Chapter 5

Sets

So far, we've been assuming only a basic understanding of sets. It's time to discuss sets systematically, including a useful range of constructions, operations, notation, and special cases. We'll also see how to compute the sizes of sets and prove claims involving sets.

5.1 Sets

Sets are an extremely general concept, defined as follows:

> Definition: A set is an unordered collection of objects.

For example, the natural numbers are a set. So are the integers between 3 and 7 (inclusive). So are all the planets in this solar system or all the programs written by students in CS 225 in the last three years. The objects in a set can be anything you want.

The items in the set are called its elements or members. We've already seen the notation for this: $x \in A$ means that x is a member of the set A.

There's three basic ways to define a set:

- describe its contents in mathematical English, e.g. "the integers between 3 and 7, inclusive."

- list all its members, e.g. $\{3, 4, 5, 6, 7\}$
- use so-called set builder notation, e.g. $\{x \in \mathbb{Z} \mid 3 \leq x \leq 7\}$

Set builder notation has two parts separated with a vertical bar or a colon. The first part names a variable (in this case x) that ranges over all objects in the set. The second part gives one or more constraints that these objects must satisfy, e.g. $3 \leq x \leq 7$. The type of the variable (integer in our example) can be specified either before or after the vertical bar. The separator (| or :) is often read "such that."

Here's an example of a set containing an infinite number of objects

- "multiples of 7"
- $\{\ldots - 14, -7, 0, 7, 14, 21, 18, \ldots\}$
- $\{x \in \mathbb{Z} \mid x \text{ is a multiple of } 7\}$

We couldn't list all the elements, so we had to use "...". This is only a good idea if the pattern will be clear to your reader. If you aren't sure, use one of the other methods.

If we wanted to use shorthand for "multiple of", it might be confusing to have | used for two different purposes. So it would probably be best to use the colon variant of set builder notation:

$\{x \in \mathbb{Z} : 7 \mid x\}$

5.2 Things to be careful about

A set is an unordered collection. So $\{1, 2, 3\}$ and $\{2, 3, 1\}$ are two names for the same set. Each element occurs only once in a set. Or, alternatively, it doesn't matter how many times you write it. So $\{1, 2, 3\}$ and $\{1, 2, 3, 2\}$ also name the same set.

We've seen ordered pairs and triples of numbers, such as $(3, 4)$ and $(4, 5, 2)$. The general term for an ordered sequence of k numbers is a k-tuple.[1] Tuples are very different from sets, in that the order of values in a tuple matters and duplicate elements don't magically collapse. So $(1, 2, 2, 3) \neq (1, 2, 3)$ and $(1, 2, 2, 3) \neq (2, 2, 1, 3)$. Therefore, make sure to enclose the elements of a set in curly brackets and carefully distinguish curly brackets (set) from parentheses (ordered pair).

A more subtle feature of tuples is that a tuple must contain at least two elements. In formal mathematics, a 1-dimensional value x is just written as x, not as (x). And there's no such thing in mathematics as a 0-tuple. So a tuple is simply a way of grouping two or more values into a single object.

By contrast, a set is like a cardboard box, into which you can put objects. A kitty is different from a box containing a kitty. Similarly, a set containing a single object is different from the object by itself. For example, $\{57\}$ is not the same as 57. A set can also contain nothing at all, like an empty box. The set containing nothing is called the empty set or the null set, and has the shorthand symbol $\emptyset$.[2]

The empty set may seem like a pain in the neck. However, computer science applications are full of empty lists, strings of zero length, and the like. It's the kind of special case that all of you (even the non-theoreticians) will spend your life having to watch out for.

Both sets and tuples can contain objects of more than one type, e.g. $(\text{cat}, \text{Fluffy}, 1983)$ or $\{a, b, 3, 7\}$. A set can also contain complex objects, e.g. $\{(a, b), (1, 2, 3), 6\}$ is a set containing three objects: an ordered pair, an ordered triple, and a single number.

5.3 Cardinality, inclusion

If A is a finite set (a set containing only a finite number of objects), then $|A|$ is the number of (different) objects in A. This is also called the **cardinality**

[1] There are latinate terms for longer sequences of numbers, e.g. quadruple, but they aren't used much.

[2] Don't write the emptyset as $\{\}$. Like a spelling mistake, this will make readers think less of your mathematical skills.

of A. For example, $|\{a, b, 3\}| = 3$. And $|\{a, b, a, 3\}|$ is also 3, because we count a group of identical objects only once. The notation of cardinality also extends to sets with infinitely many members ("infinite sets") such as the integers, but we won't get into the details of that right now.

Notice that the notation $|A|$ might mean set cardinality or it might be the more familiar absolute value. To tell which, figure out what type of object A is. If it's a set, the author meant cardinality. If it's a number, the author meant absolute value.

If A and B are sets, then A is a subset of B (written $A \subseteq B$) if every element of A is also in B. Or, if you want it formally: $\forall x, x \in A \to x \in B$. For example, $\mathbb{Q} \subseteq \mathbb{R}$, because every member of the rationals is also a member of the reals.

The notion of subset allows the two sets to be equal. So $A \subseteq A$ is true for any set A. So $\subseteq$ is like $\leq$. If you want to force the two sets to be different (i.e. like $<$), you must say that A is a **proper** subset of B, written $A \subset B$. You'll occasionally see reversed versions of these symbols to indicate the opposite relation, e.g. $B \supseteq A$ means the same as $A \subseteq B$.

5.4 Vacuous truth

If we have a set A, an interesting question is whether the empty set should be considered a subset of A. To answer this, let's first back up and look at one subtlety of mathematical logic.

Consider the following claim:

Claim 27 *For all natural numbers n, if $14 + n < 10$, then n wood elves will attack Siebel Center tomorrow.*

I claim this is true, a fact which most students find counter-intuitive. In fact, it wouldn't be true if n was declared to be an integer.

Notice that this statement has the form $\forall n, P(n) \to Q(n)$, where $P(n)$ is the predicate $14+n < 10$. Because n is declared to be a natural number, n is never negative, so $n + 14$ will always be at least 14. So $P(n)$ is always false.

Therefore, our conventions about the truth values for conditional statements imply that $P(n) \to Q(n)$ is true. This argument works for any choice of n. So $\forall n, P(n) \to Q(n)$ is true.

Because even mathematicians find such statements a bit wierd, they typically say that such a claim is *vacuously* true, to emphasize to the reader that it is only true because of this strange convention about the meaning of conditionals. Vacuously true statements typically occur when you are trying to apply a definition or theorem to a special case involving an abnormally small or simple object, such as the empty set or zero or a graph with no arrows at all.

In particular, this means that the empty set is a subset of any set A. For $\emptyset$ to be a subset of A, the definition of "subset" requires that for every object x, if x is an element of the empty set, then x is an element of A. But this if/then statement is considered true because its hypothesis is always false.

5.5 Set operations

Given two sets A and B, the intersection of A and B ($A \cap B$) is the set containing all objects that are in both A and B. In set builder notation:

$$A \cap B = \{S \mid S \in A \text{ and } S \in B\}$$

Let's set up some sample sets:

- $M = \{\text{egg}, \text{bread}, \text{milk}\}$
- $P = \{\text{milk}, \text{egg}, \text{flour}\}$

Then $M \cap P$ is $\{\text{milk}, \text{egg}\}$.

If the intersection of two sets A and B is the empty set, i.e. the two sets have no elements in common, then A and B are said to be **disjoint**.

The union of sets A and B ($A \cup B$) is the set containing all objects that are in one (or both) of A and B. So $M \cup P$ is $\{\text{milk}, \text{egg}, \text{bread}, \text{flour}\}$.

The set difference of A and B ($A - B$) contains all the objects that are in A but not in B. In this case,

$$M - P = \{\text{bread}\}$$

The complement of a set A ($\overline{A}$) is all the objects that aren't in A. For this to make sense, you need to define your "universal set" (often written U). U contains all the objects of the sort(s) you are discussing. For example, in some discussions, U might be all real numbers. U doesn't contain everything you might imagine putting in a set, because constructing a set that inclusive leads to paradoxes. U is more limited than that. Whenever U is used, you and your reader need to come to an understanding about what's in it.

So, if our universe is all integers, and A contains all the multiples of 3, then $\overline{A}$ is all the integers whose remainder mod 3 is either 1 or 2. $\overline{\mathbb{Q}}$ would be the irrational numbers if our universe is all real numbers. If we had been working with complex numbers, it might be the set of all irrational real numbers plus all the numbers with an imaginary component.

If A and B are two sets, their Cartesian product ($A \times B$) contains all ordered pairs (x, y) where x is in A and y is in B. That is

$$A \times B = \{(x, y) \mid x \in A \text{ and } y \in B\}$$

For example, if $A = \{a, b\}$ and $B = \{1, 2\}$, then

$$A \times B = \{(a, 1), (a, 2), (b, 1), (b, 2)\}$$
$$B \times A = \{(1, a), (2, a), (1, b), (2, b)\}$$

Notice that these two sets aren't equal: order matters for Cartesian products.

5.6 Set identities

It's easy to find (e.g. on the net), long lists of identities showing when two sequences of set operations yield the same output set. For example:

DeMorgan's Law: $\overline{A \cup B} = \overline{A} \cap \overline{B}$

I won't go through these in detail because they are largely identical to the identities you saw for logical operations, if you make the following correspondences:

- $\cup$ is like $\vee$
- $\cap$ is like $\wedge$
- $\overline{A}$ is like $\neg P$
- $\emptyset$ (the empty set) is like F
- U (the universal set) is like T

The two systems aren't exactly the same. E.g. set theory doesn't use a close analog of the $\rightarrow$ operator. But they are very similar.

5.7 Size of set union

Many applications require that we calculate[3] the size of the set created by applying set operations. These sets are often sets of options for some task.

If two sets A and B don't intersect, then the size of their union is just the sum of their sizes. That is: $|A \cup B| = |A| + |B|$. For example, suppose that it's late evening and you want to watch a movie. You have 37 movies on cable, 57 DVD's on the shelf, and 12 movies stored in I-tunes. If these three sets of movies don't intersect, you have a total of $37 + 57 + 12 = 106$ movies.

If your input sets do overlap, then adding up their sizes will double-count some of the objects. So, to get the right number for the union, we need to correct for the double-counting. For our movie example, suppose that the only overlap is that 2 movies are on I-tunes and also on DVD. Then you would have $(37 + 57 + 12) - 2 = 104$ movies to choose from.

[3] Or sometimes just estimate.

The formal name for this correction is the "Inclusion-Exclusion Principle". Formally, suppose you have two sets A and B. Then

Inclusion-Exclusion Principle: $|A \cup B| = |A| + |B| - |A \cap B|$

We can use this basic 2-set formula to derive the corresponding formula for three sets A, B, and C:

$$\begin{aligned}
|A \cup B \cup C| &= |A| + |B \cup C| - |A \cap (B \cup C)| \\
&= |A| + |B| + |C| - |B \cap C| - |A \cap (B \cup C)| \\
&= |A| + |B| + |C| - |B \cap C| - |(A \cap B) \cup (A \cap C)| \\
&= |A| + |B| + |C| - |B \cap C| - (|A \cap B| + |A \cap C| - |(A \cap B) \cap (A \cap C)|) \\
&= |A| + |B| + |C| - |B \cap C| - |A \cap B| - |A \cap C| + |A \cap B \cap C|
\end{aligned}$$

In addition to the inclusion-exclusion principle, this derivation uses the distributive law (third step) and the fact that intersection is commutative and associative (last step).

5.8 Product rule

Now, suppose that we form the Cartesian product of two sets A and B, where $|A| = n$ and $|B| = q$. To form an element (x, y) in the product, we have n choices for x and q choices for y. So we have nq ways to create an element of the product. So $|A \times B| = nq$.

In general:

> **The product rule**: if you have p choices for one part of a task, then q choices for a second part, and your options for the second part don't depend on what you chose for the first part, then you have pq options for the whole task.

This rule generalizes in the obvious way to sequences of several choices: you multiply the number of options for each choice. For example, suppose that T-shirts come in 4 colors, 5 sizes, 2 sleeve lengths, and 3 styles of neckline, there are $4 \cdot 5 \cdot 2 \cdot 3 = 120$ total types of shirts.

We could represent a specific T-shirt type as a 4-tuple (c, s, l, n): c is its color, s is its size, l is its sleeve length, and n is its neckline. E.g. one T-shirt type is $(\text{red}, \text{small}, \text{long}, \text{vee})$ The set of all possible T-shirt types would then be a 4-way Cartesian product of the set of all colors, the set of all sizes, the set of all sleeve lengths, and the set of all necklines.

5.9 Combining these basic rules

These two basic counting rules can be combined to solve more complex practical counting problems. For example, suppose we have a set S which contains all 5-digit decimal numbers that start with 2 one's or end in 2 zeros, where we don't allow leading zeros. How large is S, i.e. how many numbers have this form?

Let T be the set of 5-digit numbers starting in 2 one's. We know the first two digits and we have three independent choices (10 options each) for the last three. So there are 1000 numbers in T.

Let R be the set of 5-digit numbers ending in 2 zeros. We have 9 options for the first digit, since it can't be zero. We have 10 options each for the second and third digits, and the last two are fixed. So we have 900 numbers in R.

What's the size of $T \cap R$? Numbers in this set start with 2 one's and end with 2 zeros, so the only choice is the middle digit. So it contains 10 numbers. So

$$|S| = |T| + |R| - |T \cap R| = 1000 + 900 - 10 = 1890$$

5.10 Proving facts about set inclusion

So far in school, most of your proofs or derivations have involved reasoning about equality. Inequalities (e.g. involving numbers) have been much less common. With sets, the situation is reversed. Proofs typically involve reasoning about subset relations, even when proving two sets to be equal. Proofs that rely primarily on a chain of set equalities do occur, but they are much less common. Even when both approaches are possible, the approach based on subset relations is often easier to write and debug.

As a first example of a typical set proof, consider the following claim:

Claim 28 *For any sets A, B, and C, if $A \subseteq B$ and $B \subseteq C$, then $A \subseteq C$.*

This property is called "transitivity," just like similar properties for (say) $\leq$ on the real numbers. Both $\subseteq$ and $\leq$ are examples of a general type of object called a *partial order*, for which transitivity is a key defining property.

First, remember our definition of $\subseteq$: a set A is a subset of a set B if and only if, for any object x, $x \in A$ implies that $x \in B$.

> Proof: Let A, B, and C be sets and suppose that $A \subseteq B$ and $B \subseteq C$.

Our ultimate goal is to show that $A \subseteq C$. This is an if/then statement: for any x, if $x \in A$, then $x \in C$. So we need to pick a representative x and assume the hypothesis is true, then show the conclusion. So our proof continues:

> Let x be an element of A. Since $A \subseteq B$ and $x \in A$, then $x \in B$ (definition of subset). Similarly, since $x \in B$ and $B \subseteq C$, $x \in C$. So for any x, if $x \in A$, then $x \in C$. So $A \subseteq C$ (definition of subset again). $\square$

5.11 Example proof: deMorgan's law

Now, suppose that we want to to prove that two sets A and B are equal. It is sometimes possible to do this via a chain of equalities (e.g. using set operation identities), but this strategy works less often for sets than it does for (say) facts about the real numbers.

A more common method is to show that $A \subseteq B$ and $B \subseteq A$, using separate subproofs. We can then conclude that $A = B$. This is just like showing that $x = y$ by showing that $x \leq y$ and $y \leq x$. Although it seems like more trouble to you right now, this is a more powerful approach that works in a wider range of situations, especially in upper-level computer science and mathematics courses (e.g. real analysis, algorithms).

As an example, let's look at

> Claim (DeMorgan's Law): For any sets A and B, $\overline{A \cup B} = \overline{A} \cap \overline{B}$

We'll prove this in two parts:

> Proof: Let A and B be sets. We'll do this in two parts:
>
> $\overline{A \cup B} \subseteq \overline{A} \cap \overline{B}$: Let $x \in \overline{A \cup B}$. Then $x \notin A \cup B$. So it's not the case that ($x \in A$ or $x \in B$). So, by deMorgan's Law for logic, $x \notin A$ and $x \notin B$. That is $x \in \overline{A}$ and $x \in \overline{B}$. So $x \in \overline{A} \cap \overline{B}$.
>
> $\overline{A} \cap \overline{B} \subseteq \overline{A \cup B}$: Let $x \in \overline{A} \cap \overline{B}$. Then $x \in \overline{A}$ and $x \in \overline{B}$. So $x \notin A$ and $x \notin B$. So, by deMorgan's law for logic, it's not the case that ($x \in A$ or $x \in B$) So $x \notin A \cup B$. That is $x \in \overline{A \cup B}$.

In this example, the second half is very basically the first half written backwards. However, this is not always the case. The power of this method comes from the fact that, in harder problems, the two halves of the proof can use different techniques.

5.12 An example with products

Let's move on to some facts that aren't so obvious, because they involve Cartesian products.

Claim 29 *If A, B, C, and D are sets such that $A \subseteq B$ and $C \subseteq D$, then $A \times C \subseteq B \times D$.*

To prove this, we first gather up all the facts we are given. What we need to show is the subset inclusion $A \times C \subseteq B \times D$. To do this, we'll need to pick a representative element from $A \times C$ and show that it's an element of $B \times D$.

> Proof: Let A, B, C, and D be sets and suppose that $A \subseteq B$ and $C \subseteq D$.
>
> Let $p \in A \times C$. By the definition of Cartesian product, p must have the form (x, y) where $x \in A$ and $y \in C$.
>
> Since $A \subseteq B$ and $x \in A$, $x \in B$. Similarly, since $C \subseteq D$ and $y \in C$, $y \in D$. So then $p = (x, y)$ must be an element of $B \times D$.
>
> We've shown that, for all p, if $p \in A \times C$ then $p \in B \times D$. This means that $A \times C \subseteq B \times D$. $\square$

The last paragraph is optional. When you first start, it's a useful recap because you might be a bit fuzzy about what you needed to prove. As you get experienced with this sort of proof, it's often omitted. But you will still see it occasionally at the end of a very long (e.g. multi-page) proof, where even an experienced reader might have forgotten the main goal of the proof.

5.13 Another example with products

Here's another claim about Cartesian products:

Claim 30 *For any sets A, B, and C, if $A \times B \subseteq A \times C$ and $A \neq \emptyset$, then $B \subseteq C$.*

Notice that this claim fails if $A = \emptyset$. For example, $\emptyset \times \{1, 2, 3\}$ is a subset of $\emptyset \times \{a, b\}$, because both of these sets are empty. However $\{1, 2, 3\}$ is not a subset of $\{a, b\}$.

This is like dividing both sides of an algebraic equation by a non-zero number: if $xy \leq xz$ and $x \neq 0$ then $y \leq z$. This doesn't work we allow x to be zero. Set operations don't always work exactly like operations on real numbers, but the parallelism is strong enough to suggest special cases that ought to be investigated.

A general property of proofs is that the proof should use all the information in the hypothesis of the claim. If that's not the case, either the proof has a bug (e.g. on a homework assignment) or the claim could be revised to make it more interesting (e.g. when doing a research problem, or a buggy homework problem). Either way, there's an important issue to deal with. So, in this case, we need to make sure that our proof does use the fact that $A \neq \emptyset$.

Here's a draft proof:

> Proof draft 1: Suppose that A, B, C, and D are sets and suppose that $A \times B \subseteq A \times C$ and $A \neq \emptyset$. We need to show that $B \subseteq C$.
>
> So let's choose some $x \in B$. ...

The main fact we've been given is that $A \times B \subseteq A \times C$. To use it, we need an element of $A \times B$. Right now, we only have an element of B. We need to find an element of A to pair it with. To do this, we reach blindly into A, pull out some random element, and give it a name. But we have to be careful here: what if A doesn't contain any elements? So we have to use the assumption that $A \neq \emptyset$.

> Proof draft 1: Suppose that A, B, C, and D are sets and suppose that $A \times B \subseteq A \times C$ and $A \neq \emptyset$. We need to show that $B \subseteq C$.
>
> So let's choose some $x \in B$. Since $A \neq \emptyset$, we can choose an element t from A. Then $(t, x) \in A \times B$ by the definition of Cartesian product.
>
> Since $(t, x) \in A \times B$ and $A \times B \subseteq A \times C$, we must have that $(t, x) \in A \times C$ (by the definition of subset). But then (again by the definition of Cartesian product) $x \in C$.
>
> So we've shown that if $x \in B$, then $x \in C$. So $B \subseteq C$, which is what we needed to show.

5.14 A proof using sets and contrapositive

Here's a claim about sets that's less than obvious:

Claim 31 *For any sets A and B, if $(A - B) \cup (B - A) = A \cup B$ then $A \cap B = \emptyset$.*

Notice that the conclusion $A \cap B = \emptyset$ claims that something does not exist (i.e. an object that's in both A and B). So this is a good place to apply proof by contrapositive.

> Proof: Let's prove the contrapositive. That is, we'll prove that if $A \cap B \neq \emptyset$, then $(A - B) \cup (B - A) \neq A \cup B$.
>
> so, let A and B be sets and suppose that. $A \cap B \neq \emptyset$. Since $A \cap B \neq \emptyset$, we can choose an element from $A \cap B$. Let's call it x.
>
> Since x is in $A \cap B$, x is in both A and B. So x is in $A \cup B$.
>
> However, since x is in B, x is not in $A - B$. Similarly, since x is in A, x is not in $B - A$. So x is not a member of $(A - B) \cup (B - A)$. This means that $(A - B) \cup (B - A)$ and $A \cup B$ cannot be equal, because x is in the second set but not in the first. $\square$.

5.15 Variation in notation

Tuples containing tuples are, formally, different from longer tuples. E.g. $((a, b), c)$ is formally a different object from (a, b, c), and $(a, (b, c))$ is different from both. Most mathematical writers make a note of this formal difference, but then immediately say they will treat all of these forms as interchangeable. However, you will find an occasional author who is picky about the difference.

Linked lists used in computer science behave very differently from mathematical tuples. E.g. the linked list $((a, b), c)$ is completely different from the list (a, b, c). This can cause confusion because computer science lists are often written just like mathematical tuples. When in doubt, carefully examine the author's examples to see which convention they are using.