Jacobson, Pauline. 2014. Compositional semantics: An Introduction to the Syntax/Semantics Interface. Oxford University Press.

1.5. Appendix: Sets and Functions

1.5.1. Sets, members, and subsets

Since the notions of sets and of functions are crucial throughout this book, some formal definitions and discussion are provided here for readers not entirely familiar with these notions. We begin with the notion of a *set*. A set is simply any collection of objects (it can have a finite number of objects, an infinite number, or none at all). For example, we can talk about the set of positive integers less than 10; sets can be notated by listing the members and enclosing the list in curly brackets: {1,2,3,4,5,6,7,8,9}. The order in which they are listed makes no difference; a set is just a collection of things without any order. So if we were to write {2,5,3,4,9,7,8,1,6}, this names the same set. Each integer in this set is called a *member* or an *element* of the set. If we were to name this set A, then the notation $4 \in A$ means that 4 is a member (or element) of A. Something either is or is not in a set; it makes no sense to say it occurs twice (or more) in the set. Note also that a set can have a single member; this is called a *singleton set*. Thus {4} is the set with only one member; this set is distinct from 4 itself. (4 is a member of {4}.)

A set can have an infinite number of members; the set of positive integers for example is infinite. Obviously this can't be named by listing the members. One can in this case specify the set by a recursive procedure. Call the set I, then one can specify I by two statements: (a) (what is known as the base step): $1 \in I$, and (b) (the recursion step) if $n \in I$ then $n+1 \in I$. (It is understood when one lists things this way that nothing else is in I.) One will also often see a notation which describes rather than lists the members. For example, we can write the following set, call it B: {x|x is a New England state}. This names a finite set, and so we could also give B in list form as follows: {Maine, New Hampshire, Vermont, Massachusetts, Rhode Island, Connecticut}. These two are just different notations for naming the same set. This can also be used, of course, for infinite sets. Take, for example, the set

$\{x|x$ is an integer and $x > 9\}$. This names the set of integers greater than 9. And, a set can have no members. There is only one such set; its name is *the null set* or *the empty set*, and is generally written as Ø. Of course, there are other ways one can describe the null set. For example, the set of integers each of which is greater than 9 and less than 10 is the empty set. The *cardinality* of some set refers to the number of elements in that set; the notation |B| means the cardinality of B. Hence, given our set B above, |B| is six.

Take some set A. Then a *subset* of A is any set all of whose members are also in A. Suppose, for example, we begin with a set C which is $\{1,2,3\}$. Then $\{1,2\}$ is a subset of C, as is $\{1,3\}$ and so forth. The notation for the subset relation is $\subseteq$. The full definition of subset is as follows: $B \subseteq A$ if and only if every member of B is a member of A. From this it follows that every set is a subset of itself (so for the set C above, one of its subsets is the set $\{1,2,3\}$). It is, however, sometimes convenient to refer to those subsets distinct from the original set; in that case we can talk about a *proper subset* of some set. The symbol for this is $\subset$, so $B \subset A$ if and only if $B \subseteq A$ and $B \neq A$. Since the definition of subset says that B is a subset of A if and only if everything that is in B is also in A, it follows that if nothing is in B then B is a subset of A. thus the null set is a subset of every other set. Sets themselves can have sets as members, and so one can talk about the set of all subsets of a set A. This is called the *power set* of A, written as $\mathscr{P}(A)$. For example, given the set C above, $\mathscr{P}(A) = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1,2\}, \{1,3\}, \{2,3\}, \{1,2,3\}\}$.

***1.3.** If a set A has n members, then the number of subsets of A is 2^n. Try to see why this is true. *Hint:* for every member x of some set A, then for each subset B of A, x is either in B or is not in B.

1.4. How many members does the following set have: $\{\emptyset\}$?

***1.5.** What is $\mathscr{P}(\emptyset)$?

We will also have occasion to talk about the reverse of the subset relation—i.e., the *superset* relation. A is a superset of B if and only if B is a subset of A. The notation for this is $A \supseteq B$. Once again this is defined in such a way that every set is a superset of itself; a superset of B which is not identical to B is called a *proper superset*, and the notation for this is $\supset$.

1.5.2. Union, intersection, and complement

Take any two sets A and B. Then there is a set C which consists of everything that is in A and everything that is in B. This is called the *union* of A and B, and is written $A \cup B$. For example, if A is {1,2,3} and B is {2,4,6} then $A \cup B$ is {1,2,3,4,6}. Moreover, for any two sets A and B the *intersection* of A and B is the set of all things that are in both A and B. This is written $A \cap B$. So, for example, in the case directly above, the intersection of A and B is {2}. Or, if we were to intersect the set of integers which can be evenly divided by 2 (the set of even integers) with the set of integers which can be evenly divided by 3, we end up with the set of integers that can be evenly divided by 6.

1.6. a. For any two sets A and B such that $A \subseteq B$, what set is $A \cup B$?
b. For any two sets A and B such that $A \subseteq B$, what set is $A \cap B$?

One final useful notion here is the *complement* of a set. The complement of some set A is the set of all things which are not in A (this is sometimes notated as A′). Usually one talks about this notion with respect to some larger domain. Strictly speaking, the complement of {1,2,3} would include not only all integers greater than 3 but also all sorts of other numbers (like 1/3), the sun, my dog Kiana, and the kitchen sink. Rarely are we interested in that sort of set; so in practice when one talks about "the complement of some set A" this is generally with respect to some larger set B of which A is a subset. Then the complement of A refers to all things in B that are not in A (this is notated as B-A). For example, when restricting the discussion to the set of positive integers, the complement of {1,2,3} is the set of all integers greater than 3.

1.5.3. Ordered pairs, relations, equivalence relations, and partitions

Sets are unordered collections of objects. But it is quite useful (as will become very apparent as this book proceeds) to be able to talk about pairs of objects that are ordered in some way. An *ordered pair* is just that:

it is two objects with some ordering between them. If the two objects are a and b, then (a,b) is an ordered pair; (b,a) is a different ordered pair. An ordered pair need not contain distinct items: (a,a) is an ordered pair. In applying this to actual natural relations that exist in the world we are generally interested in *sets of ordered pairs*. (One can generalize this notion to ordered triples and so forth; an *ordered n-tuple* means an ordered list of n items.)

This notion is easiest to grasp with some concrete examples. Take again the set {1,2,3}, and take the relation "is greater than." Then this can be seen as a set of ordered pairs; if we are restricting this to items from our little 1-2-3 set, this would be the set {(2,1), (3,1), (3,2)}. Now suppose we instead take the following set of ordered pairs: {(2,1), (3,1), (3,2), (1,1), (2,2), (3,3)}. Then (restricting this again to our 1-2-3 set) we have now actually listed the relation "is greater than or equal to." Or, take the set {(1,1), (2,2), (3,3)}. That is the relation "is equal to" (defined for the set of integers {1,2,3}).

In other words, what we are calling a *relation* is just some set of ordered pairs. In the example above, both the first and second member of each ordered pair was drawn from the same set (the set {1,2,3}). But this is not necessary; we can have a set of ordered pairs each of whose first member is drawn from some set A and the second member from some set B where A and B are different (they can, but need not, have some of the same members). For example, the relation "is the capital of" is a relation between cities and states; it can be expressed as a set of ordered pairs of the general form {(Providence, Rhode Island), (Boston, Massachusetts), (Springfield, Illinois), (Pierre, South Dakota), . . . } (the . . . here is a shorthand for the remaining 46 pairs).

Take two sets A and B (they could be the same set or different). Then A x B refers to the set of all ordered pairs whose first member is in A and whose second member is in B. (This is also called the *Cartesian product* of A and B.) As in the case above, it is helpful to give the intuition of this by coming up with some concrete example. Suppose we take as our set A some group of professors—say, Professor Magoo, Professor Carberry, and Professor Glazie. Call that set P (for shorthand, let's call its members m, c, and g, so P is the set {m,c,g}). Now suppose we have a set S which consists of three students who we will just indicate as x, y, and z (so S = {x,y,z}). Then P x S = {(m,x}, (m,y), (m,z), (c,x), (c,y), (c,z), (g,x), (g,y), (g,z)}. Suppose that Magoo wrote a letter of recommendation for all three students, Carberry wrote one for only y, and Glazie wrote one for y and z. Then the relation "wrote a recommendation for" is a subset of P x S, and is the set of ordered pairs {(m,x}, (m,y), (m,z), (c,y), (g,y), (g,z)}.

More generally, we define a *relation* (between members of A and B) as any subset of A x B. There are some special and interesting properties that can be defined when the two sets are the same. That is, we are now looking at subsets of A x A. Consider a relation R (some subset of A x A) which is such that for all x in A, (x,x) is in R. Such a relation is called a *reflexive* relation. (These need not be the only kinds of pairs to be in R for R to be reflexive; other pairs can be in there too.) For example, if talking about the set of integers again, the relation "is greater than or equal to" is reflexive; for all numbers n, (n,n) is in the set of ordered pairs described by that relation. A relation R is called *irreflexive* if for all x in A, (x,x) is *not* in R. Further, consider any two members x and y (both members of A). Then if it's the case that for all x and y if (x,y) is in R then (y,x) is also in R, the relation is called *symmetric.* Imagine, for example, a lovely world with no unrequited love. Then *is in love with* is symmetric in that world. If our set were {m, c, g, and p}, then if the pair (m,c) were in our relation R (i.e., "is in love with") the fact that R is symmetric means that (c,m) is also in R. (Notice that our definition neither requires (c,c) to be in R nor excludes that; either is possible.) Or, to look at a relation which is symmetric by definition: consider the relation *is a sibling of*. (While *is a sibling of* is symmetric, *is a sister of* is not. Why not?) One final useful definition is a *transitive relation*. A transitive relation R is one for which for every x, y, and z, if (x,y) is in R and (y,z) is in R, then (x,z) is in R. (The relation "is greater than" is transitive, as is the relation "is greater than or equal to").

Any relation R which is reflexive, transitive, and symmetric is called an *equivalence relation*. As an example of such a relation, consider the set of students (call it S) at an elementary school that services grades 1 through 6. Then "is in the same grade as" is an equivalence relation in S x S. (While it is unusual to use the phrase "in the same grade as" when referring to the same person it seems false to say *Johnny is not in the same grade as himself* so we can see that this relation is reflexive.) It is also obvious that it is symmetric and transitive. Note that this—and any other equivalence relation—divides up the original set (here, S) into a group of non-overlapping subsets. The set of these subsets is called a *partition*. Thus, a *partition* of any set S is a set of subsets of S such that for each distinct subset A and B, $A \cap B = \emptyset$, and the union of all the subsets is S. To show that any equivalence relation induces such a partition, take any x in S and define S_x as {y|(y,x) is in R}. Since R is reflexive, we know that x is in S_x (and hence we know that S_x is guaranteed not to be empty). Moreover, the fact that R is reflexive means that each member of S is guaranteed to be in at least one such subset, so we know that the union of all of these is S. We can further show that for any two such subsets S_a and S_b, they either have no members in common (i.e., they have a null intersection) or they are the same. Thus, take any c which is in both S_a and S_b. By definition, this means that (c,a) is in R and (c,b) is in R. By the fact that R is transitive and symmetric, it follows that (a,b) and (b,a) are in R (the reader can work through the necessary steps). But then, for all x such that (x,a) is in R, (x,b) is also in R. To show this note again that R is transitive. If (x,a) is in R and (a,b) is in R then (x,b) is also in R. Hence given the initial premise that there is a non-empty intersection between S_a and S_b, it follows that everything in S_a is in S_b. That everything in S_b is also in S_a follows in the same way, and so the two are the same set. Each subset in a partition is called a *cell* in that partition.

In the example above, the cells correspond to the different grades. (There don't have to be six cells—it could be that one of the grades has no student in it. But there can be no more than six; recall that by definition a cell can't be empty.) Just as any equivalence relation induces a partition, given any partition one can give an equivalence relation that corresponds to any partition; this is the relation that holds between any two a and b in S such that a and b are in the same cell in the partition.

1.5.4. Functions

A function takes every member of some set A and assigns it a value from a set B (B could be the same set as A, but need not be). This can also be formalized using the notion of a set of ordered pairs. Thus, consider two sets A and B (which again could be the same but need not be). Then, a (total) function from A to B is any set of ordered pairs (i.e., any subset of A x B) such that for each a in A, there is one and only one ordered pair with a as first member. Thus if we think of the function f as assigning to each a in A some values in B, note that the criterion above ensures that each member of A is indeed assigned a value, and is assigned a unique value. A is referred to as the *domain* of the function, and B is referred to as the *co-domain*. For any function f and any a in the domain of f, we write f(a) to indicate the value that f assigns to a. (To use other common terminology, f(a) means the result that one gets by applying the function f to a.) There is no restriction that each member of B must appear as second member of some ordered pair; the term *range* of the function f is the set of all b in B such that there is some a such that f(a) = b. Note that these definitions are such that the range of a function is a subset of the co-domain. In practice (at least in works within linguistics) the terms "range" and "co-domain" are often not distinguished.

As noted above, there is no restriction that each member of B appear as second member of an ordered pair. Nor is there a restriction that it appear only once. If each member of B is used as a value only once (that is, for each b in B, there is a unique a such that f(a) = b) then B obviously can be no smaller than a. It can have more members, or it can be the same size. If the latter is the case, then it also follows that for every b in B, there is some a such that f(a) = b. When both conditions above hold (i.e., for each b in B, there is one and only one a such that f(a) = b, we say that there is a *one-to-one correspondence* between A and B. Note that for any function f which is a one-to-one correspondence, there is a corresponding function f^{-1} which is just the reverse: it is a function mapping each member of B to a member of A such that for all a in A and b in B, if f(a) = b then f^{-1}(b) = a.[10]

We will have some occasion to talk about the notion of a *partial function*. A partial function is one where not every member of A is actually assigned a value by f; f is *undefined* for some subset of A. (Of course any partial function f is also a total function with a smaller domain.) We can illustrate this by returning to our earlier example of ordered pairs of US cities and states, where the first member of each ordered pair is the capital of the second. This is a *partial function* from the set of US cities to states (not every US city is a capital). We can reverse it, and have each state as the first member of the ordered pair and the second as its capital (this function could be expressed in prose as *has as its capital*). This is now a total function from

[10] Incidentally, the notion of the availability of a one-to-one correspondence can be used to define what it means for two sets to have the same cardinality. Obviously for two finite sets it is clear what it means to have the same cardinality, since we can count the members. But consider the case of infinite sets. Take the following two sets: A = the set of positive integers {1,2,3, . . . } and B = the set of positive even integers {2,4,6, . . . }. Both are infinite. Surprisingly (when one first hears this) they are also of the same cardinality, because one can establish a one-to-one correspondence between them (each member of A is paired with a member of B by multiplying by 2: we will never run out of members in B).

the set of states (every state does have a capital) to the set of US cities. But it is not a one-to-one correspondence for the same reason that our original relation is not a total function; there are many cities without the honor of being a capital.

Occasionally in this text it will be useful to list out some actual functions—that is, to name every member in the domain and name what the function at issue maps that member to. There are a variety of ways one could do this. To illustrate, take a domain of four children {Zacky, Yonnie, Shelley, and Baba} (call that set C) and four men {Abe, Bert, Carl, David} (call that set M). Suppose there is a function f from C to M which maps each child to their father. Assume that Abe is the father of Zacky and Yonnie, Bert is the father of Shelley, and David is the father of Baba. Then one can write this information out in various ways. One would be to simply give the set of ordered pairs: {(Zacky, Abe), (Yonnie, Abe), (Shelley, Bert), (Baba, David)}. Usually this notation, however, is not terribly easy to read. We could also write this out in either of the ways shown in (20):

(20)	a. f(Zacky) = Abe		b. Zacky → Abe
	f(Yonnie) = Abe		Yonnie → Abe
	f(Shelley) = Bert		Shelley → Bert
	f(Baba) = David		Baba → David

Or, sometimes it is more convenient to list out the domain on the left and the co-domain on the right and connect them with arrows as in (21):

(21)

$$\begin{tikzcd}
\text{Zacky} \arrow[r] & \text{Abe} \\
\text{Yonnie} \arrow[ur] & \text{Bert} \\
\text{Shelley} \arrow[ur] & \text{Carl} \\
\text{Baba} \arrow[r] & \text{David}
\end{tikzcd}$$

Which notation is chosen makes no difference; the choice should be dictated by clarity.