Now, Over Me

Another useful and well-known approach to construct a set that combines other sets is to take the *union* of a collection of sets, i.e., a set of sets. The *Axiom of Union* supports this:

For every set S there is a set G that contains as elements the contents of every set that S contains.

i.e.,

$$\forall S. \exists G. \forall A. A \in G \Leftrightarrow \exists B. (A \in B \wedge B \in S).$$

To make sense of the formal notation, let’s consider an example. Let $S = \{\{1, 2\}, \{3, 4\}\}$. Then from the axiom, we know that there is some set G such that each of the given numbers is an element of G , i.e., $1 \in G$, $2 \in G$, and so on.

Using the axiom of extensionality, we can prove that sets justified by the axiom of union are unique. Uniqueness ushers in the traditional notation for set union:

$$\bigcup S \equiv \iota G. \forall A. A \in G \Leftrightarrow \exists B. A \in B \wedge B \in S$$

We can use general union and pairing to define the common notation for binary union:

$$A \cup B \equiv \bigcup \{A, B\}$$

Why define the binary union in terms of the big one, and not the other way around? Well, because we can’t form the union of an infinite number of sets by doing it two at a time: we’ll get bored eventually. However,

if we have a set with an infinite number of things (see below!), then we can form an infinite union whiz bang!

As usual, the definite description immediately induces a reasoning principle for sets formed using union.

$$\forall S. \forall A. A \in \bigcup S \Leftrightarrow \exists G. A \in G \wedge G \in S.$$

Thus we've turned the axiom of union into a reasoning principle about *the* union of a family of sets.

10 Extensional Set Definitions

With the introduction of the Axiom of Union, we finally have enough principles to justify the construction of sets with an arbitrary *finite* number of elements, not just zero, one, and two! The name of the axiom of extensionality leads us to the idea of an *extensional definition* of a set. In informal practice with sets, when you want to describe a set with a finite number of elements, you can just enumerate list of elements as a way to refer to the set. For example, any recent text that introduces some set theory or sets will introduce *set builder* notation, in particular expressions like:

$$\{ \text{cheese, crackers, vegemite} \}$$

to describe the *unique* set that contains exactly three food products: cheese, crackers, and vegemite. The notation above, in which we describe a set by enumerating a list of descriptions of elements, is so common that if you've seen it before, you were probably told to consider it a primitive syntactic construct. We are not going to take it as primitive. Instead we'll introduce it in two steps. First, we introduce the notation using a notational definition:

$$\{ A, B, \dots, C \} \equiv \iota S_1. \forall S_2. S_2 \in S_1 \Leftrightarrow S_2 = A \vee S_2 = B \vee \dots \vee S_2 = C.$$

This means that we expand the above expression describing food items into the complex logical expression:

$$\iota S_1. \forall S_2. S_2 \in S_1 \Leftrightarrow S_2 = \text{cheese} \vee S_2 = \text{crackers} \vee S_2 = \text{vegemite}.$$

This particular description can be read as "the (unique) set S_1 that can be described by the property that any set S_2 is an element of S_1 if and only if S_2 is identical to the set described by cheese, the set described by crackers, or the set described by vegemite."

Regarding our new notation for extensional set definitions, and the underlying definite descriptions that (notationally) define them, we're still on the hook to prove that such definitions are proper. To do so, we must prove a theorem of the form:

Proposition 2 (Extensional Descriptions are Unique). *Let S_1, S_2, \dots, S_n be sets. Then there is a unique set Q such that for all sets R , $R \in Q$ if and only if $R = S_1$, $R = S_2$, \dots , or $R = S_n$.*⁸

Once we've proven that extensional set definitions in general are proper descriptions, then we can justifiably apply the defining property. Here is the property induced immediately by our description of the set of food products.

$$\forall S. S \in \{ \text{cheese, crackers, vegemite} \} \Leftrightarrow S = \text{cheese} \vee S = \text{crackers} \vee S = \text{vegemite}.$$

Typical presentations of sets take this kind of element-hood property for granted. For us, this property is derived from expanding the set notation into a definite description that is written in terms of logical operations and elementhood. The longer form makes evident that we are working with a description of a set, not the set itself.

⁸Formally $\exists! Q. \forall R. R \in Q \Leftrightarrow R = S_1 \vee R = S_2 \vee \dots \vee R = S_n$. We'll explain the notation $\exists!$ soon.

10.1 Extensional set definitions from the ground up

So now we have enough machinery to justify extensional set definitions. Given sets a, b , and c , do how we know that the set $\{a, b, c\}$ exists? In pedantic form the statement is:

$$\exists A. \forall S. S \in A \Leftrightarrow S = a \vee S = b \vee S = c,$$

One way to justify its existence is as follows:

1. use the axiom of pairing to form $\{a, b\}$;
2. use the axiom of pairing to form $\{c\}$;
3. use the axiom of pairing to form $\{\{a, b\}, \{c\}\}$;
4. use the axiom of union to form $\{a, b, c\}$.

Well that was exhausting! But we could further extend this set by constructing $\{d\}$, pairing it with our present set, and applying union again. Singleton, pair, union; singleton, pairing, union: wash, rinse, and repeat as needed.

We will never bother performing this construction again. But note, that we have a reasoning principle for our set based simply on its description:

$$\forall S. S \in \{a, b, c\} \Leftrightarrow S = a \vee S = b \vee S = c,$$

Recall that such reasoning also makes sense for a set name like $\{a, a\}$, though it's mighty redundant: we like the name $\{a\}$ better. But this observation is a cause for care particularly when making an abstract argument. For instance, if we assume some sets a and b , and then mention "the set $\{a, b\}$ ", we can't be sure that the set $\{a, b\}$ has two elements: what if $a = b$? On the other hand, if we know a priori that $a \neq b$, then the set $\{a, b\}$ definitely has two elements. Notice, though, that in both cases it's true that $a \in \{a, b\}$ and $b \in \{a, b\}$.

10.2 Tool Building: A General Principle of Cases

The definitional property of extensional sets is quite useful, but would likely lead to proofs that are wordier than necessary. Instead of settling for this property, we can instead construct a quite similar property that takes a little work to establish, but will streamline later proofs. In particular, we can use the definite description of extensional sets to state and prove the following principle.

Proposition 3 (Principle of Cases on $\{S_1, S_2, \dots, S_n\}$). *Let Φ be a property of sets and S_1, S_2, \dots, S_n be sets. Then $\Phi(S)$ holds for all $S \in \{S_1, S_2, \dots, S_n\}$ if $\Phi(S_1), \Phi(S_2), \dots$, and $\Phi(S_n)$.*⁹

We will be equipped to prove this proposition once we study the proof techniques that apply to our logical operations. This property is pretty obvious intuitively, so one may be surprised how much work it takes to prove this rigorously. But think about it: how many easy-to-understand procedures do you know of where what you want seems straightforward, and it's obvious that it can be implemented, but the finicky details you had to go through to get it done right inspire you to search the internet for the code or (gasp!) use the standard library implementation? Our proof of this is like that. For now let's take it as given.

This is our first example of *using our tools to build better tools*, much like one can use the primitives of a programming language to design powerful and productive libraries of routines.

Right from the start, we see an incident of a formal definition (extensional sets) affecting our reasoning principles: an extensional definition of a set licenses you to reason about its elements by cases on its (finite) members, and lets us build a general principle for proving properties of a sets elements. When we get to defining sets that contain an infinitude of elements, we'll need new tools. Luckily our definitions of those sets will help with that.

⁹Formally, $\Phi(S_1) \wedge \Phi(S_2) \wedge \dots \wedge \Phi(S_n) \Rightarrow \forall S \in \{S_1, S_2, \dots, S_n\}. \Phi(S)$. One seeming mismatch of textual vs. formal statements of propositions is that in text the "for all" statement often comes *after* the proposition to which it applies, I suspect to draw the reader's attention to the heart of the statement rather than the quantifiers.

[RG: replace the below with a fresh example of an equality of set expressions]

A point that is worth re-emphasizing: our set-builder notation is just a *description* of a set, sometimes called a *name* of a set, not the set itself. For instance, consider the set expressions

$$\{ \text{cheese} \} = \lambda S. \forall S_1. S_1 \in S \Leftrightarrow S_1 = \text{cheese}$$

and

$$\{ \text{cheese}, \text{cheese} \} = \lambda S. \forall S_1. S_1 \in S \Leftrightarrow S_1 = \text{cheese} \vee S_1 = \text{cheese}.$$

Despite the repeated appearance of cheese in the expression, each describes exactly the same set: the one containing cheese and nothing else. Do not assume that the number of elements in a finite set is always one more than the number of commas in a set expression! Also, notice that despite describing the same set, each description leads to a subtly different reasoning principle. Later in the course we will discover that having multiple different definitions of the same set can give us *usefully* different reasoning principles, derived from different definite descriptions of the same set.

Exercise 1. Write down the reasoning principles for the two cheesy set descriptions. How do they differ? Which one do you think we'll prefer and why?

11 Infinite stuff!

Given the tools we have so far, we can only build finite sets! That's because all of our logical arguments have to have a finite number of steps, which can really be a pain sometimes, but that's how logicians roll. This is also how computer programmers roll: ever seen an infinite-length program?

So how do we get infinite stuff? Well, like we did with the empty set, we *assume it*, via the *Axiom of Infinity*:

There's a set that (1) contains the empty set and (2) contains the set $a \cup \{ a \}$ whenever it contains a .

i.e.,

$$\exists S. (\emptyset \in S) \wedge (\forall A. A \in S \Rightarrow A \cup \{ A \} \in S).$$

Notice how I used $A \cup \{ A \}$ in the above formula, rather than spelling the whole thing out in terms of \in , \exists , uniqueness, etc. That would be way too painful! Luckily we can already exploit our paucity of abstractions to make life more pleasant. You may as an exercise try to boil this all the way down to just statements about elementhood.

Set theorists like to call the set that has *only* the elements above ω , and they use it to represent the set of natural numbers:¹⁰

$$\begin{aligned} 0 &\equiv \emptyset \\ 1 &\equiv \{ 0 \} \cup 0 = \{ \emptyset \} \\ 2 &\equiv \{ 1 \} \cup 1 = \{ \{ \emptyset \}, \emptyset \} \\ &\vdots \end{aligned}$$

First, notice the leaky abstraction: when was the last time you took the union of a set with a number??? Let's not do that any more! instead, we will just consider the nice abstract set of *natural numbers* as the set,

$$\mathbb{N} = \{ 0, 1, 2, \dots \},$$

and say that questions like $1 \in 2$ are off limits...that would be exploiting the *representation* of numbers as sets, rather than the properties of numbers *as an abstraction* that we care about. Notice that 0 is a natural number. All right-thinking computer scientists will use this definition.¹¹

¹⁰This encoding is due to the great logician and computer scientist John von Neumann.

¹¹I once went to a wedding of two computer scientists, and during dinner they sat at "Table 0" because, as the table manifest said: "computer scientists start counting from zero."

12 All the Subsets

Now we add yet another axiom that lets us build bigger sets: the *Axiom of Power Set*:

Given any set A , there exists a set that contains all subsets of A .

i.e.,

$$\forall A. \exists S. \forall B. B \in S \Leftrightarrow (\forall C. C \in B \Rightarrow C \in A).$$

Good grapes, just looking at that makes my eyes bleed! In the above formalism, B is a subset of A because every element C of B is also an element of A . Because we use subsets so often, there's a convenient notation for them:

$$X \subseteq Y \equiv \forall C. C \in X \Rightarrow C \in Y.$$

With this we could have written:

$$\forall A. \exists S. B \in S \Leftrightarrow \forall B. B \subseteq A.$$

Once again, this axiom doesn't say that the set with this property is unique, but we can whack it with the axiom of extensionality to establish uniqueness. We call this set the *powerset* of A , and give it a notation:

$$\mathcal{P}(A) \equiv \iota S. B \in S \Leftrightarrow \forall B. B \subseteq A.$$

Some books use the notation 2^A , which sort of explains the name "powerset" (2 to the power of A): the reason for that notation is that if some set A has n elements, then $\mathcal{P}(A)$ has 2^n elements.

The uniqueness and description of the powerset give us another reasoning principle:

$$X \in \mathcal{P}(A) \Leftrightarrow X \subseteq A.$$

13 Bait and Switch

With the axioms and principles introduced so far, we can get very far, using powerset and union to build up sets, and specification to hone them down. But sometimes, surprisingly rarely, we need to describe a set by relating it directly to the properties of another set. It was only after some years of working with Zermelo's axioms that logicians realized that the axioms outlined above were not sufficient for this task. As a result, Abraham Fraenkel introduced a new defining principle, which Thoralf Skolem refined into a new formal axiom schema. This contribution is the reason that our set theory is referred to as *Zermelo-Fraenkel*.¹² It is called the *Axiom Schema of Replacement*:

If $\Phi(A, B)$ is a proposition, S is some set, and for every $A \in S$, there is a unique set B such that $\Phi(A, B)$. Then there is a set Q whose elements are exactly those B such that $\Phi(A, B)$ for some $A \in S$. i.e.,

$$\forall S. (\forall A \in S. \exists! B. \Phi(A, B)) \Rightarrow \exists Q. \forall B. B \in Q \Leftrightarrow \exists A \in S. \Phi(A, B).$$

Let's unpack this a bit. The set S is one that we constructed previously...somehow. The goal is to describe a new set Q by relating each element A of S to exactly one element B of Q , allowing for the possibility that multiple elements of S may correspond to the same element of Q . The proposition $\Phi(A, B)$ captures the relationship between elements, and the axiom allows any proposition Φ that yields the appropriate kind of relationship.

With a little more reasoning, incorporating the axiom of extensionality, we can prove that S and Φ induce a *unique* Q , which gives us license to introduce a new notation and corresponding definite description.

$$\{ B \leftarrow \Phi(A, B) \text{ for } A \in S \} \equiv \iota Q. \forall B. B \in Q \Leftrightarrow \exists A \in S. \Phi(A, B).$$

An example from arithmetic might be the set of all multiples of 7:

$$\{ m \leftarrow m = 7 * n \text{ for } n \in \mathbb{Z} \}$$

¹²Poor Skolem!

Here, S is the set of integers, \mathbb{Z} , and $\Phi(n, m)$ is $m = 7 * n$. This particular case, where the proposition $\Phi(A, B)$ says that B is equal to some expression involving A , is sufficiently common that we introduce a more compact and direct notation for it:

$$\{ F(A) \text{ for } A \in S \} \equiv \{ B \leftarrow B = F(A) \text{ for } A \in S \},$$

So the multiples of 7 can be written

$$\{ 7 * n \text{ for } n \in \mathbb{Z} \}.$$

The general replacement notation is then reserved for those descriptions that benefit from its full expressive power. In contrast to our previous set description notations, this one is only proper if the prerequisite $\forall A \in S. \exists! B. \Phi(A, B)$ holds.

Since replacement notation requires some pre-existing set S and some “function-like” proposition Φ , it does not allow us to express Russell’s paradox.

14 Sets May Not Eat Themselves!

Based on our reasoning principles, we may find ourselves constructing some set Q , and then subsequently constructing a set R such that $Q \in R$. Can we ever construct such a set like this where $R \in Q$ also holds? Nothing we have said so far precludes this possibility. However, our set theory does have an axiom, called the *Axiom of \in -Induction* whose express purpose is to forbid sets from containing themselves.¹³ So ultimately Q and R cannot be the same set. Since we do not need to explicitly use the axiom of foundation, we discuss it no further in this class. But recognizing that we have to explicitly disallow cyclic sets if we do not want them can give some more insight about logic.¹⁴

14.1 Coda

This concludes the axioms of set theory that we will focus on. IZF has more axioms, but they are mostly geared toward details we will barely touch on, if at all, so we can ignore them.

You may want to mess around with the axioms that we’ve covered so far to try and build various sets and see what a pain it is to hack in machine code.

Note that each mechanism that we have for defining/naming sets comes with its own reasoning principle. This is a recurring theme in mathematics and in this course in particular: *the structure of my definitions determines the structure of my reasoning*. Often we can describe reasoning principles directly in terms of some definition for picking out a set of interest.

Next, we’ll consider some of the pragmatics of using set theory.

15 The Practice of Working with Sets

In this section, I’ll describe how we will actually informally work with sets. Some of this rehashes concepts that we covered above, but ideally this version is written in a more high-level, actionable way, corresponding to how we will use set theory in real practice.

15.1 Assumed Sets

[RG: introduce the cofinite property explicitly as a means of extracting fresh atoms]

From time to time we will simply *assume* the existence of a set.

It will look like:

Suppose $a \in \text{ATOM}...$

¹³In Zermelo’s original set theory, this property is introduced using the Axiom of Foundation. The Axiom of \in -Induction yields the same effect without compromising constructivity [Crosilla, 2020]. In a non-constructive set theory, the two are equivalent.

¹⁴This design is a decision! There exist set theories that replace the axiom of foundation with an *axiom of anti-foundation* for the express purpose of allowing sets to contain themselves. Wild!

What we mean by this is that there must be some set with an infinite number of things in it that we can tell apart, and distinguish from the elements of other sets. So I can always pull out some a_i and then grab some other $a_j \neq a_i$. When we want to talk about specific members, we will usually use the same letter as the metavariable but typeset it in blue, as in a_1 . In the case of program variables, like $x \in \text{VAR}$, we may be more creative with our choice of elements.

We can justify this because we already know that there is a set with an infinite number of things, and one of the nice/crazy things about such a set is that you can take an infinite number of things out of that set and still have *another* infinite amount of things left over. For example, if I take the odd numbers out of the set of natural numbers, how many numbers are left over? An infinite amount: the even numbers! My head explodes when I think about this kind of thing.¹⁵

16 Set formations

Aside from naming a set by explicitly saying which elements are in it, we assume a number of ways of building sets. We're not trying to be super-primitive, or make sure that there's only one way to build a set. Rather, we just want to establish a set of primitive rules that we can always use for building new sets.

16.1 Union

We can always combine sets to form new sets, forming the union of two or any number of sets. If A is a set, and B is a set, then there is a set that we call $A \cup B$ with the properties that $c \in A \cup B$ if and only if $c \in A$ or $c \in B$.

Example 1. If $A = \{\text{meat, cheese}\}$ and $B = \{\text{vegetables, fruits}\}$ then $A \cup B = \{\text{meat, cheese, vegetables, fruits}\}$.

We can extend this to collections of sets. If F is a set of sets, then we can form the set $\bigcup F$, which is the union of all the sets in F : $c \in \bigcup F$ if and only if $c \in C$ for *some* set $C \in F$.

Example 2. If $F = \{\{\text{cow, bull, buffalo}\}, \{\text{sheep, elpaca, horse}\}, \{\text{frog, toad, turtle}\}\}$ then $\bigcup F = \{\text{cow, bull, buffalo, sheep, elpaca, horse, frog, toad, turtle}\}$.

Notice that $\bigcup \emptyset = \emptyset$, since we are taking the union of no sets, and $A \cup B = \bigcup \{A, B\}$.

16.2 Intersection

In addition to forming a set by grabbing element that appear in *any* of a given set of sets, we can also consider the set of elements that appear in *every* set in a given set. That's the intersection. If A and B are sets, then there is a set we call $A \cap B$ with the property that $c \in A \cap B$ if and only if $c \in A$ and $c \in B$.

Example 3. If $A = \{\text{paper, cut}\}$ and $B = \{\text{news, paper}\}$ then $A \cap B = \{\text{paper}\}$.

Generalizing this idea, if F is a *non-empty* set of sets, then we can form the set $\bigcap F$, which is the intersection of all the sets in F : $c \in \bigcap F$ if and only if $C \in F$, and $c \in C$ for *every* set $C \in F$. Analogous to unions, $A \cap B = \bigcap \{A, B\}$, but in stark contrast, there is no such set $\bigcap \emptyset$! This means that when you are reasoning about sets, you must establish that F is not empty before you take its intersection. In the case of both intersection and union, you must establish that F contains only sets before you take their union, otherwise it's an "abstraction error".¹⁶

¹⁵In the 19th century, the mathematician Richard Dedekind proposed roughly the above property as the defining characteristic of an infinite set.

¹⁶By "abstraction error" I mean that technically it's fine to take the union, because at the lowest level all we have are sets, but you would surely be violating an abstraction principle. For example, $\bigcup 1$ is nonsensical for numbers, but not for their von Neumann representation where $\bigcup 1 \equiv \bigcup \{\emptyset\} = \emptyset \equiv 0$. Yuk!

16.3 Powerset

Given some set C , you can form its *powerset*, which we call $\mathcal{P}(C)$, which is the set of all subsets of C : $A \in \mathcal{P}(C)$ if and only if $A \subseteq C$.

Example 4. $\mathcal{P}(\{\text{tic, tac, toe}\}) = \{\emptyset, \{\text{tic}\}, \{\text{tac}\}, \{\text{toe}\}, \{\text{tic, tac}\}, \{\text{tic, toe}\}, \{\text{tac, toe}\}, \{\text{tic, tac, toe}\}\}$.

Notice that the C has 3 elements, while $\mathcal{P}(C)$ has $2^3 = 8$ elements. Some texts will write 2^C as a different notation for $\mathcal{P}(C)$.

As a side note, we can characterize subset completely in terms of containment: If A and C are sets, then $A \subseteq C$ if and only if whenever $a \in A$, it follows that $a \in C$. We differentiate between subset \subseteq and *strict subset* \subset , meaning that if $A \subset C$, then it cannot be the case that $A = C$.

16.4 Sequences and Products

In addition to sets, unordered collections of unique elements, we will often want to consider *sequences* of elements that have a specific order and are not necessarily unique. Some examples of this are *pairs* like $\langle 1, 2 \rangle$ and $\langle 1, 1 \rangle$, *triples* like $\langle 4, 5, 3 \rangle$, and more generally *tuples*, meaning some ordered sequence of n elements. Note that for our purpose, sequences are *not* sets, so we don't talk about $1 \in \langle 1, 2 \rangle$: that's an "abstraction error."¹⁷

Now that we have a notion of sequences/tuples, we can form sets of them. If A is a set, and B is a set, then we can refer to the *product* of A and B , which we name $A \times B$, and it has the property that $c \in A \times B$ if and only if $c = \langle a, b \rangle$ for some $a \in A$ and $b \in B$.

Example 5. If $A = \{\text{eat, throw}\}$ and $B = \{\text{chicken, darts}\}$ then $A \times B = \{\langle \text{eat, chicken} \rangle, \langle \text{eat, darts} \rangle, \langle \text{throw, chicken} \rangle, \langle \text{throw, darts} \rangle\}$.

There are several generalizations of products. You can write $A \times B \times C$ for the set of triples from these three sets: if you *really* want to write the product of $A \times B$ with C , you should parenthesize as $(A \times B) \times C$. This is just a convention since we are far more likely to want triples than pairs that include pairs in the first position.

Also, in analogy to arithmetic, the set $A^2 = A \times A$ and $A^3 = A \times A \times A$ and so on. Note that $A^0 = \{\langle \rangle\}$, the set containing only the empty sequence (since it's the only "0-long" tuple).

Given a set A , we use A^* to refer to the set of all finite sequences of elements of A . We sometimes use the name ε to denote the empty sequence $\langle \rangle$.

$$A^* = \bigcup_{n \in \mathbb{N}} A^n = A^0 \cup A^1 \cup A^2 \cup \dots$$

Technically, I haven't told you how to construct this set (we need one more axiom of set theory (Replacement) to do it, but we're not going to use it often, so I'm eliding it)

16.5 Total Functions, Partial Functions, and Relations

From your days taking math classes you have a rough idea of what a function is: a mapping that takes one value to another. We can formalize that idea in set theory: if A is a set, and B is a set, then we have a set $A \rightarrow B$ of *total functions* from A to B . It is characterized as follows: $f \in A \rightarrow B$ if and only if

1. $f \subseteq A \times B$.
2. If $\langle a, b_1 \rangle \in f$ and $\langle a, b_2 \rangle \in f$ then $b_1 = b_2$. This says that f maps elements of A uniquely to elements of B . Notice that what this is saying is that $\langle a, b_1 \rangle = \langle a, b_2 \rangle$, so more informally, there is at most one pair in f with a in its first position.

¹⁷Bear in mind that since all we *really* have to work with are sets, we end up encoding pairs and other sequences using, well, sets. See https://en.wikipedia.org/wiki/Ordered_pair#Defining_the_ordered_pair_using_set_theory for a variety of pair encodings in set theory that have been developed by a variety of mathematicians and logicians.

3. If $a \in A$ then $\langle a, b \rangle \in f$ for some $b \in B$. This means informally that there is *at least* one pair in f with a in its first position.

We can relax the notion of total functions to get the notion of *partial functions*. Intuitively, a partial function is a mapping that doesn't necessarily map *every* element. From A and B , we can form the set of partial functions $A \rightarrow B$, where elements $f \in A \rightarrow B$ need only satisfy the first two properties of total functions: they don't have to map every element of A to some element in B , but any element that they do map, they map uniquely. It should be clear from this description that every total function is a partial function. Note as well, when we just say "function", we mean "total function."

A little bit of terminology about functions. Given a partial function $f : A \rightarrow B$, where this is just a common notation for $f \in A \rightarrow B$, we call A the function's *domain*; we call B the function's *codomain* (others call it the function's *range*). Sometimes when we need to talk about it, we may write $dom(f)$ and $cod(f)$ for the function's domain and codomain respectively.

Even more general than partial functions is the idea of *relations*. A relation captures many-to-many correspondences between elements of sets. The most direct analogue to a partial function is a *binary relation*, which is technically any set of pairs: a binary relation R on A and B is some subset $R \subseteq A \times B$. Relations need not be only binary: you can have a relation between three sets, i.e. $R \subseteq A \times B \times C$, and even a unary relation $R \subseteq A$, which is basically a way of representing some property of A 's by simply giving the subset of A that satisfies that property.

Notation We treat functions, partial functions, and relations with special notation sometimes. In particular, we introduce *function application* notation:

$$F(A) \equiv \iota S. (\exists D. \exists C. F \in D \rightarrow C) \wedge \langle A, S \rangle \in F.$$

In words, $F(A)$ denotes the unique set S such that F is a partial function from some set D to C and F maps A to S . The two existential quantifications constitute a sanity check. $F(A)$ is only a proper description if F is a partial function *and* F is defined for A . This in turn implies that A is in the F 's domain.

Thus, in the case of a partial function f , we write $f(a) = b$ to mean that $\langle a, b \rangle \in f$. We also write $f(a)$ to refer to the element b such that $\langle a, b \rangle \in f$. In the case of total functions, the expression $f(a)$ always makes sense, because there must be some b that fits the bill, but for partial functions, that may be nonsense, because the partial function is undefined at a . If a partial function f is undefined at a , we write $(f(a))\uparrow$ which literally means $\forall b \in cod(f). \langle a, b \rangle \notin f$. The opposite, that f is defined at a is written $(f(a))\downarrow$.

We often write binary relations using *infix* notation: we write $a R b$ to mean that $\langle a, b \rangle \in R$. For a familiar example, $1 \leq 2$ just means $\langle 1, 2 \rangle \in \leq$ where $\leq \subseteq \mathbb{N} \times \mathbb{N}$. For relations among more than two sets, we will often write them in "mixed-fix" notation. For example, when we discuss type systems, we will write the *typing judgment* $\Gamma \vdash t : T$ to mean that $\langle \Gamma, t, T \rangle \in \cdot \vdash \cdot : \cdot$ where $\cdot \vdash \cdot : \cdot \subseteq TENV \times TERM \times TYPE$. As with this example, we often use dots \cdot in the names of functions and relations to indicate the positions where arguments go. So we will refer to the typing judgment either by a usage, $\Gamma \vdash t : T$, or by the dotted name $\cdot \vdash \cdot : \cdot$. The nice thing about the usage version is that you can deduce the sets that make up the relation from its metavariables. The dot notation, though, is clearer about which elements are really parameters.

Mathematical Functions Don't Run! One thing that it might take time to get used to in Set Theory is that functions are not like functions in programming languages...they don't take an input, churn on it, and then spit out a result after a bunch of computation. They're NOT procedures! Rather, they are (possibly infinite) lookup tables. Often when we write down a function definition, using a bunch of *equations*, it looks almost exactly like a program in a functional programming language like Haskell,¹⁸ but really that is just a way of *naming* a particular function, just like we treat $\{a, b, c\}$ as the *name* of a set. Such names, or descriptions, give us certain reasoning principles, and it's that process of reasoning that corresponds to the kind of computation that a function definition in a computer does. So two mottos that arise are:

- *Functions are like database tables or spreadsheets, not like procedures.*
- *Computation arises from deduction, not from definition.*

¹⁸That's no coincidence: they purposely stole the notation!

We'll return to this discussion later because it's a common sticking point for programmers when it comes to math. At least it was a sticking point for *this* programmer!

16.6 Set Comprehensions

Most of the set constructions that we have been describing so far make “bigger” sets from smaller sets. We are very careful about the ways that we can build sets so that we don't write down a description of a set that can't possibly exist.¹⁹

However, once you have a set, you can always build a smaller set by choosing some of the elements of the set. We call this *set comprehension*: Given some set A , and some *predicate* $P(a)$ that describes properties of elements a of A , we can form the set $\{a \in A \mid P(a)\}$ which is the set of elements of A for which the predicate P holds.

One of the constructions we already developed, the intersection of two sets, can be described this way:

$$A \cap B = \{a \in A \mid a \in B\}.$$

In this case, the predicate $P(a) \equiv a \in B$. Notice that you can't use set comprehension to describe the *union* of two sets without already having something that contains the union. That's why we have an axiom for union, but not one for comprehension.

Another example, which we are introducing for the first time here, is the notion of *set difference*:

$$A \setminus B = \{a \in A \mid a \notin B\}.$$

Notice that set difference is always well defined for any two sets A and B .

At first you may not be clear on what counts as a legal predicate. In general, a predicate must be a formula in first-order logic about sets and their elements. Here I'll give some more examples of set descriptions using comprehensions. I recommend that you work your way through the logical statements to make sure that you understand how they work.

Partial Functions

$$A \rightarrow B = \{f \in \mathcal{P}(A \times B) \mid \forall a \in A. \forall b_1, b_2 \in B. \langle a, b_1 \rangle \in f \wedge \langle a, b_2 \rangle \in f \Rightarrow b_1 = b_2\}$$

Total Functions

$$\begin{aligned} A \rightarrow B &= \{f \in \mathcal{P}(A \times B) \mid \forall a \in A. \exists b \in B. \langle a, b \rangle \in f\} \\ &= \left\{ f \in \mathcal{P}(A \times B) \mid \begin{array}{l} (\forall a \in A. \exists b \in B. \langle a, b \rangle \in f) \wedge \\ (\forall a \in A. \forall b_1, b_2 \in B. \langle a, b_1 \rangle \in f \wedge \langle a, b_2 \rangle \in f \Rightarrow b_1 = b_2) \end{array} \right\} \end{aligned}$$

16.7 Set Replacements

In addition to element-by-element set-builder notation, we also have replacement-based notation and its expression-focused shorthand:

$$\begin{aligned} \{ B \leftarrow \Phi(A, B) \text{ for } A \in S \} \\ \{ F(A) \mid A \in S \} \end{aligned}$$

Each of which lets us describe a new set Q in terms of a pre-existing set S by relating elements A of S to elements B (or $F(A)$) of Q .

¹⁹For more on this, look at the Wikipedia page on “Paradoxes of set theory.”

17 Where is this coming from?

So....what hole did I pull these rules out of? Well, technically, “set theory” is not just one thing. There are lots of sets of rules that logicians and mathematicians over time have cobbled together, just like X86, Z80, and Power architectures have different instruction sets. However they all tend to have a pretty common core of things: you can build sets, take their union, intersection, yadayada.

The guiding principles behind these informal rules come from what is called Zermelo-Fraenkel set theory, after two logicians who came up with the rules²⁰.

18 The Syntax of First-Order Set Theory

So far we have been simply using logical notation without explanation, in the hope that seeing it in context would provide *some* understanding, much like an immersive natural language course. However logical notation is precise, and a precise understanding is necessary for mastery. Let me try for moment to be a bit more precise about at least the low-level parts of first-order set theory. We will not be formal about all of it, but let me introduce some of it.

A word of warning though. In the following I use formal notation like Backus-Naur Form (BNF) simply to suggest the structure of the language, assuming that you have seen BNF notation before and can intuitively interpret it. Later I will give a very formal interpretation to BNF for programming languages. One can apply the same analysis to the following development as well, but doing so would amount to “defining set theory using set theory”, kind of like writing a Python interpreter in Python.²¹

Representing a language inside itself does not amount to a fully-formal definition, but thanks to our human intuition and prior learning, we can still learn a lot from it. That’s the goal of this section.

The core of the syntax of logic can be described as follows:

$$\begin{aligned} \Xi &\in \text{PROPCONSTANT}, & \Psi &\in \text{ATOMICPROP}, & \Phi &\in \text{PROP}, & S &\in \text{SETVAR}, & \mathcal{E} &\in \text{SETEXP} \\ \Psi &::= \Xi \mid \mathcal{E} = \mathcal{E} \mid \mathcal{E} \in \mathcal{E} \\ \Phi &::= \Psi \mid \top \mid \perp \mid \Phi \wedge \Phi \mid \Phi \vee \Phi \mid \Phi \Rightarrow \Phi \mid \forall S. \Phi \mid \exists S. \Phi \\ \mathcal{E} &::= S \mid \iota S. \Phi \end{aligned}$$

Here the metavariables S denotes a set variable, while \mathcal{E} denotes a set expression. Both are means of giving a *name* for a set. Some example names that we use are TERM, or $\{1, 2, 3\}$. The latter is not a set in and of itself, but is a *descriptive* name: in the latter case we deduce properties of the set being named based on the structure of that name (whereas TERM is not immediately telling without knowing that the name TERM is equivalent to a more contentful set description). The propositional constants Ξ can be viewed as formal placeholders for real propositions. We use them to help explain the propositional connectives of our logic without having to worry about the particular content or meaning of an atomic proposition. This comes into play especially when we discuss our rules for logical deduction.

Now for a few common abbreviations. First, since we are not set-theorists—exploring the properties of *arbitrary sets*—but rather programming language theorists—exploiting set theory to study sets of programs and related structures—we will mostly be considering elements of *particular sets*. This leads to a desire to qualify our quantifiers. We do that by using the following notations that expand as described:

$$\begin{aligned} \forall S_1 \in S_2. \Phi &\equiv \forall S_1. S_1 \in S_2 \Rightarrow \Phi; \\ \exists S_1 \in S_2. \Phi &\equiv \exists S_1. S_1 \in S_2 \wedge \Phi. \end{aligned}$$

Take a moment to convince yourself that the expansions will have the intended meaning. It’s indeed a bit curious that the expansion is not the same for universal and existential quantification, but this is necessary to induce the desired interpretation.

²⁰Technically another logician named Skolem came up with the same part that Fraenkel did, but he got left out. Poor guy, that happens a lot in science and other forms of knowledge production.

²¹Nothing wrong with that! The PyPy project (<http://pypy.org>) has amply demonstrated benefits to doing this, and the LISP family of programming languages has a strong tradition of defining *metacircular interpreters* [McCarthy, 1960].

Also, since we often introduce many elements of the same set, we abbreviate this in terms of the above abbreviations:

$$\begin{aligned}\forall S_1, S_2, \dots, S_n \in S. \Phi &\equiv \forall S_1 \in S. \forall S_2 \in S. \dots \forall S_n \in S. \Phi; \\ \exists S_1 \in S_2. \Phi &\equiv \exists S_1 \in S. \exists S_2 \in S. \dots \exists S_n \in S. \Phi.\end{aligned}$$

Furthermore, we can abbreviate using the same kind of quantifier for different kinds of elements:

$$\begin{aligned}\forall S_1, S_2, \dots, S_n \in S, T_1, T_2, \dots, T_n \in T. \Phi &\equiv \forall S_1, S_2, \dots, S_n \in S. \forall T_1, T_2, \dots, T_n \in T. \Phi; \\ \exists S_1, S_2, \dots, S_n \in S, T_1, T_2, \dots, T_n \in T. \Phi &\equiv \exists S_1, S_2, \dots, S_n \in S. \exists T_1, T_2, \dots, T_n \in T. \Phi;\end{aligned}$$

Sometimes we wish to quantify elements of some n-ary relation. For example, Given some binary relation $\odot \subseteq S \times T$, we write $S_1 \odot T_1 \equiv \langle S_1, T_1 \rangle \in (\cdot \odot \cdot)$ (occasionally using the \cdot placeholders to indicate that we use *infix*, or more generally *mixed fix* notation) to indicate membership. Given such a relation, we use the following abbreviation to quantify elements of the relation in terms of elements of the subcomponents:

$$\begin{aligned}\forall (S_1 \odot T_2). \Phi &\equiv \forall S_1 \in S, T_2 \in T. S_1 \odot T_2 \Rightarrow \Phi; \\ \exists (S_1 \odot T_2). \Phi &\equiv \exists S_1 \in S, T_2 \in T. S_1 \odot T_2 \wedge \Phi.\end{aligned}$$

Finally, we introduce some richer logical forms that are interpreted using the basic ones above:

$$\begin{aligned}\neg \Phi &\equiv \Phi \Rightarrow \perp; \\ \Phi_1 \Leftrightarrow \Phi_2 &\equiv (\Phi_1 \Rightarrow \Phi_2) \wedge (\Phi_2 \Rightarrow \Phi_1); \\ \exists! S. \Phi(S) &\equiv (\exists S. \Phi(S)) \wedge (\forall S_1, S_2. \Phi(S_1) \wedge \Phi(S_2) \Rightarrow S_1 = S_2).\end{aligned}$$

The first one represents “not Φ ”, the second represents “if and only if”; and the third represents “there exists a unique set S such that Φ .” The notation $\Phi(S)$ says that Φ is a formula that *may* refer to S freely in it. Once this notation is introduced, then $\Phi(T)$ refers to the same formula but with T substituted for S .

References

- J. Avigad. Reliability of mathematical inference. *Synthese*, 2020. doi: 10.1007/s11229-019-02524-y. <https://doi.org/10.1007/s11229-019-02524-y>.
- J. Bagaria. Set theory. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, winter 2014 edition, 2014. URL <http://plato.stanford.edu/archives/win2014/entries/set-theory/>.
- A. Bauer. Proof of negation and proof by contradiction. *Mathematics and Computation Blog*, March 2010. <http://math.andrej.com/2010/03/29/proof-of-negation-and-proof-by-contradiction/>.
- L. Crosilla. Set Theory: Constructive and Intuitionistic ZF. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2020 edition, 2020.
- J. Ferreirós. The early development of set theory. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2019 edition, 2019. URL <https://plato.stanford.edu/archives/sum2019/entries/settheory-early/>.
- P. R. Halmos. *Naive Set Theory*. Springer-Verlag, first edition, Jan. 1960. ISBN 0387900926. A classic introductory textbook on set theory.
- P. Maddy. Believing the axioms. I. *The Journal of Symbolic Logic*, 53(02):481–511, 1988. An interesting (though complicated) analysis of why set theorists believe in their axioms.

- J. McCarthy. Recursive functions of symbolic expressions and their computation by machine, part I. *Commun. ACM*, 3(4):184–195, 1960.
- J. Spolsky. The law of leaky abstractions. Joel on Software Blog, November 2002. <https://www.joelonsoftware.com/2002/11/11/the-law-of-leaky-abstractions/>.
- S. Stenlund. Descriptions in intuitionistic logic. In S. Kanger, editor, *Proceedings of the Third Scandinavian Logic Symposium*, volume 82 of *Studies in Logic and the Foundations of Mathematics*, pages 197 – 212. Elsevier, 1975. URL <http://www.sciencedirect.com/science/article/pii/S0049237X08707328>.
- M. Tiles. Book Review: Stephen Pollard. *Philosophical Introduction to Set Theory*. *Notre Dame Journal of Formal Logic*, 32(1):161–166, 1990.
A brief introduction to the philosophical issues underlying set theory as a foundation for mathematics.