The Study of Meaning

An Introduction to Semantics and Pragmatics

Nathan Klinedinst

Yasutada Sudo

University College London

Chapter 2

Extensions and Set Theory

The main topic of this chapter is the idea of *extension*. We will start with the distinction between *sense* and *reference*, which is fundamental to semantic theory. After that, we will introduce *Set Theory* as our first formal tool, and show how to use it to analyze the meanings of nouns, verbs and adjectives.

2.1 Sense and reference

If you have read Lewis Caroll's *Through the Looking Glass, and What Alice Found There*, you might remember the scene where Humpty Dumpty asks Alice what her name means. Perplexed at this question, Alice asks him 'Must a name mean something?', to which Humpty Dumpty confidently replies 'Of course it must'.

Humpty Dumpty's question might sound strange, but it is not absurd at all from the perspective of the study of meaning. A proper name, as a grammatical expression in a natural language, must have some meaning. Certainly, a proper name like 'Alice' does not say much about the bearer of this name, perhaps other than that she is likely to be female and English speaking (if it's pronounced as [ælis]), and that she is called Alice, at least by some people. This is in fact the funny thing about Humpty Dumpty's question. He seems to believe that names must be more descriptive and convey some specific information about what kind of people their bearers are, when he says of himself, "my name means the shape I am—and a good handsome shape it is, too. With a name like yours, you might be any shape, almost".

Humpty Dumpty's assumption that names must be descriptive is certainly wrong, but what is the meaning of a proper name, then? This is a hard question that has kep a number of great minds awake at night. In this introductory textbook, we will not be able to provide a full answer to it, or even a proper review of different ideas that have been put forward. Instead, we will zoom in on a particular aspect of the meaning of proper names by adopting an influential guiding principle put forward by the American philosopher David Lewis.

> In order to say what a meaning *is*, we may first ask what a meaning *does*, and then find something that does that. (Lewis 1970: p. 22)

Following Lewis' advice, let us ask: What does a proper name do? One thing it certainly does is to *refer*. Or more strictly speaking, speakers use a proper name to refer to someone or something in particular. Of course, the *referent* of 'Alice' can change depending on the context of use, but every time it is used, there is someone in particular it is intended to refer. It is, therefore, reasonable to assume that reference is one crucial aspect of the meaning of a proper name.

Interestingly, and very importantly for the study of meaning, reference is not all there is to the meaning of proper names. The German philosopher Gottlob Frege (1892) made this point using the following famous pair of sentences. These sentences contain two proper names, 'Hesperus' and 'Phosphorus', both of which are of Greek origin and refer to the planet Venus.

(2.1) a. Hesperus is Hesperus.
b. Hesperus is Phosphorus.

Whenever used, the two proper names refer to one and the same heavenly body. If you assume that their common referent captures everything about their meaning, then these proper names should have completely identical meaning, and if so, the sentences in (2.1) should mean exactly the same thing. Frege points out that this prediction is incorrect, and the sentences in (2.1) mean different things. Specifically, (2.1a) is tautological, while (2.1b) is not. Or to put it differently, everyone knows that (2.1a) is trivially true, but you can imagine someone not knowing that (2.1b) is the case. In fact, Ancient Greeks (unlike Babylonians) didn't know that (2.1b) was true, and that's why they had these two proper names, but of course they all knew that (2.1a) was true.

Thus, the meanings of these two sentences need to be different, and given that the only difference between them is in the second proper name, their difference should be coming from some difference in meaning between the two proper names. In order to make sense of this, Frege proposed to postulate two levels of meaning for proper names. The *referent* (*Bedeutung* in German; sometimes also called *reference*) of a proper name is the entity that it is meant to refer to in a given occasion, whereas its *sense* (*Sinn* in German) is a more abstract aspect of meaning that is used to determine the referent. The two proper names, 'Hesperus' and 'Phosphorus', have the same referent, namely Venus, but different senses. The grammatical construction used in the above examples says that the sense of the first proper name points to the same referent as the sense of the second proper name. This captures the fact that (2.1a) is trivial, because it is trivially true that one and the same sense determines the same referent, while (2.1b) conveys a contingent piece of information.

2.2 Extension and model theory

In order to talk about referent and sense in more concrete terms, formal semanticists often use what is called *model theory*. Model theory was originally developed for stuyding formal languages by logicians such as Alfred Tarski and Rudolf Carnap in the first half of the 20th century, but has had a tremendous impact on the birth and development of *formal semantics*, the formal study of meaning. The pioneers of formal semantics, most notably Richard Montague and David Lewis, adopted model theoretic techniques in analyzing the semantics of natural language expressions. Although this is not the only form of formal semantics, we will follow this tradition in this textbook. Semantic theory that is built on model theory is sometimes called *model-theoretic semantics*.

In model theory, expressions are given *extensions* in some *model*. A model can be seen as a mathematical object that represents some specific state of affairs. For instance, one can construct a model for the situation in the Oval Office at the White House at noon on 15 August 2019. The state of affairs that a model represents does not need to be real. One could make a model for the Oval Office with Hilary Clinton as the president of United States, for example.

A model is formally a mathematical object and has a rigorous definition, but we need not go into that here. All you have to understand at this point is that it represents a particular state of affairs. More specifically, it contains people and objects that one could refer to, which are called *entities* (or *individuals*), and these entities have certain properties such as being the president of the United States, or being a big desk, and stand in certain relations, such as the desk's belonging to the president.

By assumption, a model contains just enough information for determining referents. Therefore, if a model for a natural language is given, it should be able to tell you what a proper name refers to. By way of illustration, let us take an example model, say a model for some (potentially fictional) situation in the Oval Office at noon on 15 August 2019 with Donald Trump and Mike Pence chatting in it. To facilitate the discussion, let's call this model O. In the situation that O is meant to represent, the proper name 'Mike' will refer to Mike Pence, whenever used. We represent this using the idea of *extension*. That is, the model O assigns Mike Pence as the extension of the expression 'Mike'. Semanticists write this as:

$$[\![\text{Mike}]\!]^{O} = \text{Mike Pence}$$

More generally, we write '$[\![\alpha]\!]^{M}$' to mean *the extension of expression α with respect to model M*.

For the study of meaning, individual models are almost never of particular interest. Rather, we want to understand how extensions change, or do not change, across models, as that will give us a way to explain what inferences people draw from linguistic expressions and how people reason based on them. This might not be very clear at this point, but it will be demonstrated with concrete examples in the the rest of this textbook.

Note that one and the same proper name, say 'Mike', can be used to refer to different people in different situations, depending on who is in them. If we are in a situation with, for example, Mike Myers and Dana Carvey, and no one else, then you might refer to Mike Myers by the proper name 'Mike'. We can represent this state of affairs using the above notation as follows. Let W be the model that represents this situation under discussion.

$$[\![\text{Mike}]\!]^{W} = \text{Mike Myers}$$

Note that since any states of affairs, real or not, can be represented by models, there are infinitely many models we can construct. One and the same proper name can have different extensions in different models, corresponding to its referents in different states of affairs. Importantly, the sense of the proper name is always constant and does not depend on the specific state of affairs. Model-theory gives us a way to represent the sense concretely by abstracting over the model parameter in the above representation. That is, if a theory of the sense of the proper name 'Mike' can be stated by filling in the blank of the following statement: 'For any arbitrary model M, $[\![\text{Mike}]\!]^{M}$ is ____'.

As a matter of fact, the semantics of proper names is one of the topics that has continuously been debated throughout the history of the study of meaning, and the literature contains a plethora of different theoretical views and open questions. For instance, it is actually not agreed that 'Mike' in O and 'Mike' in W are the same proper name. Some scholars think that each proper name has the same referent in every model, and 'Mike' referring to Mike Pence and 'Mike' referring to Mike Myers are simply homophonous. Also, different theories offer different ways of dealing with situations containing two or more people named Mike, or situations containing no one called Mike. As the literature is extremely copious in this domain, we will not delve into the question of the senses of proper names any further in this textbook and refer the interested reader to the further readings section of this chapter. All you have to know about proper names in order to follow the rest of the textbook is the distinction between referent and sense and the model-theoretical way of representing extensions, including the above notation.

Note that we could easily extend the model-theoretical analysis to other expressions that can be used to refer to entities. For example, with respect to the above model O, the expression 'the floor', which is a type of expression called a *definite description*, will refer to the floor of the Oval Office. We can represent this as:

$$[\![\text{the floor}]\!]^{O} = \text{the floor of the Oval Office in the White House}$$

You might be wondering how useful model theory is. Indeed, if we are only talking about referring expressions like proper names and definite descriptions, it might look rather uninspiring. It will become more interesting when we consider the extensions of other expressions as well. In the following sections, we will talk about the extensions of nouns, verbs, and adjectives, but in order to do so, we have to introduce an important theoretical tool, *Set Theory*.

2.3 Set Theory

Set Theory was invented at the end of the 19th century and has been used since to give formal foundations to mathematics and logic. Since then it has been widely used in a number of other disciplines, including linguistics, to give rigorous definitions and representations to formal theoretical concepts.

That sounds very scary, but as a matter of fact, the core notions of Set Theory are very intuitive and easy to grasp. Certainly there are some rules you need to learn, but they are all matters of definition, and these definitions are very good to know, because Set Theory is so widespread that you are likely to encounter it in other domains as well. If you have already learned Set Theory, you can skip this section and proceed directly to the next section.

2.3.1 Basics

First of all, Set Theory assumes that there are sets. A *set* is essentially an abstract collection of things. You can take any two things and form a set. For example, take our colleague Klaus Abels and his office at UCL, we can form a set containing them and nothing else. To represent this set, we use curly bracket and a comma as in (2.2).

(2.2) { Klaus Abels, Klaus Abels' office }

Those things in the set are called *members* (or *elements*) of the set. There can be more than two members, as in the following set.

(2.3) { Klaus Abels, Klaus Abels' office, Hyde Park }

Or, there can be fewer than two members, as in (2.4).

(2.4) { Klaus Abels' office }

Sets are defined solely by their members. If two sets contain the exact same members, they are said to be *equivalent*. If m is a member of a set S, we write '$m \in S$'. If not, we write '$m \notin S$'. For example,

$$\text{Hyde Park} \in \{\text{ Klaus Abels, Klaus Abels' office, Hyde Park }\}$$

but

$$\text{Hyde Park} \notin \{\text{ Klaus Abels' office }\}$$

Importantly, the members of a set are not ordered. The only thing that matters is what is a member of the set. This means that one and the same set can be represented in multiple ways. For example, the following representations all denote the same set.

(2.5)
a. { Klaus Abels, Klaus Abels' office, Hyde Park }
b. { Klaus Abels' office, Klaus Abels, Hyde Park }
c. { Hyde Park, Klaus Abels' office, Hyde Park, Klaus Abels, Klaus Abels }

The last set might look like it has more members, but that is not the case. There are exactly three distinct members, as in the other sets. Needless to say, it's not advisable to use a representation like (2.5c), unless there is a particular reason to do so, because it's very confusing.

A set does not need to have a member. Such a set is called an *empty set*. In fact, there is only one empty set, because we define sets in terms of their members, and there can be only one set with no members. The empty set can be represented as '{ }', but it is more common to use the special symbol '$\varnothing$'. Note that the empty set is a set and the symbol '$\varnothing$' doesn't mean 'nothing' (although that is how this symbol is sometimes used in other subareas of linguistics). $\varnothing$ is a set.

A set can contain another set, including the empty set, as a member, as in the following example.

(2.6) { Klaus Abels, { Hyde Park, Brazil } }

Note that Hyde Park is *not* a member of this set, although the set { Hyde Park, Brazil } is. Set membership ($\in$) is only about direct members and does not apply to members of members. So the following sets are all distinct from (2.6) and from each other, because they have different members.

(2.7)
a. { Klaus Abels, Hyde Park, Brazil }
b. { { Klaus Abels } , Hyde Park, Brazil }
c. { { Klaus Abels, Hyde Park } , Brazil }
d. { { Klaus Abels, Hyde Park } , { Brazil } }
e. { Klaus Abels, Hyde Park, Brazil, $\varnothing$ }

Similarly, $\varnothing$ and { $\varnothing$ } are distinct sets. $\varnothing$ has no member, while { $\varnothing$ } has one member, namely, $\varnothing$. Again, remember that $\varnothing$ doesn't mean 'nothing'. It's a set with no members, so there *is* a member in { $\varnothing$ }.

It is often convenient to talk about how many distinct members there are in a given set. The number of distinct members of a given set S is called the *cardinality* of S, and represented by $|S|$. For example:

(2.8)
a. | { Klaus Abels, Klaus Abels' office } | $= 2$
b. | { Klaus Abels, { Hyde Park, Brazil } } | $= 2$
c. | { Klaus Abels, Hyde Park, Brazil } | $= 3$
d. | { Klaus Abels, Hyde Park, Brazil, $\varnothing$ } | $= 4$
e. $|\varnothing| = 0$
f. $|\{\varnothing\}| = 1$

The idea of cardinality is given a rigorous definition, and becomes very interesting for sets with infinite members. If you are interested in this, please check the further readings section at this end of this chapter.

2.3.2 Enumeration vs. abstraction

There are generally two ways of representing sets. One is by *enumerating* the members and putting a pair of curly brackets around them, as we have been doing. But this way of representing sets gets cluttered when there are many members, and also cannot represent sets that have infinitely many members. Also, there are situations where you do not know exactly what is in the set, but still want to talk about it, e.g. the set of all prime numbers smaller than $2^{10^{10}}$.[1]

For these purposes, there is a way to represent a set with a *variable* and a membership condition. For example, the set of all prime numbers smaller than $2^{10^{10}}$ is represented as:

(2.9) $\{\, x \mid x$ is a prime number smaller than $2^{10^{10}} \,\}$

Or the set of all restaurants in London will be:

(2.10) $\{\, x \mid x$ is a restaurant in London $\}$

These are obviously finite sets, but you might not be able to completely identify their members. If so you cannot name these sets by enumerating their members. Moreover, we can similarly represent infinite sets, such as the following two.

(2.11) a. $\{\, x \mid x$ is an even number $\}$
b. $\{\, x \mid x$ is a real number between 0 and $1 \,\}$

In this notation, a set generally looks like '$\{\, \xi \mid \phi \,\}$', where 'ξ' is a variable and 'ϕ' is a statement.[2] Normally, we use 'x', 'y' and 'z' for variables. Keep in mind that what is on the right of '|' always needs to be a complete statement that is either true or false. The way to understand this notation is the following. '$\{\, \xi \mid \phi \,\}$' is the set containing each and every thing that makes 'ϕ' true when 'ξ' refers to it, and nothing else. Let us go through an example to get used to it.

According to the rule, (2.10) represents the set containing each and every thing that will make the statement 'x is a restaurant in London', true when x refers to it, and the set contains nothing else. That is to say, this is the set of all restaurants in London. Take Nobu in Shoreditch, London, for example. It's a restaurant in London, so when x refers to it, the statement 'x is a restaurant in London' is true. This means that Nobu is in this set. Similarly, if x refers to Ottolenghi in Spittalfields, London, the statement 'x is a restaurant in London' is true, so Ottolenghi is in the set, as well. On the other hand, take the main building of UCL. It is not a restaurant, so when x refers to it, 'x is a restaurant in London' is false. This means that the main building of UCL is not in this set.

[1] The largest prime number known as of December 2018 is $2^{82589933} - 1$.

[2] Some authors use ':', instead of '|', as in '$\{\, \xi : \phi \,\}$'. Also, it is standardly considered that not every statement can be used to define a set. We'll come back to this point later in Section 2.3.6.

Similarly, if x refers to the set containing Nobu in Shoreditch and Ottolenghi in Spittalfields, and nothing else, then the statement 'x is a restaurant in London' will be false, because x is a set, and not a restaurant in London (although it is a set of restaurants in London). Therefore, this set is not in the set in (2.10).

Strictly speaking, the condition ϕ does not need to contain the variable ξ. In that case, the set is going to contain either everything or nothing, depending on whether ϕ is true or false. For example,

$$\{\, x \mid \text{London is the capital of England} \,\}$$

is the set of everything, because the statement 'London is the capital of England' is true, and every value of x will make this statement trivially true, for these values don't change the statement. On the other hand, the set

$$\{\, x \mid \text{Moscow is the capital of England} \,\}$$

is the empty set $\varnothing$, because the statement 'Moscow is the capital of England' is false, as no value of x can make this statement true.

There are several caveats, especially regarding the use of variables in this notation. Firstly, ξ in $\{\, \xi \mid \phi \,\}$ is a variable, and it itself is not necessarily a member of the set. For example, in the case of (2.10), the variable x itself is not a member of this set, because a variable is not a restaurant. Secondly, notice that the natural language paraphrase of this set is 'the set of all restaurants in London', and does not contain x. In fact, we could state the same set using different variables as in (2.12).

(2.12) a. $\{\, y \mid y \text{ is a restaurant in London} \,\}$
b. $\{\, z \mid z \text{ is a restaurant in London} \,\}$

This is because we are only interested in the values of these variables, and not the variables themselves.[3] Thirdly, we can represent a finite set in this notation as well. For example, the following set is equivalent to the set in (2.2).

(2.13) $\{\, x \mid x \text{ is Klaus Abels or } x \text{ is Klaus Abels' his office} \,\}$

Note the use of 'or' here. If 'and' were used, as in (2.14) it would be a different set. Specifically, since nothing is both Klaus Abels and his office at the same time, (2.14) is $\varnothing$.

(2.14) $\{\, x \mid x \text{ is Klaus Abels and } x \text{ is Klaus Abels' his office} \,\}$

You might find the use of variables a bit difficult at first, but understanding

[3]Note that if we are talking about a set containing a variable as a member, then the variable might be in it, e.g. $\{\, x \mid x \text{ is a variable used in this textbook} \,\}$ contains the variable x. But in this case too, one can name the same set without using x, as in $\{\, y \mid y \text{ is a variable used in this textbook} \,\}$

this notation itself should be easy enough. In the rest of this textbook, we will be mostly talking about simple sets like the ones in (2.15).

(2.15) a. $\{ x \mid x \text{ is a cat} \}$
(the set of all cats)

b. $\{ x \mid x \text{ is a black cat} \}$
(the set of all black cats)

c. $\{ y \mid y \text{ is a person and } x \text{ danced but did not sing} \}$
(the set of all people who danced but did not sing)

That being said, Set Theory, as well as the idea of variables, is a standard tool in a number of different disciplines and it's worth mastering everything mentioned in this section. Keep in mind that Set Theory is a formal language, and is basically a second language for everyone, and it's normal to struggle at the beginning. In fact, you can only become fluent in it through exposure and practice, just like any other second language. You can find some simple exercises at the end of this chapter to familiarize yourself with this language.

2.3.3 Operations on sets

Now we know how to define individual sets: We either enumerate the members, or use a variable and state the membership condition. We can furthermore create sets out of two sets using certain operations. The ones below are particularly important for the study of meaning. Let A and B any two arbitrary sets.

(2.16) a. The *intersection* of A and B, written '$A \cap B$', is $\{ x \mid x \in A \text{ and } x \in B \}$.

b. The *union* of A and B, written '$A \cup B$', is $\{ x \mid x \in A \text{ or } x \in B \}$.

c. The *relative complement* of B with respect to A, written '$A - B$' (or '$A \backslash B$'), $\{ x \mid x \in A \text{ and } x \notin B \}$.

$A \cap B$ is the set containing all the common members of A and B (and nothing else). $A \cup B$ is the set containing all members of A and all members of B (and nothing else). $A - B$ is just like A, except that those members that are also in B are missing.

It helps to visualize these operations using *Venn diagrams* as in Figure 2.1. In each diagram here, the two circles labeled A and B represent arbitrary sets A and B. If something falls in the circle, it means that it is a member of the set, and if something is outside the circle, it is not a member of the set. The shaded parts of the diagrams represent $A \cap B$, $A \cup B$ and $A - B$, respectively. In these diagrams, A and B are depicted as overlapping sets, but in the general case, this is not necessary. In such a case, we say A and B are *disjoint.* If A and B are disjoint, then $A \cap B = \varnothing$, and $A - B = A$.

More operations than these three can be defined, but they are not very common. We will discuss one such relatively rare operation in an exercise at the end

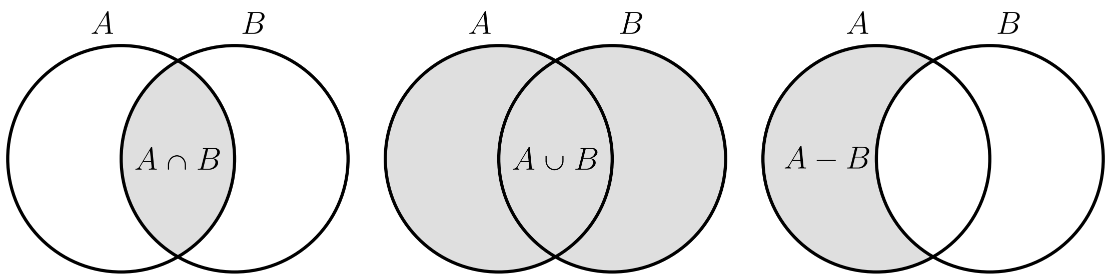

Figure 2.1: Venn diagrams representing $A \cap B$ (left), $A \cup B$ (middle), and $A - B$ (right).

of this chapter.

2.3.4 Relations between sets

We also often talk about certain relations between sets. The most important one is the *subset* relation, defined as (2.17). Again, A and B are any two arbitrary sets.

(2.17) A is a *subset* of B, written $A \subseteq B$, iff every member of A is also a member of B.

For any set A, $A \subseteq A$, because every member of A is obviously a member of A. The little horizontal bar underneath signifies that a subset of a set can be the same set. Note that perhaps this is not the way we normally use the word 'subset' in everyday speech; it tends to mean a smaller set (i.e., a proper subset; see below). In the context of Set Theory, every set is a subset of itself.

Another thing you need to remember is that $\varnothing$ is a subset of every set, including $\varnothing$ itself. That is, we regard the statement 'every member of $\varnothing$ is a member of B' to be vacuously true, when there is no member to talk about. This is a matter of convention, and you need to remember it.

We can state equivalence between two sets in terms of $\subseteq$. Recall that two sets are equivalent iff they contain exactly the same members. The same thing can also be stated as follows, for any two sets A and B.

(2.18) A and B are equivalent, written $A = B$, iff $A \subseteq B$ and $B \subseteq A$.

When two sets are not equivalent, we write $A \neq B$.

When $A \subseteq B$ and $A \neq B$, we say A is a *proper subset* of B and write $A \subset B$.[4] Obviously every proper subset of A is a subset of A, for any set A.

For any set A, the set of all subsets of A is called its *power set*, $\wp(A)$ (or Pow(A)).

[4]Confusingly, some authors use '$\subset$' to mean what we mean by '$\subseteq$' and '$\subsetneq$' to mean what we mean by $\subset$. This alternative notation is not very widely used in linguistics today, but you might encounter it in old papers.

For example, $\wp(\{a, b\}) = \{\varnothing, \{a\}, \{b\}, \{a, b\}\}$. Note that $\varnothing$ is always in a power set, because it is a subset of every set. Note also that each member of a power set is a set itself. Generally, if there are n distinct members in a set A, i.e. $|S| = n$, $\wp(A)$ has 2^n members, i.e. $|\wp(A)| = 2^n$. This is because $\wp(A)$ contains all sets consisting of zero or more members of A, and there are 2^n ways to make such sets (i.e. each member of A is either a member or not).

2.3.5 Ordered pairs and n-tuples

Recall that the members of a set are not ordered. So $\{a, b\}$ and $\{b, a\}$ represent the same set. However, it is sometimes convenient to be able to talk about ordered versions of sets. When there are two members to be ordered, we use an *ordered pair*, which look like $\langle a, b\rangle$ with angled brackets, instead of curly brackets. Crucially, $\langle a, b\rangle \neq \langle b, a\rangle$, whenever a and b are distinct (they are identical if a and b are the same thing).

An ordered-pair can be understood as a special kind of set. There are many different ways of 'modeling' ordered pairs as sets, but Kazimierz Kuratowski's version is most widely used. According to it, an ordered pair $\langle a, b\rangle$ is defined as the set $\{\{a\}, \{a, b\}\}$. Then $\langle b, a\rangle$ will be $\{\{b\}, \{b, a\}\}$, and indeed, $\langle a, b\rangle \neq \langle b, a\rangle$, unless $a = b$. Recall that $\{a, a\} = \{a\}$, because sets are only defined in terms of its members. Therefore, when $a = b$, $\{\{a\}, \{a, b\}\} = \{\{a\}, \{a\}\} = \{\{a\}\}$. By the same reasoning, $\{\{b\}, \{b, a\}\}$ will be the same set in that case.

We can use ordered pairs to define ordered triples. One way to define it is $\langle a, b, c\rangle \coloneqq \langle a, \langle b, c\rangle\rangle$. We could as well define it as $\langle\langle a, b\rangle, c\rangle$ or $\langle\langle a, c\rangle, b\rangle$, but their differences are inconsequential, and we just need to make an arbitrary choice. Note that $\langle a, \langle b, c\rangle\rangle$ is an ordered pair, containing an ordered pair as one of its members. If we look at an ordered pair as a set, it is natural to expect that ordered pairs can contain ordered pairs, because sets can contains sets.

In the same way, we can define an ordered sequence of n members, or n-*tuples*, $\langle a_1, a_2, \ldots, a_n\rangle$ as $\langle \cdots \langle a_1, \langle a_2, \langle \cdots, a_n\rangle \cdots\rangle\rangle \cdots\rangle$. For example, the following is an ordered 5-tuple (or quintuple): $\langle \text{London}, \text{Paris}, \text{Berlin}, \text{Barcelona}, \text{Milan}\rangle$.

2.3.6 Russell's paradox

To conclude our short introduction to Set Theory, we would like mention an interesting fact that not all expressions of the form '$\{\xi \mid \phi\}$' is a set, as pointed out by the British philosopher Bertrand Russell in 1901. If any variable ξ and any statement ϕ can be used to form a set, then the following will be a set.

(2.19) $\{x \mid x \text{ is a set and } x \notin x\}$

This is meant to be the set of all sets that do not contain themselves. Let's call it R (for Russell). Russell discovered that R is not a well-formed set. Why? Because it gives rise to a paradox. The paradox arises when we ask if R is a member of itself. Recall that for a given set, everything in the universe of Set Theory is either a member of it or not a member of it. So we should be able to ask whether

$R \in R$ or $R \notin R$. Suppose that $R \in R$. If $R \in R$, then by the membership condition of this set, it must be the case that $R \notin R$. But this is a contradiction. So it cannot be the case that $R \in R$. Then maybe we have $R \notin R$? Well, if $R \notin R$, then R meets the membership condition of this set, so it must be the case that $R \in R$. This is again a contradiction. So R is neither a member or not a member of itself, which is a contradiction in Set Theory.

This led set theoreticians to abandon the idea that any statement can be used to define a set. Several solutions were proposed but what came to be known as Zermelo-Fraenkel Set Theory (with the axiom of choice), often referred to as ZFC, is particularly well known and studied, and is standardly taken to be the foundation of mathematics. But we will not delve into this very technical but interesting literature here.

2.4 Nouns

2.4.1 Noun extensions as sets

Set Theory allows us to talk about the extensions of nouns in formal terms. For example, take a simple noun, 'laptop'. What is its extension? Let us recall David Lewis' slogan from the beginning of this chapter. We should first figure out what the noun 'laptop' does and then find something that does it. What does it do? Intuitively, it classifies entities into laptops and non-laptops, and singles out the former.

To illustrate, let us look at an example situation with exactly three laptops, exactly two desktops, and exactly one cat. Call these laptops l_1, l_2 and l_3, the desktops d_1 and d_2, and the cat c (These labels or names in our metalanguage for the objects are not important; They can be anything). In this situation, the noun 'laptop' singles out the three laptops out of the six entities.

We can represent this analysis in model-theoretic terms as follows. Let us assume that the model L represents this state of affairs with five computers and a cat. With respect to this model, the noun 'laptop' is given the set of all three laptops as its extension. Using the same notation as before, we can write this as (2.20).

(2.20) $[\![\text{laptop}]\!]^L = \{ l_1, l_2, l_3 \}$

Thus, the extension of 'laptop' is the set of entities in the model that are laptops. Similarly, we can write the extension of 'computer' in the same model as (2.21).

(2.21) $[\![\text{computer}]\!]^L = \{ l_1, l_2, l_3, d_1, d_2 \}$

You can similarly represent the extensions of 'desktop' and 'cat'. In this model, there are no dogs or unicorns, so the extension of 'dog' and 'unicorn' will be empty.

(2.22) $[\![\text{dog}]\!]^L = [\![\text{unicorn}]\!]^L = \varnothing$

As in the case of proper names, the extensions are dependent on the model. With respect to a different model, say I, where there is only one laptop k and no other computers and where a unicorn u is browsing Instagram using k, we have the following extensions.

(2.23) a. $[\![\text{laptop}]\!]^I = \{\, k \,\}$
b. $[\![\text{computer}]\!]^I = \{\, k \,\}$
c. $[\![\text{unicorn}]\!]^I = \{\, u \,\}$

Note that the extensions of 'laptop' and 'computer' are identical in this model I, but this doesn't mean that their meanings are identical in all respects. This is parallel to the example of Hesperus and Phosphorus. They may have the same extension in some model, but they don't need to in every model, which is to say, they have different senses.

These analyses can be stated in general terms with respect to all models as follows.

(2.24) For any model M,
a. $[\![\text{laptop}]\!]^M = \{\, x \mid x \text{ is a laptop in } M \,\}$
b. $[\![\text{computer}]\!]^M = \{\, x \mid x \text{ is a computer in } M \,\}$
c. $[\![\text{unicorn}]\!]^M = \{\, x \mid x \text{ is a unicorn in } M \,\}$

Note that these equations look trivial, but this triviality is only an apparent one that arises from our use of English as both the object language and metalanguage. Ultimately we want to understand how exactly these extensions are determined, because after all, there are infinitely many possible models, so we cannot know the senses of these nouns by simply listing their extensions with respect to different models. To answer this, we believe we need a better understanding of the concepts encoded by these nouns, which we will not delve into in this textbook for reasons mentioned in the previous chapter. However, there are many interesting things we can do without a deep analysis of individual content words like these, as we will see for the rest of this textbook.

It should be noted that when a noun 'N_1' is a hyponym of a noun 'N_2', then for any model M, we always have the following subset relation:

$$[\![N_1]\!]^M \subseteq [\![N_2]\!]^M$$

For example, this is the case for 'laptop' and 'computer'. Every laptop is always a computer, so $[\![\text{laptop}]\!]^M$ is always a subset of $[\![\text{computer}]\!]^M$ with respect to any model M. Note that the subset relation stays true even if there is no laptop or there is no computer in some models, because $\varnothing$ is a subset of every set, including $\varnothing$ itself.

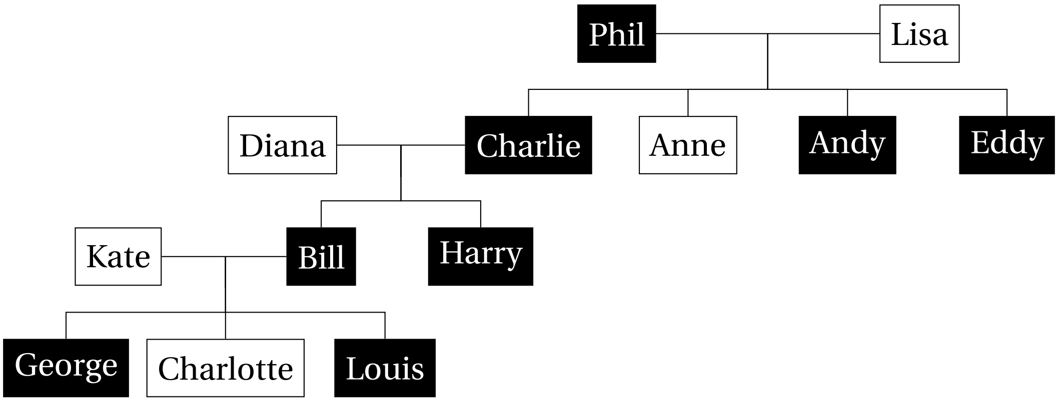

Figure 2.2: A hypothetical family tree. The white boxes are women and the black boxes are men.

2.4.2 Relational nouns

One interesting fact about languages like English is that there are nouns that seem to express *relational concepts*. For example, consider the noun 'parody'. This noun expresses a relational concept, because whenever something is a parody, it's always a parody *of something*. Other such relational nouns include 'sequel', 'part', and kinship terms like 'mother' and 'brother'. We can model the extensions of relational nouns in terms of sets as well. This time, however, they will not be sets of entities, but the sets of ordered pairs of entities.

To illustrate, let us consider a concrete model, for example, the family tree in Figure 2.2. We will only consider people mentioned here. There are several relations that hold among these people. For example, consider the mother-of relation. Lisa is the mother of Charlie, Anne, Andy and Eddy, Diana is the mother of Bill and Harry, and Kate is the mother of George, Charlotte, and Louis. There is not other mother-of relation here. We can represent this state of affairs as the following set of ordered pairs.

(2.25) $$\left\{ \begin{array}{l} \langle \text{Lisa, Charlie} \rangle, \langle \text{Lisa, Anne} \rangle, \langle \text{Lisa, Eddy} \rangle, \\ \langle \text{Diana, Bill} \rangle, \langle \text{Diana, Harry} \rangle, \\ \langle \text{Kate, George} \rangle, \langle \text{Kate, Charlotte} \rangle, \langle \text{Kate, Louis} \rangle \end{array} \right\}$$

This set contains eight ordered pairs. In each of these ordered pairs, the first member is the mother of the second member. We can represent the same set as (2.26).

(2.26) $\{ \langle x, y \rangle \mid x \text{ is the mother of } y \text{ in Figure 2.2} \}$

We can regard this set as the extension of the noun 'mother' in the model representing the situation depicted in Figure 2.2. That is, calling this model F, we have:

(2.27) $$[\![\text{mother}]\!]^F = \left\{ \begin{array}{l} \langle \text{Lisa, Charlie} \rangle, \langle \text{Lisa, Anne} \rangle, \langle \text{Lisa, Eddy} \rangle, \\ \langle \text{Diana, Bill} \rangle, \langle \text{Diana, Harry} \rangle, \\ \langle \text{Kate, George} \rangle, \langle \text{Kate, Charlotte} \rangle, \langle \text{Kate, Louis} \rangle \end{array} \right\}$$

Note that we could have used the set in (2.28) below instead as the extension of 'mother'. This set is exactly like the set of ordered pairs above except that the members of each ordered pair are reversed.

(2.28) $$\{ \langle y, x \rangle \mid x \text{ is the mother of } y \text{ in Figure 2.2} \}$$
$$= \left\{ \begin{array}{l} \langle \text{Charlie, Lisa} \rangle, \langle \text{Anne, Lisa} \rangle, \langle \text{Eddy, Lisa} \rangle, \\ \langle \text{Bill, Diana} \rangle, \langle \text{Harry, Diana} \rangle, \\ \langle \text{George, Kate} \rangle, \langle \text{Charlotte, Kate} \rangle, \langle \text{Louis, Kate} \rangle \end{array} \right\}$$

These two sets of ordered pairs represent the same relation, namely the mother-of relation, and contain the exact same amount of information, as each of them can be uniquely reconstructed from the other by reversing the order between the members in each pair. This means that, in a way, these two sets represent 'the same thing', although they are still formally distinct sets. For our purposes of constructing an analysis of the extensions of relational nouns, we simply need to make an arbitrary choice between the two. We will use the former, as the order matches the order in the English description like the condition in (2.26), but there is no theoretically principled reason for this choice.

Generalizing the above analysis to arbitrary models, we can state the extension of 'mother' in any given model M as follows.

(2.29) $$[\![\text{mother}]\!]^M = \{ \langle x, y \rangle \mid x \text{ is the mother of } y \text{ in } M \}$$

It is easy to extend this analysis to any other relational nouns, as in (2.30).

(2.30) a. $[\![\text{brother}]\!]^M = \{ \langle x, y \rangle \mid x \text{ is a brother of } y \text{ in } M \}$

b. $[\![\text{sequel}]\!]^M = \{ \langle x, y \rangle \mid x \text{ is a sequel to } y \text{ in } M \}$

To get familiar with this idea, it will be a good exercise to actually write down the members of the extensions of these nouns in an example model. You can find some exercises of this kind at the end of this chapter.

2.5 Verbs

As we will demonstrate now, we can extend the above analysis of nouns to verbs without further ado. However, you might think that the meanings of verbs are fundamentally different from meanings of nouns, and so we might not actually want to pursue such an analysis of verbs that treats nouns and verbs on a par. Naively put, one might think that verbs are about actions, while nouns are about objects and creatures. If so, how could we represent the extensions of verbs like we did for the extensions of nouns as sets of entities or sets of order pairs of entities?

It should be noticed that while these intuitions are probably not incorrect as statistical tendencies, they are not absolutely correct. Firstly, nouns can be about actions too. For example, what is the difference between a noun like 'destruction' and a verb like 'destroy'? They are very similar in meaning, if not completely identical. Secondly, not all verbs are about actions in the intuitive sense. For instance, 'border' can be used a verb in (2.31), but intuitively does not describe an action. It is more about a *state* Hungary is in.

(2.31) Hungary borders seven countries.

There is a way to understand verbal meaning on a par with nominal meaning in terms of sets of entities or pairs of entities, while maintaining the intuition that verbs tend to be about actions or states. Specifically, we can regard the extension of the verb 'fly', for example, in a given model to be the set of entities that perform an action that counts as flying. Similarly, the extension of a verb like 'border' can be seen as the set of pairs of entities for which there is a state of bordering each other.

To see this analysis more clearly, let us consider an example model, call it T. In the situation represented by T, there are three entities, Alice, Tweedledum, and Tweedledee. Tweedledee and Alice each spoke, but Tweedledum did not. Tweedledee and Tweedledum resemble each other, but Alice does not look like either of them. Then, the extensions of the verbs 'spoke' and 'resembles' with respect to this model will be as follows.

$$\begin{aligned} [\![\text{spoke}]\!]^T &= \{\,\text{Alice}, \text{Tweedledee}\,\} \\ [\![\text{resembles}]\!]^T &= \{\,\langle\text{Tweedledee}, \text{Tweedledum}\rangle\,, \langle\text{Tweedledum}, \text{Tweedledee}\rangle\,\} \end{aligned}$$

Note that in the extension of 'resemble', we have two pairs whose members are the same but whose order is different. Generally, given the meaning of this verb, whenever $\langle a, b\rangle$ is in its extension, it must be the case that $\langle b, a\rangle$ is also in its extension. In other words, if a resembles b it must be the case that b resembles a as well. Verbs like this are called *symmetric* verbs.

Perhaps unsurprisingly, not all verbs are symmetric. For example, 'outperform' is not symmetric, because if a outperforms b, then it is not the case that b outperforms a. In fact, in such a case that we can be sure that b does not outperform a. Such verbs are called *asymmetric* verbs. Similarly 'like' is not symmetric. If a likes b, it is not necessarily the case that b likes a. Notice however that 'like' not asymmetric either, because it is possible that a likes b and b also likes a.

(A)symmetry is applicable to relational nouns as well. For instance, 'brother' is symmetric, because whenever a is a brother of b's, it is necessarily the case that that b is a brother of a's. On the other hand, 'father' is asymmetric because whenever a is b's father, b cannot be a's father.[5]

[5]This is true for biological father but not necessarily for step-father. For instance, if I married

We can define these semantic properties in model-theoretic terms as follows. Let α be an expression whose extension is a set of ordered pairs of entities. Remember that by convention distinct variables like x and y may take the same entity as their values.

(2.32) a. α is *symmetric* iff the following holds: For each model M and for any entities x and y in M, if $\langle x, y\rangle \in [\![\alpha]\!]^M$, then $\langle y, x\rangle \in [\![\alpha]\!]^M$.
b. α is *asymmetric* iff the following holds: For each model M and for any entities x and y in M, if $\langle x, y\rangle \in [\![\alpha]\!]^M$, then $\langle y, x\rangle \notin [\![\alpha]\!]^M$.

Expressions like 'like' that are neither symmetric or asymmetric are called *non-symmetric*. There is another class of expressions that are called *anti-symmetric*, but we will reserve it for an exercise for this chapter.

Importantly, these definitions require their respective conditional statements to hold with respect to *all* models. It is not enough to satisfy them in one model. That is, there can be a model where each pair $\langle x, y\rangle$ in the extension of 'like' has its inverse $\langle y, x\rangle$ in it as well, but that doesn't make this verb symmetric, as there are also models where this is not the case.

There are a number of other similar properties of such relational expressions. *Reflexivity* is a particularly well discussed concept in linguistics, due to its relevance in certain linguistic phenomena. We can define it and its converse, *irreflexivity*, in terms of extensions as follows. Let α be an expression whose extension is a set of pairs of entities.

(2.33) a. α is *reflexive* iff for each model M, and for any entity x in M, $\langle x, x\rangle \in [\![\alpha]\!]^M$.
b. α is *irreflexive* iff for each model M, and for any entity x in M, $\langle x, x\rangle \notin [\![\alpha]\!]^M$.

A relational expression that is neither reflexive nor irreflexive is said to be *non-reflexive*. For example, 'subset' in the technical sense of the term is reflexive, as every set is a subset of itself. On the other hand, 'mother' (in the biological sense) is irreflexive, and *supporter* is non-reflexive. Among verbs, 'divides' in the mathematical sense is reflexive, as every number divides itself. On the other hand, 'differ' is irreflexive, and 'love' is non-reflexive.

Another property often mentioned in linguistics is *transitivity*. This property can be defined in model theory as follows. Here again, x, y and z can be pairwise identical, i.e. any two of them or all of them can be identical.

(2.34) α is *transitive* iff the following holds: For each model M, and for any entities x, y, and z in M, if $\langle x, y\rangle \in [\![\alpha]\!]^M$ and $\langle y, z\rangle \in [\![\alpha]\!]^M$, then $\langle x, z\rangle \in [\![\alpha]\!]^M$.

someone with a grown-up child, and if my biological father married this child, then I would be my father's step-father. Then my biological father and I would be each other's father!

For instance, 'brother' is a transitive noun, 'father' is not. Similarly, among verbs, 'outperform' is transitive, while 'like' is not.

We can define more properties, but they are not used as often in linguistic contexts as the above ones. We will discuss some of them in exercises at the end of this chapter.

2.6 Argument structure

Verbs like 'walk' and 'smile' normally appear with a subject. For instance in 'Alice walked', Alice is the subject. This sentence is about an event of walking and Alice is described as the walker. Such an entity that is related to the event or state being described by a verb is called an *argument* of the verb. Verbs like 'walk' and 'smile' that have only one argument are called *intransitive verbs*. On the other hand, verbs like 'touch' and 'resemble' usually have two arguments expressed by the subject and the object, as in 'The little girl touched the jar'. Such verbs with two arguments are called *transitive verbs*. In our analysis above, the extension of an intransitive verb is a set of entities, while the extension of a transitive verb is a set of pairs of entities.

Confusingly, this notion of transitivity is distinct from the one we mentioned above and defined in (2.34). These two uses of the term co-exist in linguistics, so it is important to pay attention to which use is intended.

There are also verbs that have three arguments, such as *give* and *introduce*. These verbs are called *ditransitive* verbs. The extensions of ditransitive verbs can be understood as sets of triples of entities. For instance, let us consider a concrete situation where Alice introduced herself to Bob, Bob introduced himself and Cathy to Alice, and there was no other introduction. With respect to the model W that represents this state of affairs, 'introduced' will have the following extension.

$$[\![\text{introduced}]\!]^W = \left\{ \begin{array}{l} \langle \text{Alice}, \text{Alice}, \text{Bob} \rangle, \\ \langle \text{Bob}, \text{Bob}, \text{Alice} \rangle, \\ \langle \text{Bob}, \text{Cathy}, \text{Alice} \rangle \end{array} \right\}$$

2.6.1 Alternations

Once one starts looking at actual verbs, it becomes immediately clear that the number of arguments is sometimes not uniquely determined for some verbs. For example, take the verb 'speak'. This verb can appear with just a subject, or a subject and object.

(2.35) a. Lewis spoke.
b. Lewis spoke German.

One could understand this as a case of ambiguity. Given that the two occurrences of the verb in (2.35) express very similar ideas, this is probably a case of polysemy, rather than a case of accidental ambiguity like the two senses of 'bank'. In fact, in many languages, verbs that correspond to 'speak' exhibit similar intransitive-transitive ambiguity, or *alternation*.

Transitive-intransitive alternation is not specific to 'speak'. In fact, there are a plethora of such verbs in English, e.g. 'eat', 'study', 'sing', 'blink', 'pour', 'break', 'open', among many others.

(2.36) a. The boy ate (the tasty food).
b. The girl studied (mathematics) at MIT.
c. The professor sang (a beautiful song).

In English, even verbs like 'die' can be used transitively, as in 'The soldier died a painful death', and 'kill' can be used intransitively, as in 'Cigarettes kill'. If these cases of transitive-intransitive alternation are all cases of polysemy, we want to understand what this alternation really consists in. Importantly, there are intransitive verbs that do not have transitive uses, for example, 'disappear' and 'exist', and there are also verbs that are always transitive, for example, 'desire' and 'devour', so an analysis of the intransitive-transitive alternation needs to also explain why it does not apply to all verbs.

To make the matter more complicated, it turns out that not all cases of intransitive-transitive alternation are the same process. To see this, compare (2.36) above with the following examples.

(2.37) a. My laptop broke.
b. My sister broke my laptop.

(2.38) a. The door opened.
b. I opened the door.

In these examples, the subject of the intransitive use becomes the object of the transitive use, unlike in (2.36), where the intransitive version obtains by simply omitting the object of the transitive version. Notice also that the meaning expressed by the transitive sentences, (2.37b) and (2.38b), are express the idea of *causation*. That is, (2.37b) essentially means that my sister did something and that caused my laptop to break, and (2.38b) means that I caused the door to open. This type of alternation is called *inchoative-causative* alternation.

2.6.2 How arguments are realized

In addition to the number of arguments, verbs also show variation as to how their arguments are realized. In all examples we have seen so far, the arguments are noun phrases, and that is how intransitive vs. transitive verbs are normally defined. However semantically speaking, we also want to include arguments expressed in other ways, most notably, prepositional phrases. In order to see this, consider the following examples.

(2.39) a. Hans heard the music.

b. Hans listened to the music.

(2.40) a. Adrian saw the bird.
b. Adrian looked at the bird.

Clearly, 'hear' and 'see' have transitive uses, but do we want to say that 'listen' and 'look' are intransitive, just because their non-subject argument is expressed by a prepositional phrase? Semantically, 'hear' and 'listen' seem to express related meanings, and if the former can be understood as a set of pairs of entities, then the latter can be as well. Similarly for 'see' and 'look'.

Generally, *how* the arguments are expressed is idiosyncratic, at least to an extent. There is really no semantic reason why the 'object' of 'listen' is a 'to'-phrase, rather than a noun phrase. In fact, in Italian, the verb that corresponds to 'listen'—namely, 'ascoltare'—appears with a subject and object, just like the counterpart of 'hear'—'sentire'. For semanticists, this superficial difference between 'listen' and 'ascoltare' is of no more interest than their phonological difference, and if the extension of the latter is to be understood in terms of a set of pairs of entities, then the extension of the former as well.

A similar point can be made with ditransitive verbs like 'give'. This verb can appear in two superficially different constructions.

(2.41) a. The teacher gave the pupil a book.
b. The teacher gave a book to the pupil.

The construction in (2.41a) is often called the *double object construction*, because it looks like it has two objects. On the other hand, the indirect object argument in (2.41b) is expressed as a prepositional phrase. It seems wrong to analyze 'give' in (2.41a) as ditransitive and 'give' in (2.41b) as transitive, because the meaning expressed by these sentences is essentially identical. Then, again, this suggests that the 'to'-phrase in (2.41b) should be seen as an argument.

However, including prepositional phrases opens a can of worms, because it becomes difficult to decide where to stop. For instance, is 'walk' intransitive in (2.42a) but transitive in (2.42b) and (2.42c)?

(2.42) a. The boy walked.
b. The boy walked the dog.
c. The boy walked to the park.

It seems to make sense to say that 'walk' in (2.42b) is a transitive verb, but it is less intuitive to call 'walk' in (2.42c) a transitive verb. How do we determine which phrase is an argument of a verb what is not? It is sometimes mentioned that whether the phrase is somehow essential to the meaning of the verb is one criterion. For instance, for the transitive use of 'walk' in (2.42b), the object is an essential part, because when someone walks something, there must be two entities, where as for the use of 'walk' in (2.42c) expresses the same meaning as in (2.42a), and the only essential part of this meaning is the walker, as there might or might not be a destination. However, this is a very loose criterion to be usable in the general case. To see this, consider the following example.

(2.43) a. The chef sliced the ham.
b. The chef sliced the ham with a sharp knife.

No one would disagree that 'slice' in (2.43a) is a transitive verb, but is 'slice' in (2.43b) a ditransitive verb, because there is an instrument argument? Notice that whenever someone slices something, there must be an instrument (which could be their hands). So in this sense, the instrument argument is an essential part of the meaning of the verb.[6]

2.6.3 Thematic roles

In this textbook, we will not give a definitive answer to the question of how to determine what arguments a given verb has, or its *argument structure*, as there does not seem to be a consensus in the literature. Instead of answering this question, a lot of existing research on argument structure focuses on alternations, including the ones mentioned above. In order to describe and understand alternations, it is useful to introduce the idea of *thematic roles* (or *θ-roles*). Thematic roles are semantic generalizations about certain arguments. For example, for (2.43a), the subject 'the chef' is said to be the *agent* of the verb 'sliced' and the object 'the ham' is said to be the *theme* (or *patient*).[7] Generally, the agent of an event is the argument that intentionally performs the action being described, and the theme is something that is affected by the action. Other thematic roles include *causer*, *instrument*, *goal*, *source*, etc.

More often than not, thematic roles are not very strictly defined, and there are cases where it is difficult to determine what label to use. For instance, the subject of (2.43a) is standardly considered to be an agent, but one could insist that it is a causer because this argument causes the ham to be in the state of being sliced. It should also be mentioned that different authors use different lists of thematic roles (cf. fn. 7).

Despite this somewhat disorderly nature, thematic roles are particularly useful in describing alternations in semantic terms. For instance, the inchoative-causative alternation in (2.37) and (2.38) can be understood as the addition of a causative argument to a sentence that only has a theme argument (with a change in the position of the theme argument from the subject to the object).

[6] Another criterion often used to determine argumenthood is syntactic displacement. Since this would take us too far afield, we will simply refer the reader to the works cited in the further readings section of this chapter.

[7] The difference between theme and patient is especially confusing and controversial in the literature. A theme is typically involves some dislocation or change of some kind, while a patient does not. But some scholars deny the relevance of this difference for linguistic theory, and do not distinguish them. We will not delve into this debate in this textbook. For more information, please check the works mentioned in the further readings section of this chapter.

It makes sense to describe this phenomenon in these abstract semantic terms, because there are many verbs that participate in this alternation and they seem to form a semantic natural class; the verb describes some change of the theme argument's state. Many other alternations have also been described in similar terms in the literature.

Furthermore, thematic roles are essential in understanding how arguments are expressed in a given language more generally. That is, in addition to alternations, we also need thematic roles to account for word order constraints. For instance, in English, if a verb has an agent and a theme, then the agent is always realized as the subject, and the theme is generally realized as the object, although in some exceptional cases, it appears as a prepositional phrase, as we saw in (2.39b) and (2.40b).

In the same way, we can talk about thematic roles of nouns. For example, in a complex noun phrase like 'the student's completion of the task', 'the student' is the agent and 'the task' is the theme. Moreover, this construction shows an alternation similar to the intransitive-transitive alternation for verbs, as in 'the task's completion' or 'the completion of the task'.

To stress, thematic roles are essential in understanding how arguments are realized by the morphology and syntax of the language, and the core intuition is that meaning plays a crucial role in determining argument realization. In this sense, thematic roles are part of the theory of the syntax-semantics interface. For those of you who are interested in knowing more about this topic, we have included some references in the further readings section at the end of this chapter.

2.7 Adjectives and nominal modification

Finally, let us turn to adjectives. They are amenable to an analysis that is analogous to the above analyses of nouns and verbs. That is, their extensions can be understood as sets of entities or sets of pairs of entities, depending on whether they are intransitive or transitive. For example, the the adjective *pregnant* divides the entities in the state of affairs into those who are pregnant and those who are not pregnant and highlight the former. So for any model M, we have:

$$[\![\text{pregnant}]\!]^M = \{\, x \mid x \text{ is pregnant in } M \,\}\,.$$

We can similarly analyze adjectives like 'brown', 'spicy', 'American' and so on. These are intransitive adjectives so their extensions are all sets of entities, but there are also transitive adjectives such as 'related' and 'dependent', which are conceptually about two entities, similarly to transitive verbs. Their extensions can therefore be seen as pairs of entities, for instance:

$$[\![\text{related}]\!]^M = \{\, \langle x, y\rangle \mid x \text{ is related to } y \text{ in } M \,\}\,.$$

So far, there is no essential difference between how we analyze nouns and verbs and how we analyze adjectives. However, one big difference is that adjectives can modify nouns as in 'brown dog', and the semantic theory that we have built so far gives us a nice way of understanding what is going on with adjectival modification.

Notice that the noun phrase 'brown dog' has a very similar semantic function as a simple noun 'dog'. That is, just as the extension of 'dog' can be seen as the set of all things that this noun describes, i.e. the set of all dogs, we can see the extension of 'brown dog' as the set of all things that this noun phrase describes, i.e. the set of all brown dogs. For an arbitrary model M,

$$[\![\text{brown dog}]\!]^M = \{\, x \mid x \text{ is a brown dog in } M \,\}\,.$$

It is natural to expect that the extension of 'brown dog' has something to do with the extension of 'brown' and the extension of 'dog'. In fact, as one can see, it is the set of entities that are common members of the latter two sets, because the set of brown dogs is the intersection of the set of brown entities and the set of dogs. In other words, we have the following equation for any model M. Recall that $A \cap B$ is the *intersection* of A and B, which is the set of common members of the two sets, i.e., $\{\, x \mid x \in A \text{ and } x \in B \,\}$.

$$[\![\text{brown dog}]\!]^M = [\![\text{brown}]\!]^M \cap [\![\text{dog}]\!]^M$$

We can similarly analyze 'pregnant woman' in terms of set union. The extension of 'pregnant' is the set of all pregnant entities and the extension of 'woman' is the set of all women. Although perhaps the former is always a subset of the latter, we can still understand the extension of 'pregnant woman' as the intersection of the two sets. That is, for any model M,

$$[\![\text{pregnant woman}]\!]^M = [\![\text{pregnant}]\!]^M \cap [\![\text{woman}]\!]^M.$$

Note also that multiple adjectives can modify the same noun, as in 'smart brown dog'. If we see the extension of 'smart' as the set of all smart entities, then the extension of 'smart brown dog' will be:

$$\begin{aligned} [\![\text{smart brown dog}]\!]^M &= [\![\text{smart}]\!]^M \cap [\![\text{brown dog}]\!]^M \\ &= [\![\text{smart}]\!]^M \cap ([\![\text{brown}]\!]^M \cap [\![\text{dog}]\!]^M) \end{aligned}$$

From these observations, we can hypothesize that the semantics of adjectival modification can be modeled by set union, as in (2.44).

(2.44) Let 'A' be an adjective and 'NP' a noun phrase. Then, for any model M, $[\![\text{A NP}]\!]^M = [\![\text{A}]\!]^M \cap [\![\text{NP}]\!]^M$.

Interestingly, however, this does not work for some adjectives. For example, take an adjective 'alleged', as in 'alleged murderer'. There are two issues here. Firstly, it is hard to understand the extension of 'alleged' as a set of entities. This is an adjective that cannot semantically stand alone. Secondly, observe that $A \cap B$ is always a subset of each of these two sets. Thus, if the extension of ‘alleged murderer’ were the intersection of the extensions of the adjective and the noun, then every member of it should be a member of the extension of ‘murderer’, meaning they should all be murderers. However, it is not necessarily the case that an alleged murderer is a murderer.

Thus, ‘alleged’ is an adjective that yields a set of entities that is not a subset of the extension of the noun. There’s a term for this. We call such adjectives *non-subsective adjectives*. On the other hand, adjectives that always result in a subset of the extension of the noun are called *subsective adjectives*. We can define these terms as follows.

(2.45) An adjective A is called *subsective* if the following is the case: For any model M and for any noun phrase NP, $[\![\text{A NP}]\!]^M \subseteq [\![\text{NP}]\!]^M$. A is called *non-subsective* otherwise.

To summarize the discussion so far, ‘brown’ and ‘pregnant’ are subsective, while ‘former’ and ‘purported’ are non-subsective. The hypothesis in (2.44) is simply wrong for non-subsective adjectives.

Moreover, there is a further issue with this hypothesis. That is, subsective adjectives are not homogeneous and there are ones that cannot be captured by it. Consider, for example, ‘hard worker’. The extension of the adjective ‘hard’ can be understood as the set of hard entities, but the extension of ‘hard worker’ is not a subset of it, contrary to what the above hypothesis predicts. Or in other words, a hard worker is not (necessarily) hard! However, this is still a subsective adjective, because a hard worker is always a worker.

Adjectives whose modificational meaning can be captured by the hypothesis in (2.44) are called *intersective adjectives*, and they are only a particular class of adjectives in English. Summing up the terminology, we have the following classification of adjectives.

(2.46)

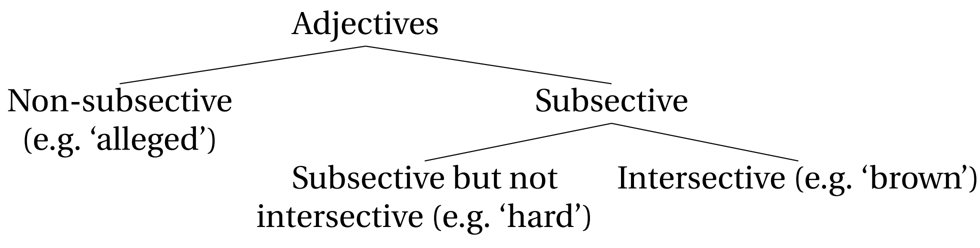

We will not give an analysis of non-intersective adjectives here, as that would require a lot more technical machinery, but it is an interesting fact that English and other natural languages usually do not seem to care much about this classification in the sense that they do not have morphological or syntactic markers for these different classes of adjectives.

2.8 Chapter summary

The key theoretical concepts introduced in this chapter are:

- The distinction between *sense* and *reference*.
- The idea of *extension* in model theory
- Concepts related to *Argument Structure*, especially *argument alternations* and the idea of *thematic roles*
- Classification of adjectives with respect to intersectivity and subsectivity.

We also introduced *Set Theory* as a formal tool. As we mentioned, it is used everywhere in linguistics as well as in other fields, so it is worth mastering it.

2.9 Further reading

Set Theory is so fundamental to many fields, numerous introductory textbooks and other materials can be found. We particularly recommend Part A of Partee, ter Meulen & Wall (1990), which goes a bit more deeply into the role of Set Theory in mathematics than we covered here.

- Barbara H. Partee, Alice ter Meulen & Robert E. Wall. 1990. *Mathematical methods in linguistics*. Dordrecht: Kluwer

Argument Structure has been quite intensively investigated in theoretical linguistics, especially in syntax. Correspondingly, a lot has been written by syntacticians on this topic. Levin & Rappaport Hovav (2009) is an advanced textbook containing in-depth description of relevant linguistic phenomena and different theoretical approaches. More recent survey articles include Ramchand (2014).

- Beth Levin & Malka Rappaport Hovav. 2009. *Argument realization*. Cambridge: Cambridge University Press
- Gillian Ramchand. 2014. Argument structure and argument structure alternations. In Marcel den Dikken (ed.), *The Cambridge handbook of Generative Syntax*, 265–321. Cambridge: Cambridge University Press

If you want to learn more about different types of adjectives and other modifiers, Morzycki (2015) is an excellent advanced textbook on the topic of modification, which should be accessible after mastering the current textbook.

- Marcin Morzycki. 2015. *Modification*. Cambridge: Cambridge University Press

Exercises

Q1. Let S_1 and S_2 as follows:

$$S_1 = \{\, \text{Bob}, \{\, \text{Alice} \,\}, \varnothing, \{\, \text{Alice}, \text{Bob} \,\} \,\}$$
$$S_2 = \{\, \text{Alice}, \{\, \varnothing \,\}, \text{Bob} \,\}$$

Given these sets, are the following statements true or false?

i) Bob $\in S_1$

ii) $\{\,\text{Bob}, \text{Alice}\,\} \in S_1$

iii) $\{\,\text{Alice}\,\} \subset S_1$

iv) $\varnothing \subset S_1$

v) $\varnothing \in S_2$

vi) $\varnothing \subseteq S_2$

vii) $\{\,\text{Bob}\,\} \in (S_1 \cup S_2)$

viii) $\{\,\varnothing\,\} \in (S_1 \cup S_2)$

ix) $\varnothing \subset (S_1 \cup S_2)$

x) $\{\, S \mid S \subseteq S_2 \,\}$ has exactly seven distinct subsets.

Q2. Represent sets that meet the following descriptions by enumerating their members between { }. If there are multiple sets that meet a description, you only need to represent one.

i) The set of colours, each of which appears both on the Union Jack and on the flag of Wales.

ii) A set that has exactly three distinct members, each of which in turn has exactly three members, each of which is a whole number less than 5.

iii) Let $S = \{\, a, b, c, d, e, f, g \,\}$. Find two subsets of S—call them X and Y—such that X has exactly three members, Y has exactly four members and $X \cap Y$ has exactly one member.

iv) Let $T = \{\, a, \{\, a \,\}, \varnothing \,\}$. Rewrite $\{\, A \mid A \subseteq T \text{ and } A \neq \varnothing \,\}$.

Q3. Represent the following sets in abstraction notation. Your answers should look like $\{\, x \mid ... \,\}$ without symbols like $\cap$ or $\cup$. You can use any variable name in place of x, but avoid variable names that are potentially confusing.

i) The set of all people who own more than one bike and live in London.

ii) The intersection of the set of all black things and the set of all cats.

iii) The union of the set of all dogs and the set of all fish.

iv) The union of the set of all things that are in France and the set of all things that are not in Paris.

v) The set of all sets that contain exactly three members.

Q4. In addition to the three set operations introduced in this chapter—union, intersection and relative complement—we can define more operations for sets. For instance, the *symmetric difference* of two sets S_1 and S_2 (written $S_1 + S_2$) is defined as the set whose members are in S_1 or in S_2 but not in both S_1 and S_2.

$$S_1 + S_2 = (S_1 \cup S_2) - (S_1 \cap S_2)$$

Let $A = \{\, 1, 2, 3, 4 \,\}$, $B = \{\, 1, 2 \,\}$ and $C = \{\, 3, 4, 5 \,\}$. What are the following sets?

i) $B + B$

ii) $B + C$

iii) $B + A$

iv) $C + \varnothing$

v) $((A - B) + (B - A))$

Q5. We analyzed the extensions of nouns and verbs as sets of certain kinds. Analyze the following nouns and verbs in a similar way for an arbitrary model

M, using the abstraction notation. Pay attention to their morphosyntactic categories (nouns, verbs), and if they can have different categories, analyze them separately. Also pay attention to semantic ambiguities, including the possibility of assigning them different argument structures.

i) car
ii) solution
iii) run
iv) chase
v) fear
vi) assign
vii) ambassador
viii) bet

Q6. Consider the family tree in Figure 2.2 again. Rewrite the following sets by listing their members.

i) $\{ x \mid x$ has a brother in Figure 2.2 $\}$
ii) $\{ \langle x, y \rangle \mid x$ is the father of y in Figure 2.2 $\}$
iii) $\{ \langle y, x \rangle \mid x$ is an uncle of y's in Figure 2.2 $\}$
iv) $\{ \langle z, x \rangle \mid x$ has a brother and z has a sister in Figure 2.2 $\}$

Q7. In this chapter, we introduced some formal properties that sets of ordered pairs of entities can have.

i) Give one noun and one verb whose extensions are asymmetric.
ii) Give one noun and one verb whose extensions are irreflexive.
iii) Give one noun and one verb whose extensions are transitive.

Q8. We can define more properties of sets of ordered pairs of entities, e.g.

(2.47) α is *anti-symmetric* iff the following holds: For each model M and for any entities x and y in M, if $\langle x, y \rangle \in [\![\alpha]\!]^M$, then $\langle y, x \rangle \notin [\![\alpha]\!]^M$, unless $x = y$.

(2.48) α is *intransitive* iff the following holds: For each model M and for any entities x, y, and z in M, if $\langle x, y \rangle \in [\![\alpha]\!]^M$ and $\langle y, z \rangle \in [\![\alpha]\!]^M$, then $\langle x, y \rangle \notin [\![\alpha]\!]^M$.

(2.49) α is *euclidean* iff the following holds: For each model M and for any entities x, y, and z in M, if $\langle x, y \rangle \in [\![\alpha]\!]^M$ and $\langle x, z \rangle \in [\![\alpha]\!]^M$, then $\langle y, z \rangle \notin [\![\alpha]\!]^M$.

Come up with one example noun or verb for each of these properties

Q9. We considered adjectival modification in this chapter, but that's not the only form of modification we can find in English. Among others, we have adverbial modification, which is very similar to adjectival modification except that adverbs modify verb phrases, while adjectives modify noun phrases. These two kinds of modification must share some common component, as similar words are often used for both purposes, e.g. 'slow' and 'slowly', and sometimes even the same word, e.g. 'fast' and 'early'. Notice also that the suffix '-ly' is a productive way of forming adverbs out of adjectives, which also suggests a systematic connection between adjectival and adverbial modification.

However, it turns out that achieving a uniform analysis of adjectives and adverbs will not be so straightforward, even for intersective adjectives and their adverbial counterparts. In order to see this, let us zoom in on some specific examples, say, ‘silent’ and ‘silently’. The adjective ‘slient’ in a nouns phrase like ‘silent man’ can be analyzed as an intersective adjective, which is to say that for an arbitrary model M, we can analyze the extension of ‘silent’ as the following set of individuals

$$[\![\text{silent}]\!]^M = \{\, x \mid x \text{ is silent in } M \,\}$$

and furthermore we have the following equation:

$$[\![\text{silent man}]\!]^M = [\![\text{silent}]\!]^M \cap [\![\text{man}]\!]^M$$

Let us now see how we could extend this analysis to the adverb ‘silently’ in a phrase like ‘silently laughed’. Our working analysis of the denotation of ‘laughed’ is in terms of a set of entities.

$$[\![\text{laughed}]\!]^M = \{\, x \mid x \text{ laughed in } M \,\}$$

If we assume that ‘silently’ has the same extension as ‘silent’, we could maintain the intersective analysis. That is, on the assumption that

$$[\![\text{silently}]\!]^M = [\![\text{silent}]\!]^M = \{\, x \mid x \text{ is silent in } M \,\}$$

we could analyze the extension of the verb phrase as follows.

$$[\![\text{silently laughed}]\!]^M = [\![\text{silently}]\!]^M \cap [\![\text{laughed}]\!]^M$$

However, this analysis is problematic. Your task in this exercise is to discuss what problem or problems it has. Hint: Suppose that Andrew silently laughed and also suppose that he sang loudly. In this situation, ‘Andrew sang silently’ is intuitively false. What does the analysis predict for this sentence? Does the prediction match the intuition?

Bibliography

Alxatib, Sam & Francis Jeffry Pelletier. 2011. The psychology of vagueness: Borderline cases and contradictions. *Mind & Language* 26(3). 287–326. https://doi.org/10.1111/j.1468-0017.2011.01419.x.

Cross, Charles & Floris Roelofsen. 2020. Questions. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Fall 2020. Metaphysics Research Lab, Stanford University.

van Deemter, Kees. 2010. *Not exactly: In praise of vaguness*. Oxford: Oxford University Press.

Frege, Gottlob. 1892. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und Philosophische Kritik* 100. 25–50.

Geeraerts, Dirk. 2010. *Theories of lexical semantics*. Oxford: Oxford University Press.

Levin, Beth & Malka Rappaport Hovav. 2009. *Argument realization*. Cambridge: Cambridge University Press.

Lewis, David. 1970. General Semantics. *Synthese* 22(1/2). 18–67.

Margolis, Eric & Stephen Laurence. 2021. Concepts. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Spring 2021. Metaphysics Research Lab, Stanford University.

Morzycki, Marcin. 2015. *Modification*. Cambridge: Cambridge University Press.

Partee, Barbara H., Alice ter Meulen & Robert E. Wall. 1990. *Mathematical methods in linguistics*. Dordrecht: Kluwer.

Ramchand, Gillian. 2014. Argument structure and argument structure alternations. In Marcel den Dikken (ed.), *The Cambridge handbook of Generative Syntax*, 265–321. Cambridge: Cambridge University Press.

van Rooij, Robert. 2011. Vagueness and linguistics. In Giuseppina Ronzitti (ed.), *Vagueness: a guide*, 123–170. Springer. https://doi.org/10.1007/978-94-007-0375-9_6.

Sauerland, Uli. 2011. Vagueness in language: The case against fuzzy logic revisited. In Petr Cintula, Christian G. Fermüller, Lluis Godo & Petr Hajek (eds.), *Understanding vagueness: logical, philosophical and linguistic perspectives*, vol. 36, 185–198. Rickmansworth: College Publications.

Sennet, Adam. 2021. Ambiguity. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Smmer 2021. Metaphysics Research Lab, Stanford University.

Solt, Stephanie. 2015. Vagueness and imprecision: Empirical foundations. *Annual Review of Linguistics* 1. 107–127. https://doi.org/10.1146/annurev-linguist-030514-125150.

Sorensen, Roy. 2018. Vagueness. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Summer 2018. Metaphysics Research Lab, Stanford University.

Terkourafi, Marina. 2016. The linguistics of politeness and social relations. In Keith Allan (ed.), *The Routledge handbook of linguistics*, 221–235. Abingdon-on-Thames: Routledge.

Vicente, Agustí & Ingrid L. Falkum. 2017. *Polysemy*. https://doi.org/10.1093/acrefore/9780199384655.013.325.

Willaims, J. Robert G. 2012. Vagueness. In Gillian Russell & Delia Graff Fara (eds.), *The Routledge companion to philosophy of language*, 143–152. London: Routledge.