

# The Study of Meaning

## An Introduction to Semantics and Pragmatics

Nathan Klinedinst

Yasutada Sudo

*University College London*



# Chapter 1

## The study of meaning

### 1.1 Linguistic meaning

Meaning is one of the central objects of study in linguistics, along with sounds, words, etc. As linguists, we are especially interested in meaning conveyed by expressions in *natural languages*.<sup>1</sup> We will call such meaning *linguistic meaning* in this textbook.

Linguistic meaning is not the only kind of meaning there is. Expressions of artificial languages obviously have meaning. For instance, in the language of arithmetic, the symbols ‘+’ and ‘=’ have some meaning (which you could paraphrase in a natural language, if you like). Furthermore, things that are usually not called ‘languages’ can have meaning. For example, the green light of the traffic light, a double-finger tap on the iPad, or (at least for some people) a certain constellation of stars and planets can be said to carry meaning. This being a textbook in linguistics, we will set linguistic meaning as our primary object of study, but the theoretical techniques that will be introduced in this textbook could potentially be applied to non-linguistic meaning. It is also interesting to notice that the opposite has been done quite often as well. That is, linguists have borrowed many ideas, analytical techniques, and methodologies developed for studying artificial languages, especially logic, to study linguistic meaning. We will see concrete examples of them throughout this book.

That linguistic meaning exists is quite obvious. For instance, anyone who speaks English would agree that the word ‘cat’ has some core meaning that is more or less fixed and shared across speakers of English, and similarly, a larger phrase consisting of multiple words like ‘walk a dog in the park’ has some other meaning. What is interesting about such meanings? The goal of this textbook is to convince you that interesting research questions can be asked about them, and show you some concrete examples of interesting research done in the past. In this first chapter, we will review some fundamental notions and theoretical concepts that are useful in thinking about linguistic meaning.

---

<sup>1</sup>Natural languages are languages that are naturally acquired by human children. They contrast with *artificial languages*, which include the language of mathematics, programming languages, and constructed languages like Klingon.

## 1.2 Semantics and pragmatics

The subarea of linguistics that studies meaning is traditionally divided into *semantics* and *pragmatics*. Roughly put, semantics is about meaning encoded in linguistic expressions, while pragmatics is about meaning that arises when expressions are put in use. To illustrate the difference with an example, let's consider the expression 'My sister speaks Cantonese', which is a type of expression called a sentence. Our intuitions as speakers of English tell us that this sentence has some fixed meaning. If, for example, someone who is learning English asks you what this sentence means, you might try to express it by paraphrasing it with other expressions in English or in some other language that you and your conversational partner commonly know. That meaning, although quite abstract, can be said to be the semantic meaning of the sentence.

Pragmatic meaning is the kind of meaning that arises when expressions are used in actual conversations. For instance, consider the following short dialogue containing the same sentence 'My sister speaks Cantonese'.

- (1.1) Anna: Do you speak Cantonese or Korean?  
Ben: My sister speaks Cantonese.

From this dialogue, one tends to understand that Ben himself does not speak Cantonese.<sup>2</sup> Importantly, this arguably does not follow directly from the semantic meaning of the expression that Ben used, namely 'My sister speaks Cantonese', which should be solely about the speaker's sister's linguistic abilities, and should not imply anything about whether the speaker himself speaks Cantonese or not. In fact, in a different dialogue, the inference that Ben doesn't speak Cantonese does not arise from the exact same sentence, as demonstrated in (1.2), for instance.

- (1.2) Anna: Does any of your siblings speak Cantonese or Korean?  
Ben: My sister speaks Cantonese.

Unlike in the previous example, it doesn't necessarily follow from Ben's utterance in this dialogue that Ben does not speak Cantonese. In fact, the participants of this dialogue could even be assuming that Ben does speak Cantonese. These observations suggest that the meaning of Ben's utterance in (1.1) that Ben himself doesn't speak Cantonese is not encoded in the semantics of the sentence itself, because if it were, it should be always observed, including (1.2).

Having observed that the inference that Ben doesn't speak Cantonese arises in (1.1) but not in (1.2), we should now ask, where does this inference come from? It seems plausible to say that it arises because Ben does not directly answer Anna's question in (1.1). Contrast this with (1.2), where he does provide a direct answer to her question. This sounds like a promising analysis, but sev-

---

<sup>2</sup>Notice that Ben's utterance can also be taken to imply that neither Ben nor Ben's sister speaks Korean. This will be discussed in an exercise for this chapter.

eral further questions need to be answered, before calling it a proper analysis. What does it mean to give a ‘direct answer’ to a question? And why is it that not directly answering the question results in this particular meaning in (1.1)? Pragmatics is a field that is concerned with questions like these.

Later in this book (Chapter ??), we will introduce an influential theory of pragmatics, put forward by the British philosopher Herbert Paul Grice in the middle of the 20th century, which will give us crucial insights into pragmatic inferences like the above. Roughly put, the most important insight behind Grice’s theory of pragmatics is that rational and competent conversational agents draw extra inferences from linguistic utterances based on the assumption that their conversational partners are also rational and competent and abide by a common set of pragmatic principles called the *Conversational Maxims*. He called pragmatic inferences drawn this way *implicatures*. We will discuss the details of the Gricean theory of pragmatics in Chapter ??, and see how the above example and other similar cases can be analyzed in this theory as implicatures.

In sum, linguistic meaning can be classified into two broad categories, linguistic meaning encoded in expressions (*semantic meaning*) and linguistic meaning that arises when expressions are used in conversation (*pragmatic meaning*). Semantics is the study of the former and pragmatics is the study of the latter. This being said, it should be kept in mind that where exactly to draw the line between these two subfields is far from a settled issue. For the above example of pragmatic meaning in (1.1), there is convincing evidence to think that the relevant meaning is pragmatic in nature. In particular, the fact that it is sometimes absent, as illustrated by (1.2), strongly suggests this conclusion. However, there are cases that cannot be readily decided whether they are semantic or pragmatic meanings. We will come back to this issue, when we encounter relevant linguistic phenomena, but ultimately, where exactly to draw the line between semantics and pragmatics can only be meaningfully answered with a theoretical framework in mind. For this reason, we will regard the subarea of linguistics that studies linguistic meaning as a single area for now, and for want of a better term, call it the study of meaning.

### 1.3 The study of meaning in modern linguistics

Different subareas of linguistics study different aspects of language. *Phonology* is the study of linguistic sound patterns, *morphology* is the study of word structure, *syntax* is the study of phrase structure, and so on. Among these aspects of language, meaning might look particularly abstract. Indeed, linguistic sounds can be heard and measured, and it is quite intuitive to represent them using symbols like ‘b’ and ‘o’. Structures of words and phrases are somewhat more abstract than linguistic sounds, but can be visually represented by depicting their structural relationships in certain ways, e.g. tree diagrams.<sup>3</sup>

---

<sup>3</sup>The most common tool for visualizing syntactic structure is tree diagrams, but the bracket notation is often used as well. For example, the structure of the sentence ‘Daniel likes Eva’ can be represented as a tree diagram in (i) or in bracket notation as in (ii). (Note that this is


Such visual representations are certainly very useful and have a lot of heuristic value, but it should be kept in mind that the real objects of study in theoretical linguistics are always quite abstract in nature. More specifically, modern theoretical linguistics, which was initiated by Noam Chomsky in the 1960s, is the study of *linguistic knowledge* that humans (can) have, which is abstract, and all subareas of linguistics, including the study of meaning, should be seen as studies of different aspects of linguistic knowledge.

Phonologists certainly talk about concrete, measurable sounds, but what they aim at theorizing is not these concrete sounds themselves. A particular instance of a particular linguistic sound articulated by a particular speaker at a particular time holds no more intrinsic interest for a phonologist than does a particular movement of a particular object at a particular time for a theoretical physicist. Rather, phonology is concerned with the knowledge behind sound patterns observed in natural languages. For example, it is an interesting fact for the phonologist that suffixation of ‘-al’ in English systematically shifts stress, as in ‘párent’ vs. ‘paréntal’, but suffixation of ‘-hood’ doesn’t, as in ‘párent’ vs. ‘párenthood’, because this indicates that speakers of English have knowledge about stress assignment that is sensitive to the type of suffix.<sup>4</sup>

Notice that the crucial data here are *well-formedness intuitions* about the phonological realizations of words. That is, native speakers of a natural language can tell whether a given sound or a sequence of sounds feels correct or incorrect in that language, so to speak. For instance, in English, ‘paréntal’ sounds right, while ‘párental’ doesn’t sound right, and similarly, ‘párenthood’ sounds right, and ‘parénthood’ doesn’t sound right. Phonology aims at explaining these observed intuitions. Such intuitions themselves are certainly abstract, but are concrete enough to construct a scientific theory for, as demonstrated by past phonological research.

---

a toy example, and a more serious syntactic analysis might assign a different, perhaps more complicated, representation to this sentence.)



(ii) [s [DP Daniel] [VP [V likes] [DP Eva] ] ]

These two representations contain the same amount of information and represent the exact same structure, because one can be uniquely reconstructed from the other. This means it does not matter much which representation to use. However, which one is easier to understand for the human mind is obvious.

<sup>4</sup>A *suffix* is a kind of expression, such as ‘-al’ and ‘-hood’, that is not a word itself and always attaches to the end of some other expression (which can be a word or smaller than a word itself).

Similarly, syntacticians are not interested in the tree diagrams themselves, but in the syntactic knowledge that gives rise to the observed syntactic intuitions. Speakers of any language have syntactic intuitions about which structures are well-formed (or *grammatical*) or ill-formed (or *ungrammatical*). Here is an example from English: modal auxiliaries such as ‘can’ and ‘must’ are systematically ungrammatical in non-tensed contexts, as shown in (1.3) below. The asterisk ‘\*’ is commonly used to indicate unacceptability of expressions.

- (1.3) a. \*The woman wants to can sing better.  
b. \*I should can do it by tomorrow.  
c. \*Must pay attention to everything is tiring.

There are ways to express the intended meanings of these sentences in English, as demonstrated by the examples in (1.4), so it is not the intended meanings that make the sentences in (1.3) unacceptable. Rather, there are certain grammatical constraints on modal auxiliaries, which constitute part of the syntactic knowledge that speakers of English have.

- (1.4) a. The woman wants to be able to sing better.  
b. I should be able to do it by tomorrow.  
c. Having to pay attention to everything tiring.

Syntactic data is of the same nature as phonological data, namely, intuitions about well-formedness, or *grammaticality judgments*, as they are often called. Syntacticians collect grammaticality judgments and construct theoretical hypotheses about native speakers’ syntactic knowledge.

Morphology works essentially the same way, except that morphological data are well-formedness intuitions about word structure, rather than phrase structure. Words can be simple, like ‘desk’, or complex, like ‘un-con-stitut-ion-al’. The goal of morphology is to understand morphological knowledge that is responsible for grammatical judgments about such word structure.<sup>5</sup>

It is an interesting fact about language that one typically does not have conscious access to most knowledge about one’s native language. Moreover, a lot of this knowledge must have been learned at some point in one’s life, which makes it all the more striking that most of the linguistic knowledge is implicit. If you are a speaker of English, for example, you probably cannot list all the suffixes that trigger stress shift like ‘-al’ does. But if you think about it, you must have learned these suffixes one by one when you were small, and in do-

---

<sup>5</sup>Some recent theories of morphology hold that complex words are built based on the same principles and operations that are used to build complex phrases out of words. If this is correct, then morphology and syntax should be treated as a single subarea, *morphosyntax*. Similarly, there are aspects of morphology that have a lot to do with sounds. The aforementioned examples of stress shift are arguably not a pure phonological phenomenon, as morphological knowledge about the two suffixes *-al* and *-hood* is also relevant. We can call the study of phenomena like this *morphophonology*.

ing so, you must have determined whether they are stress-shifting like ‘-al’, or not, like ‘-hood’, because there’s no way of knowing beforehand which suffixes are stress-shifting and which are not in English. This learning process is completely unconscious, and what is learned becomes implicit knowledge. It is linguists’ job to collect such linguistic intuitions, make hypotheses about what knowledge lies behind these linguistic intuitions, and how children manage to acquire that knowledge.

In the case of the study of meaning, main empirical data are intuitions about linguistic meaning. One crucial difference from phonology, morphology and syntax is that data for the study of meaning are not intuitions about well-formedness, because meaning has no form. As we will see below, there are a number of different types of intuitions about linguistic meaning, but they are as abstract as, or rather, as concrete as, well-formedness intuitions used in phonology, morphology, and syntax, and hypotheses and theories about them can be constructed.

## 1.4 Lexical semantics

Let us start with intuitions about the semantic meaning of single words, or *lexical semantics*. As mentioned above, linguistic knowledge is generally implicit and when one is asked about the meaning of a random word, one often cannot really say what exactly it is. For example, what do you think is the meaning of the word ‘chair’? If you speak English, you certainly know what it means, but what is that meaning? The entry for this word in a dictionary might give you some ideas about what it is, but the information that you can find in a dictionary is nothing more than a paraphrase of the word, and this paraphrase, being a linguistic expression itself, has some meaning of its own, which is identical or at least sufficiently close to the meaning of ‘chair’. But that doesn’t answer the question what the meaning of ‘chair’ itself is.

One might say that the meaning of the word ‘chair’ has something to do with the *concept* that it represents. Maybe this concept *is* the meaning of this word. This sounds plausible, but then the next question is, what is a concept? This is a long-standing, unsolved issue, and many different views have been proposed. As a matter of fact, a lot is still unknown about what concepts really are (If you want to know more about this topic, there is a suggested reading in the further readings section at the end of this chapter). If we do not know what concepts are, it is not possible to construct a theory of linguistic meaning based on them. So let us put aside this question about concepts for now. This does not mean that we cannot study the meanings of words like ‘chair’, because there are certain semantic intuitions about them that are sufficiently clear. Let us look at some such intuitions.

### 1.4.1 Hyponymy and hypernymy

While we do not have very clear intuitions about the meanings of individual words, we often have much clearer intuitions about the relation between the



meanings of two words, for instance, the meaning of ‘cat’ and the meaning of ‘mammal’. Specifically, our intuitions tell us that the former is a subcategory of the latter, because every cat is a mammal. There are many such relations among words, and technical terms for describing them. Let us look at some representative ones.

We say the word ‘cat’ is a *hyponym* of the word ‘mammal’, and conversely, the word ‘mammal’ is a *hypernym* of the word ‘cat’. A hyponym is a more specific word than its hypernym. Here is another example: ‘armchair’ is a hyponym of ‘chair’ and ‘chair’ is a hypernym of ‘armchair’.

It’s worth stressing that hyponymy/hypernymy is based on semantic intuitions like the following.

- (1.5) Everything that is describable by ‘cat’ is describable by ‘mammal’, but not everything that is describable by ‘mammal’ is describable by ‘cat’.

This reflects our semantic knowledge about these two words.

By generalizing this statement, we can define the hyponym/hypernym-relation as follows:

- (1.6) A word  $w_1$  is a *hyponym* of another word  $w_2$  (and equivalently,  $w_2$  is a *hypernym* of  $w_1$ ) if and only if everything that is describable by  $w_1$  is describable by  $w_2$  but not everything that is describable by  $w_2$  is describable by  $w_1$ .

The locution ‘if and only if’ is very commonly used in linguistics, as well as in other disciplines, and it’s worth clarifying what it means at this point. When flanked by two sentences, schematically as in ‘ $\phi$  if and only if  $\psi$ ’, it means that whenever  $\phi$  is the case,  $\psi$  is also the case, and moreover, whenever  $\psi$  is the case,  $\phi$  is also the case. This essentially means  $\phi$  and  $\psi$  are equivalent in meaning, and for this reason, ‘if and only if’ is widely used in giving definitions like (1.6). In fact, it’s so often used that it’s very common to use the shorthand ‘iff’. We will adopt this convention from now on.

Now (1.6) certainly sounds very pedantic, and you might wonder why we do this for something as obvious as hyponymy/hypernymy. The answer is, it’s generally good practice to define your technical terms explicitly in order to make every detail of your theory crystal-clear. It is probably overkill in this particular case, but as a matter of fact, in doing research, we often encounter situations where we seem to have an idea or intuition that is sufficiently clear but it is still far from trivial to spell out what that idea is in a clear manner like (1.6). We will see an example of one such case in an exercise for this chapter.

There are several remarks to be made about the definition in (1.6). Firstly, we have the clause ‘but not everything that is describable by  $w_2$  is describable by  $w_1$ ’, because if it were not for it, two words that mean the same thing would be hyponyms/hypernyms of each other. One could actually use these terms

of include such pairs of words, but the widely used definitions exclude them. Secondly, the hyponym/hypernym-relation obviously does not hold between every pair of words. For instance, the word ‘dog’ is neither a hyponym or hypernym of ‘chair’. Thirdly, note that the definition in (1.6) is not limited to nouns, and this is intentional, because the relation holds between words of other categories as well. For example, the verb ‘stroll’ is a hyponym of the verb ‘walk’. There is generally a question about what verbs describe, but intuitively they describe actions. And whenever some action is describable by ‘stroll’, it is also describable by ‘walk’, but not the other way around.

### 1.4.2 Synonymy

Another semantic relation between words for which we have clear intuitions is the relation of *synonymy*. When two words have the same meaning, we say that they are *synonyms* of each other, or they are *synonymous*. For example, the word ‘disease’ and ‘illness’ are synonyms of each other. We can give a pedantic definition in the same style as (1.6) as follows.

- (1.7) A word  $w_1$  is a *synonym* of a word  $w_2$  iff everything that is describable by  $w_1$  is describable by  $w_2$  and everything that is describable by  $w_2$  is describable by  $w_1$ .

Note that according to this definition,  $w_1$  is a synonym of  $w_2$  iff  $w_2$  is a synonym of  $w_1$ , and this is as desired, because the synonym-relation intuitively always holds mutually. Thus, (1.7) can be restated as follows.

- (1.8) A word  $w_1$  and a word  $w_2$  are *synonyms* of each other iff everything that is describable by  $w_1$  is describable by  $w_2$  and everything that is describable by  $w_2$  is describable by  $w_1$ .

One can often find such equivalent ways of saying the same thing in different textbooks, research articles, and monographs. Since they are equivalent in meaning, their superficial differences do not make any practical difference, so you can pick any of them. But you should keep in mind that equivalence and non-equivalence between different statements are sometimes not immediately clear.

The idea of synonymy sounds very simple, but there are certain complications that should be mentioned. There are cases where the two meanings seem to be basically identical but have certain subtle differences at the same time. Consider, for example, the noun ‘cat’ and the noun ‘feline’. They can be seen as synonyms, and the above definition would render them synonyms, but at the same time, intuitions tell us that ‘feline’ is a formal word and less often used in everyday speech, while ‘cat’ is more or less neutral in this respect, or perhaps eschewed in certain scientific contexts in favor of ‘feline’. Such intuitions about formality and register of speech are certainly part of your linguistic knowledge. However, whether or not this type of linguistic knowledge should be consid-

ered as part of the lexical semantic knowledge is less obvious. If it is not part of the semantics of these words, we should treat these two nouns as synonyms of each other, but it also appears legitimate to say that what types of context ‘feline’ can be felicitously used in is part of the meaning of this word, in which case, the meanings of these nouns are not completely identical and so the two words should not count as synonyms (and also we should change the above definition of synonymy accordingly).

This question touches on the issue of the semantics-pragmatics divide, mentioned at the end of Section 1.2 above. One could say that which context the word ‘feline’ can be felicitously used is part of pragmatic knowledge and is not semantic in nature, because, after all, it’s about when to use the word. On the other hand, it also seems sensible to say that this knowledge is different in nature from pragmatic inferences that are drawn in conversations like the example we saw in Section 1.2. In some sense, this phenomenon has both pragmatic and semantic aspects. How to deal with different levels of formality and speech registers is an interesting and actively researched topic in the study of meaning today, but it is outside of the scope of this textbook, so we will leave these questions unanswered here (Please see the further readings section of this chapter for references). Please bear in mind that this textbook mentions a number of open issues like this one, and we would like to encourage you to investigate further whatever you find interesting. As a matter of fact, there are so many things we still don’t know about linguistic meaning, and this textbook is meant to be a guide for finding interesting questions, rather than answers to them.

### 1.4.3 Homonymy

Another clear lexical semantic intuition is *homonymy*. When two words have the same pronunciation but different meanings, we say they are *homonyms* of each other. As before, let us give it a pedantic definition. Having defined the term ‘synonym’ above, we can use it here to make the definition compact.

(1.9) A word  $w_1$  is a *homonym* of a word  $w_2$  iff by  $w_1$  and  $w_2$  have the same pronunciation and are not synonyms of each other.

Words that have homonyms are generally said to be *ambiguous*. For instance, the noun ‘bank’ has two meanings, one that is about the sides of a river, and one that is about a type of financial institution. Another example is the word ‘bat’, one meaning is about a kind of flying animal, and the other meaning is about a kind of instrument used in certain sports like baseball and cricket. We could say that in each of these cases there are actually two nouns, each having one of the two meanings, but they happen to be pronounced in the same way.

For these examples of homonymy, it seems that they are homonyms *by accident*. That is, there seems to be no principled reason why the two meanings are conveyed by the same sequence of sounds, and the two words could well be pronounced differently. In fact in a number of other languages, the corre-

sponding expressions are typically not homonyms.

On the other hand, there also are words that have multiple meanings that are somehow systematically related. Such cases are called *polysemy*. For instance, take the noun 'lunch'. This noun means a type of meal, typically eaten in the middle of the day, but it can also mean the event of eating the meal. To see this difference, consider the following examples.

- (1.10) a. The lunch was delicious!  
b. The lunch started with a boring speech by the host.

In (1.10a), 'the lunch' refers to food, but not in (1.10b), because food is not something that starts with a speech. Rather, it refers to an event of eating lunch. To describe this ambiguity in meaning, we say that the noun 'lunch' is *polysemous*.

What counts as a case of polysemy rather than a case of homonymy is often not so clear and needs to be carefully checked. One indicative criterion for polysemy is systematicity. For instance, in the case of 'lunch', similar words, such as 'breakfast', 'dinner', 'meal', and so on have a similar range of meanings. Furthermore, the same polysemy is often observed in other languages (although potentially not in all) as well. These facts suggest that there is some general mechanism that gives rise to these different meanings from some core meaning. What is importantly for linguists is that this mechanism is by hypothesis part of the linguistic knowledge, and theories of polysemy aim at understanding what the mechanism is and how it works. In the case of 'bank' or 'bat', there is no such systematicity, so we can say that they are cases of homonymy, and not of polysemy.

Also certain linguistic tests can be used to distinguish polysemy and homonymy. For instance, the two meanings of a polysemous word can be expressed by a single occurrence, as in (1.11a), a pronoun and its antecedent need not have the same meaning, as in (1.11b), and an elliptical noun and its antecedent need not have the same meaning, as in (1.11c).<sup>6</sup>

- (1.11) a. The lunch was delicious, but started with a boring speech by the host.  
b. The lunch was delicious, but it started with a boring speech by the host.  
c. My lunch was delicious. Daniel's started with a boring speech by the host.

It is obvious that these are not possible with accidentally ambiguous words like 'bat', as shown in (1.12). In these examples, only one meaning of 'bat' is avail-

<sup>6</sup>An expression like 'Daniel's' in (1.11c) is said to contain an elliptical noun (or nominal ellipsis), as it is equivalent in meaning to 'Daniel's lunch'. An *antecedent* of a pronoun or elliptical noun is an expression that is semantically related to the pronoun or elliptical noun and occurs in the same sentence or another sentence in the same discourse.




Figure 1.1: Two line segments

able at a time.

- (1.12) a. The bat was made of wood, and flying above my house.  
b. The bat was made of wood, and it was flying above my house.  
c. My bat is made of wood. Elena's is flying above my house.

Note that two senses of an ambiguous word are not always completely unrelated. For example, 'foot' in the sense of a body part and 'foot' in the sense of a unit of length are historically/etymologically related, but this is still not a case of polysemy, as there is no systematicity, and also the following example shows that an elliptical noun must have the same meaning as its antecedent 'feet', which renders the overall of meaning of the sentence incoherent and hence *infelicitous*. The hash mark ('#') at the beginning of the example indicates infelicity (as opposed to ungrammaticality, indicated by '\*').

- (1.13) #This rope is only five feet in length but is enough to tie up the duck's.

#### 1.4.4 Content vs. function words

Arguably, the semantic relations between words we discussed above reflect the relations between the concepts that the words express. Ultimately, we want to synthesize the theory of lexical semantics and the theory of concepts, but this needs to wait until we have a better grasp at what concepts are to begin with. You might find it counter-intuitive that we have clearer intuitions about relations between concepts than about concepts themselves, but something like this can be observed in many areas. For instance, if you are given two line segments, as in Figure 1.1, it's very hard to tell how long these segments are without measuring them, but you immediately know certain relations between them, for example, which one is longer than the other, whether or not they are parallel, etc. Intuitions about lexical semantics are similar to this.

It should also be pointed out that the idea of concepts is more useful for analyzing the lexical semantics of words like 'cat' and 'walk', which are often called *content words*, but less insightful for the meanings of words like 'if', 'the', 'every', and 'whether', which are more of grammatical devices used in constructing complex phrases. These words are called *function words*. Function words seem to have some meanings, but can we characterize them as concepts on a par with content words? If not, what are these meanings?

The distinction between content words and function words is not a sharp

one. Content words are the ones that tend to carry the main information of a linguistic utterance, while function words look more like building blocks that glue different parts of the expression together. Content words tend to be an *open class* in the sense that more and more new words can be added to the language, while function words tend to be a *closed class* and it is very rare, if not impossible, to find a newly added function word.

You might be surprised to learn that a large part of the semantic literature is about function words, rather than about content words. The reason is because the analytical tools and concepts that linguistics have been using are better suited for analyzing the meanings of function words than the meanings of content words. Also, although attempts have continuously been made, the fact that we still do not really understand what concepts are is a major hindrance to developing a satisfactory theory of content words. For this reason semantic theories tend to limit themselves to certain particular aspects of content words, and mainly target other questions about linguistic meaning. This of course doesn't mean that one cannot study content words, nor is it uninteresting to try to do so. Yet, as we will see in the rest of this book, there are many interesting questions we can ask even with this limitation.

## 1.5 Truth-conditions and entailment

Lexical semantics does not exhaust linguistic meaning. Linguistic expressions that are larger than a single word, for example 'red chair', also have meanings. How should we analyze such meanings? Intuitively, the meaning of 'red chair' is related to the meaning of the word 'red' and the meaning of the word 'chair'. If the meaning of 'red' is the concept of red-ness and the meaning of 'chair' is the concept of chair-ness, how do these concepts relate to the meaning of 'red chair'? And is that relation same as the relation between the meaning of 'angry' and the meaning of 'neighbor' in the phrase 'angry neighbor', which seems to have a similar syntactic structure? These are questions about *compositionality*, the idea that linguistic meanings can be combined to give rise to other meanings. We will defer the discussion of how to build a theory of *compositional semantics* until the final chapter of the book (Chapter ??), but it should be pointed out here that such a theory is built upon specific kinds of semantic intuitions that native speakers have about syntactically complex phrases, most notably, sentences. In this section, we will introduce some important notions about sentence meaning.

### 1.5.1 Truth and falsity

The most fundamental notion about sentence meaning is *truth* and *falsity*. By way of illustration, consider the following (very boring) sentence.

(1.14) There is a circle in a triangle.

Now consider the situation depicted in Figure 1.2. One has an intuition that




Figure 1.2: A triangle, a circle and a square

this sentence is *true* in this situation. Similarly, any speaker of English would say that the sentence in (1.15) is *false* in the same situation.

(1.15) There is a square in a triangle.

This might look unsurprising to you, but it should be emphasized that these intuitions are part of the semantic knowledge of a competent speaker of English. Needless to say, the same can be demonstrated for any natural language. More generally, a native speaker can typically tell if a given sentence is true or false in a given situation.

This claim, however, is to be qualified in several respects. Firstly, there are sentences for which the intuitions about their truth and falsity may be murky. For instance, perhaps the sentence in (1.16) does not sound true, but does not sound completely false either in the same situation depicted in Figure 1.2. Rather, you might say ‘It depends’, or maybe ‘It’s neither big nor small’.

(1.16) The square is big.

Many sentences in natural languages are similarly *vague*. One might think that *vagueness* is a major obstacle in studying the semantics of sentences in natural languages, but as a matter of fact, it is more of an interesting research topic than an insurmountable issue for the study of meaning. We will discuss the issue of vagueness and several ways in which it has been analyzed in greater detail in Chapter ??.

Similarly, sentences can fail to be clearly true or clearly false for other reasons than vagueness. For instance, the following sentence is definitely not true in the same situation as above, but it is also not really false.

(1.17) The pentagon is in the triangle.

This would be false if there were one pentagon outside the triangle, but in the absence of a (unique) pentagon, it is not possible to say if (1.17) is true or false. This intuition has to do with the kind of meaning we call *presupposition*, which we will discuss in more detail in Section 1.5.5.

Secondly, the idea of truth and falsity simply does not apply to some sen-

tences. The examples we have discussed so far in this section are so-called *declarative sentences*, which are a syntactically defined class of sentences that are typically used for asserting or stating something. Natural languages contain other types of sentences. For instance, *interrogative sentences* like (1.18a), which are typically used for asking questions, and *imperative sentences* like (1.18b), which are typically used for giving commands.

- (1.18) a. Is there a circle in a triangle?  
b. Draw an oval in the square!

Evidently, these sentences are not said to be true or false. Yet, semanticists have noticed that their meanings still have to do with truth/falsity in some way. Roughly, the idea is that (1.18a) is asking whether the sentence in (1.14) is true or false, and (1.18b) is essentially requesting the hearer to bring about a situation where the sentence ‘You have drawn an oval in the square’ becomes true. This way, the semantic intuitions about truth/falsity are considered to be fundamental in understanding the meanings of various types of sentences. Unfortunately, we do not have enough space in this textbook to cover such non-declarative sentences, and we therefore simply say ‘sentences’ to mean declarative sentences in the rest of the textbook.

### 1.5.2 Truth conditions

As we have just seen, our semantic knowledge about the sentences in (1.14) and (1.15) tells us they are true and false respectively in the situation depicted in Figure 1.2. But what is that semantic knowledge itself? To answer this question, it is not enough to just look at this one situation, because it is just one of many other situations where these sentences are true or false. In fact, there are infinitely many such situations to consider, because these situations that we have intuitions about do not need to be realized in reality. For instance the truth of (1.14) will stay the same in a situation that is very much like the one in Figure 1.2 except that the triangle is yellow, or one where a hexagon is added somewhere in the situation. On the other hand, the sentence becomes false if the circle is replaced by a hexagon, or any other shape, or if it is simply removed. Thus, even with respect to Figure 1.2 alone, one can easily come up with a slew of variants, and in each of them our semantic knowledge can tell us whether the sentence (1.14) is true or false in it.

This means that we cannot characterize the semantic knowledge about the truth and falsity of (1.14) by simply listing situations where it is true and situations where it is false, because such a list will need to be infinitely long and will not end. Rather, we can state it in the form of a condition that needs to be met for the sentence to be true, together with a condition that needs to be met for it to be false. For example, in the case of (1.14), we can state the conditions as follows.

- (1.19) a. The sentence in (1.14) is true if there is a circle and there is a tri-



- angle and the circle is contained in the triangle.
- b. The sentence in (1.14) is false if there is no circle, or there is no triangle, or there is a circle and a triangle but the circle is not contained in the triangle.

This is a very verbose way of stating the conditions. We could actually simplify them as follows.

- (1.20) a. The sentence in (1.14) is true if a circle is in a triangle.
- b. The sentence in (1.14) is false if it's not true that a circle is in a triangle.

Semanticists often abbreviate these conditions further as a single statement using 'iff' as follows (recall that 'iff' means 'if and only if').

- (1.21) The sentence in (1.14) is true iff a circle is in a triangle.

The following statement means exactly the same thing as (1.21), but it's customary to mention the true case as in (1.21), rather than the false case as in (1.22).

- (1.22) The sentence in (1.14) is false iff it's not true that a circle is in a triangle.

Note that these statements correctly predict that (1.14) is true in the situation depicted in Figure 1.2, and similar for other variants of this situation mentioned above. These conditions are called the *truth-conditions* of the sentence in (1.14), and are meant to characterize our semantic knowledge about the truth and falsity of this sentence. We can similarly state the truth-conditions of (1.15) as follows.

- (1.23) The sentence in (1.15) is true iff a square is in a triangle.

What is surprising about natural language is that a native speaker has such *truth-conditional knowledge* about infinitely many sentences. To illustrate, consider the following syntactically complex sentences.

- (1.24) a. It is not the case that a circle is in a triangle.
- b. It is not the case that it is not the case that a circle is in a triangle.
- c. Either a circle is in a triangle or it is not the case that a square is in a triangle.
- d. If it is not the case that a square is in a triangle, then a circle is in a triangle.
- e. It is not the case that if it is not the case that a square is in a triangle, then a circle is in a triangle.

As shown here, one can create more and more complex well-formed sentences out of (1.14) and (1.15). For example, the sentence in (1.24a) is created from (1.14) by adding ‘it is not the case that’. In fact, there is no upper-bound on the complexity of a well-formed sentence in a natural language. This is because natural languages have a property called *recursion*, which allows sentences to be part of other sentences (and more generally, phrases to be part of phrases of the same kind). Concretely, ‘it is not the case that’ can be prefixed to a sentence to produce another sentence. For example, when prefixed to (1.14), it yields (1.24a). The result is a sentence, so this operation can be repeated, yielding (1.24b). Certainly, this is not a very natural sentence, and it is perhaps not very easy to understand either, but it is still a grammatical sentence that has some truth-conditions. The other examples in (1.24) illustrate the same point.

What is particularly important about recursion is that it produces infinitely many grammatical sentences. And for the study of meaning, it is crucial that native speakers know the truth-conditions of infinitely many sentences, because this means that this knowledge cannot be characterized by a list of truth-conditions of sentences, for such a list would be infinitely long and could not be memorized. Rather, there must be a mechanism that allows one to *compute* truth-conditions. This idea is called *compositionality*, and we will discuss it in more detail in Chapter ??.

What is important for now is that there are infinitely many sentences whose truth-conditions we are interested in. It turns out that we can state the truth-conditions of any sentence  $\phi$  in English as follows.

(1.25)  $\phi$  is true iff  $\phi$ .

If you look back at (1.21) and (1.23), they are almost in this format, and it is evident that they are equivalent to (1.26).

(1.26) a. ‘There is a circle in a triangle’ is true iff there is a circle in a triangle.  
 b. ‘There is a square in a triangle’ is true iff there is a square in a triangle.

These statements might look trivial. Indeed, the truth-conditions expressed this way might not be very insightful, but they are actually not trivial, and it is important to understand this. The apparent triviality of (1.25) arises because we are analyzing English using English. Generally, the language that is being analyzed is called the *object language* and the language that is used to state the analysis is called the *metalanguage*. In the above representations, what appears inside the quotation marks is an expression in the object language and what appears outside the quotation marks belongs to the metalanguage.

Obviously, if the object language and the metalanguages are different languages, a general statement like (1.25) is impossible to form. For example, if the object language is Japanese, you will state truth-conditions as in (1.27).

(1.27) ‘Sankakuno nakani maruga aru’ is true iff there is a circle in a triangle.

This is clearly not a trivial statement and could well be wrong. The object language can be any language that you are interested in. In the general case it need not be a natural language.

Similarly, the metalanguage could be any language. It could even be a formal language. In fact, you might see something along the lines of (1.28) in published articles.

- (1.28) a. ‘There is a circle in a triangle’ is true iff  $\exists x\exists y(Cx \wedge (Ty \wedge Ixy))$   
b.  $\llbracket\text{There is a circle in a triangle}\rrbracket^M = 1 \Leftrightarrow \exists x\exists y(Cx \wedge (Ty \wedge Ixy))$

These are meant to be the same statement as (1.26a). We will explain how to read statements like these with symbols later this textbook, but for now, we will use English as our metalanguage, so the truth-conditions will be stated in the format in (1.26). The object language will be mostly English as well, except in Chapters ?? and ?? where we will discuss other languages and crosslinguistic variation. Admittedly, it is sometimes confusing to use the same language as the object language and metalanguage, but since English is the only common language among us and you readers, we will bite the bullet and stick to it. However, as long as the distinction between the object and metalanguage is clear enough, this should pose no problem.

### 1.5.3 Entailment

Truth-conditions are semantic intuitions about one statement. Based on them, we can define a very important relation between sentences called *entailment*. Its definition is in (1.29).

(1.29) A sentence  $\phi$  *entails* a sentence  $\psi$  iff whenever  $\phi$  is true,  $\psi$  is also true.

Let’s look at some examples. We have intuitions that tell us that the sentence in (1.30a) entails the sentence in (1.30b).

- (1.30) a. There are a lot of happy dogs in Hyde Park.  
b. There are animals in Hyde Park.

Also note that our intuitions tell us that (1.30b) does *not* entail (1.30a). This is because there are situations, real or hypothetical, where (1.30b) is true, but (1.30a) is not true. For example, when there are cats but no dogs in Hyde Park.

Importantly, entailment is about all possible situations, including ones that are not actually realized. Thus, it does not matter whether or not there are actually a lot of dogs in Hyde Park. In other words, whether (1.30a) is actually a true statement is simply irrelevant for entailment. This is more clearly illustrated by the following pair of sentences.

- (1.31) a. The prime minister is listening to pink elephants sing in Hyde Park.  
b. There are pink elephants.

Most certainly neither of these sentences are actually true, but we still have semantic intuitions that (1.31a) entails (1.31b), because if (1.31a) *were* true at all, (1.31b) must also be true.

It should also be stressed that strictly speaking, one cannot *prove* that a sentence entails another sentence, because in order to prove it, one would have to present all possible situations, including hypothetical ones, where the sentences are true, but there are infinitely many such situations. For example, one could come up with many hypothetical situations where (1.30a) is true that can differ with respect to such minute details as exactly how many happy dogs there are, what colors they are, whether or not each of them is running, whether any of them are hungry, and also seemingly irrelevant aspects such as how many people there are in Hyde Park, what the weather is like in London, and even exactly how many people lived in Beijing in 1973.

This means that when a sentence entails another sentence, that is essentially a matter of intuition. On the other hand, when a sentence does *not* entail another sentence, one could actually show it concretely, because in that case, one situation where the first sentence is true, but the second sentence is false is enough to show that the former does not entail the latter. This was already demonstrated for the lack of entailment from (1.30b) and (1.30a) above. Similarly, we can show that (1.31b) does not entail (1.31a) by coming up with an example (hypothetical) situation where (1.31b) is true but (1.31a) is false. For instance, there are many pink elephants in Regent's Park, but none in Hyde Park.

As we have just seen, in each of the above examples, the (a)-sentence entailed the (b)-sentence, but not the other way around. In such a case, we say the (a)-sentence *asymmetrically entails* the (b)-sentence, or the (a)-sentence is *stronger* than the (b)-sentence.

In other cases, entailment holds in both directions. This is obviously the case when the two sentences  $\phi$  and  $\psi$  are the same sentence: Whenever  $\phi$  is true,  $\phi$  is of course true. There are also pairs of distinct sentences that entail each other. This is demonstrated by the following examples.

- (1.32) a. A scientist bought a car from a philosopher.  
b. A philosopher sold a car to a scientist.

When there is a mutual entailment relation between two sentences, that means that the two sentences are true in exactly the same set of situations. This means that the two sentences mean the same thing or are *equivalent*, as far as their truth-conditions are concerned. It should be kept in mind, however, that this does not mean that all aspects of the meanings of the two sentences in (1.32a)

are completely identical, because they might differ with respect to non-truth-conditional aspects. In other words, truth-conditions are just one specific aspect of linguistic meaning.

#### 1.5.4 Pragmatic inferences

In stating truth-conditions and assessing entailments, one needs to be careful with pragmatic inferences. As we discussed in Section 1.2, there are inferences that arise in certain specific contexts but not others, which we illustrated with the example (1.1). In this example, one can infer from Ben's utterance of 'My sister speaks Cantonese' that Ben himself doesn't speak Cantonese, but we saw that there are contexts in which this inference does not arise, meaning that there is a situation where (1.33a) is true, but (1.33b) can be false. Therefore, (1.33a) doesn't entail (1.33b), because entailment is, by definition, about all situations where (1.33a) is true.

- (1.33) a. My sister speaks Cantonese.  
b. Ben does not speak Cantonese.

That the relevant inference is pragmatic in nature might look relatively obvious in this example, but not all cases of pragmatic inferences are as conspicuous. For instance, does (1.34a) entail (1.34b)?

- (1.34) a. Anna has just published a book about bird songs.  
b. Anna has just published exactly one book about bird songs.

In many contexts, an utterance of (1.34a) is used to mean essentially the same thing as (1.34b), as least with respect to the truth-conditions. After all, it is not very common for someone to publish multiple books at the same time, especially if they are about the same topic. But entailment needs to be checked without such plausibility assumptions, because it is about all situations in which the sentence is true, including implausible situations. In other words, entailment is just about the semantic aspect of the sentence with no pragmatics. If one considers a situation where Anna has just published two books about bird songs, the sentence in (1.34a) is still true, although it might be a bit misleading if you are trying to inform someone about one of Anna's books. On the other hand, the sentence (1.34b) is definitely false in that situation. Therefore, there is at least one situation where (1.34a) is true and (1.34b) is false, so (1.34a) does not entail (1.34b), and this means that in context where you can conclude from an utterance of (1.34a) that Anna has just published exactly one book about bird songs, you have some extra pragmatic inference, as your conclusion doesn't follow from the truth-conditions of the sentence alone.

There are even more counter-intuitive cases. For example, consider the following two sentences.

- (1.35) a. Alice speaks French or Italian.

- b. Alice doesn't speak both of French and Italian.

Intuitively, (1.35a) seems to entail (1.35b), but according to many scholars (if not all), this is merely a pragmatic inference—more specifically, an implicature—of an utterance of (1.36a). We will discuss motivation for this seemingly counter-intuitive analysis in greater detail in Chapter ??.

### 1.5.5 Presupposition

Lastly, we mention one interesting universal property of natural language semantics: Not all entailments are the same. In order to see this, let us consider the following example.

(1.36) Greg forgot to call his mother.

This sentence entails a number of sentences, but let us zoom in on the following two.

- (1.37) a. Greg did not call his mother.  
b. Greg was supposed to call his mother.

The sentence in (1.36) entails both of these sentences, as whenever (1.36) is true, these sentences need to be true. However, you might feel that they are of different quality when (1.36) is actually used in a conversation. That is, (1.37a) tends to be the main point of what the speaker wants to convey, while (1.37b) sounds more like what they take for granted to be the case. In fact, when someone is wondering whether (1.37a) is true or false, you can inform them by telling them (1.36), but when someone is wondering if (1.37b) is true or false, it won't be very natural to use the sentence in (1.36). These intuitions become sharper with example dialogues like the following. Here we change the question a little bit to make it more natural to use the phrase *his mother*. Recall that # means that the sentence/utterance is (grammatical but) infelicitous or anomalous.

- (1.38) Q: Did Greg call both his parents?  
A: No, he forgot to call his mother.

- (1.39) Q: Was Greg supposed to call his parents?  
A: #Yes, he forgot to call his mother.

Such entailments that seem to be taken for granted are called presuppositional entailments, or simply *presuppositions*. In the above example, we say that (1.36) presupposes the truth of (1.37b), or that (1.36) presupposes (1.37b) to be true. The non-presupposed type of entailment is variously called in the literature, but let us call them *at-issue* entailments in this book. Another term often used is *asserted* entailments or *assertions*.

These intuitions about presuppositions vs. at-issue entailments correlate with

another set of phenomena, namely, when more complex sentences are created from (1.36), the at-issue entailment to (1.37a) disappear when certain words and phrases are involve, but the presuppositional entailment to (1.37b) tends to stay. Consider, for example, the negative version of (1.36).

(1.40) Greg didn't forget to call his mother.

This sentences does not entail (1.37a). This is not very surprising, because that's exactly what negation is supposed to do, i.e. it denies the meaning of the original sentence, so the negative sentence means that the original sentence is false. However, notice that it still entails (1.37b). That is, whenever this negative sentence is true, (1.37b) still needs to be true. Or to put it differently, negation in English only negates at-issue entailments, and lets presuppositions pass.

The same point can be made with an example containing (1.36) as a conditional antecedent, as in (1.41).

(1.41) If Greg forgot to call his mother, she was unhappy.

This whole sentence does not entail (1.37a), but again, it still entails (1.37b).

Another grammatical context that similarly changes the entailment pattern is *yes/no* questions. To illustrate, let us turn (1.36) into a *yes/no* question as in (1.42).

(1.42) Did Greg forget to call his mother?

This is a question, so the notion of entailment does not apply straightforwardly, as it does not make sense to say that a question is true or false. Crucially, however, what this question is asking is whether (1.37a) is true or false, and not whether (1.37b) is true or false. Rather, whoever asks this question tends to already take (1.37b) to be true. In other words, when a sentence is turned into a question, the question asks about the truth or falsity of the at-issue entailments, but not of the presuppositions, of the original sentence.

This technique of putting a sentence in various grammatical contexts to discern which entailments of the sentence are presuppositions is a very useful test, and is often called the *family of sentences test*. It is advisable to apply multiple constructions like the above to each sentence you want to check, because not every case yields a clear entailment. For instance, consider (1.43).

(1.43) I didn't forget to call my mother! No one expected me to, including her, so I didn't.

You would expect the first sentence of (1.43) to presuppose that (1.44) is true, on a par with the previous example, but it appears that it does not do so in this example.

(1.44) I was supposed to call my mother.

You might notice that (1.43) is most naturally read with an emphatic stress on *forget*, and this prosody is probably a crucial factor that licenses this reading in this particular example. In the general case, however, there might be other factors that could affect the interpretation, and it is best if you check multiple constructions to be sure.

Let us look at another example of a presuppositional entailment.

(1.45) Rachel is in a bar with both of her brothers.

This sentence entails (1.46).

(1.46) Rachel has exactly two brothers.

We now claim that this is a presuppositional entailment. Firstly, indeed, whenever (1.45) is uttered in a conversation, (1.46) would not sound like the main point of the utterance. Secondly, let us apply the family of sentences test, which indeed indicates that the sentences in (1.47) all presuppose the truth of (1.46).

- (1.47)
- a. Rachel is not in a bar with both of her brothers.
  - b. If Rachel is in a bar with both of her brothers, her partner is there too.
  - c. Is Rachel in a bar with both of her brothers?

## 1.6 Chapter summary

We have introduced some core concepts and notions used in the study of meaning. Let us review particularly important ones here.

- **Semantics vs. pragmatics:** Semantics is about the conventional aspect of linguistic meaning, and pragmatics is about the additional inferences that arise when expressions are put in use. However, the distinction between these two is not always clear.
- **Lexical semantics:** Lexical semantics concerns meanings of individual lexical items and their relations.
- **Content vs. function words:** Content words are words that carry the main information of a linguistic utterance, e.g. nouns like *cat* and verbs like *walk*. Function words are grammatical devices that glue expressions together and their meanings are often hard to characterize in intuitive terms, e.g. *the*, *if*, *what*.
- **Truth-conditions:** The truth-conditions of a sentence state when the sentence is true and when the sentence is false. We use the following format:  $\phi$  is true iff \_\_\_\_\_. Here, 'iff' is shorthand for 'if and only if'.



- **Entailment:** Sentence  $\phi$  entails sentence  $\psi$  iff whenever  $\phi$  is true,  $\psi$  is also true.
- **Presupposition:** We say sentence  $\phi$  presupposes the truth of sentence  $\psi$ , when  $\phi$  entails  $\psi$ , and the truth of  $\psi$  tends to be backgrounded when  $\phi$  is uttered. Also, the entailment survives in the family of sentences test.

## 1.7 Further reading

As explained in this chapter, we avoid detailed discussion of concepts in this textbook, but that of course shouldn't stop you from investigating them. If you want to learn about what has been said about concepts and what kind of theoretical questions are currently open, a good place to start is the survey article about concepts in *the Stanford Encyclopedia of Philosophy*, Margolis & Laurence (2021), which contains a number of further references.

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

We introduced several lexical semantic concepts in this chapter, but there is obviously a lot more to lexical semantics than simply defining relations among words. We will not cover advanced topics in lexical semantics in this textbook, but if you wish to learn more about this topic, there is an excellent book, Geeraerts 2010, which gives a historical overview of different theories of lexical semantics.

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

Among different types of ambiguity found in natural language, we mentioned polysemy as a particularly theoretically rich topic. Vicente & Falkum (2017) provides an accessible overview of different theoretical perspectives on this complex topic. Sennet (2021) discusses polysemy and other types of ambiguity more generally.

- Agustí Vicente & Ingrid L. Falkum. 2017. *Polysemy*. <https://doi.org/10.1093/acrefore/9780199384655.013.325>
- Adam Sennet. 2021. Ambiguity. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Summer 2021. Metaphysics Research Lab, Stanford University

Another topic we only briefly mentioned in the main text is politeness and different speech styles, or registers. Most research on these phenomena has been done in the sub-branch of linguistics called sociolinguistics, which studies socio-cultural variation, but in recent years there has been a considerable amount of interaction between sociolinguists and researchers interested in the study of meaning, although such recent research is beyond the scope of this introductory textbook. If you want to learn about major theoretical views about politeness in natural language, Terkourafi (2016) is a recent survey article.

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

Finally, although we confine our attention to declarative sentences in this textbook, there is a considerable amount of work done in semantics and pragmatics about non-declarative sentences. Although this topic tends to be highly technical and should be regarded as an advanced topic, you can get a rough sense of what it is like by reading the survey article by Cross & Roelofsen (2020).

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

## Exercises

Q1. Give a hyponym and a hypernym of each of the following words.

- |              |         |
|--------------|---------|
| i) flower    | iv) go  |
| ii) airplane | v) say  |
| iii) idea    | vi) big |

Q2. Give one pair of synonymous nouns and one pair of synonymous verbs in English, or any other language you know.

Q3. A pair of words that express opposite meanings are called *antonyms* of each other. For instance, ‘truth’ and ‘falsity’ are antonyms.

English has productive ways to construct antonyms with affixes. For example, the prefix ‘in-/im-/il-/ir-’, as in ‘incorrect’, ‘impossible’, ‘illogical’, and ‘irrelevant’, forms antonyms from adjectives like ‘correct’, ‘possible’, ‘logical’, and ‘relevant’. A similar piece of morphology exists for verbs such as ‘un-’, as in ‘unbutton’ and ‘de-’ as in ‘deconstruct’.

The idea of antonymy should be quite intuitive. In this exercise you will discuss how to explicitly define this idea.

- i) As a first pass, let us consider the following definition.

(1.48) Definition 1: A word  $w_1$  is an *antonym* of a word  $w_2$  iff not everything that is describable by  $w_1$  is describable by  $w_2$ .

This is quite similar to how synonymy is defined in (1.8). In fact, (1.48) is its negation. Notice that it will correctly render a pair of antonyms like ‘truth’ and ‘falsity’ as antonyms.

However, it turns out that this is not a good definition because it does not accurately capture the idea of antonymy. This problem is illustrated by a hyponym-hypernym pair, e.g. ‘cat’ and ‘mammal’, which are intuitively not antonyms. Explain in words why this pair constitutes a problem for the definition of antonymy given in (1.48).

- ii) Another way of looking at the problem of (1.48) is that it does not cap-

ture the *symmetric* nature of the antonymy relation. That is, whenever  $w_1$  is an antonym of  $w_2$ ,  $w_2$  should be an antonym of  $w_1$ . But this is not predicted by (1.48).

Let us fix this by changing the definition as follows.

(1.49) *Definition 2:* A word  $w_1$  is an *antonym* of a word  $w_2$  iff not everything that is describable by  $w_1$  is describable by  $w_2$  and not everything that is describable by  $w_2$  is describable by  $w_1$ .

This definition still renders actual antonyms as antonyms, and furthermore, it correctly captures that this relation is symmetric. Thanks to the symmetry, moreover, hyponym-hypernym pairs like ‘cat’ and ‘mammal’ will no longer count as antonyms, as you can easily verify. However, it turns out that this definition still has counter-examples. Find a pair of words that are intuitively not antonyms, but (1.49) will wrongly characterise as antonyms.

- iii) The purpose of this exercise is, on the one hand, to give you an idea of how to evaluate theoretical ideas (in this case definitions of a technical term) against empirical data, and on the other hand, to let you realise that something that you can intuitively understand, such as the idea of antonymy, can turn out to be difficult to clearly state. We will not offer you a definition of antonymy that will actually work here, but try to formulate one of your own, and discuss whether it has counter-examples.

Q4. Consider the dialogue in (1.2) again.

(1.2) Anna: Do you speak Cantonese or Mandarin Chinese?  
Ben: My sister speaks Cantonese.

As discussed in this chapter, the sentence Ben uses in this dialogue has some semantic meaning, but his utterance has additional pragmatic meaning to the effect that Ben himself doesn’t speak Cantonese or Mandarin Chinese. Note that his utterance also implies that his sister doesn’t speak Mandarin Chinese either. Construct an argument that this is also a case of pragmatic meaning, and not semantic meaning.

Q5. For each of the following pairs of sentences, determine whether the first sentence entails the second. If the answer is no, justify your answer by describing an example situation where the first sentence is true, but the second sentence is false.

- i) Alice does not live in Paris.  
Alice lives in France.
- ii) Becky is tall.  
Becky is not short.
- iii) Cate is a doctor.  
Cate is not a lawyer.
- iv) Dave said that two is an even number.

Two is an even number.

- v) Ellen is considered to be a genius.  
Ellen is a genius.
- vi) Fred got fired and moved to Tahiti.  
Fred moved to Tahiti and got fired.
- vii) George knows that the earth is round.  
The earth is round.
- viii) Harry saw the man with binoculars.  
Harry had binoculars.
- ix) Not all polar bears live in the UK.  
At least one polar bear lives in the UK.

Q6. When two sentences cannot be true simultaneously, we say they *contradict* each other.

(1.50) A sentence  $\phi$  contradicts a sentence  $\psi$  iff  $\phi$  and  $\psi$  cannot be true simultaneously.

Discuss which of the following pairs of sentences contradict each other.

- i) Andrew lives with two brown dogs.  
There are no dogs at all.
- ii) Becky drives a BMW.  
Becky owns a car.
- iii) Chris is a tall guy.  
Chris was a short boy.
- iv) Dave is not British.  
Dave is not European.
- v) Three linguists are not in this room.  
There are more than two people who are linguists.
- vi) Most candidates are female.  
There is only one female candidate.
- vii) All unicorns have one horn.  
There are unicorns.

Q7. Apply the family of sentences test to test if the first sentence of each of the following pairs of sentences presupposes the truth of the second. Write down the sentences you checked and what conclusions you drew from them.

- i) There is a mosquito in the bedroom again.  
There was a mosquito in the bedroom before.
- ii) There is a mosquito in the bedroom again.  
There is a mosquito.
- iii) There is a mosquito in the bedroom again.  
There is one or more bedrooms.

- iv) It is three expensive cars that were stolen from the famous linguist's garage.  
There are three cars.
- v) It is three expensive cars that were stolen from the famous linguist's garage.  
There is exactly one famous linguist.
- vi) It is three expensive cars that were stolen from the famous linguist's garage.  
The famous linguist has a garage.
- vii) It is three expensive cars that were stolen from the famous linguist's garage.  
More than one item was stolen from the famous linguist's garage.
- viii) Every chef is happy.  
There are more than two chefs.

Q8. An adjective is said to be *autological* if it is correctly predicated of itself. For example, 'English' and 'short' are autological, because the adjective 'English' is (an) English (word) and the adjective 'short' is short. We say that an adjective is called *heterological* if it is not autological. For example, 'French' and 'long' are heterological, because 'French' is not (a) French (word), and 'long' is not long.

- i) Find another autological adjective.
- ii) Discuss whether the adjective 'autological' is autological or heterological.
- iii) Find another heterological adjective.
- iv) Discuss whether the adjective 'heterological' is autological or heterological.



# Bibliography

- Alxatib, Sam & Francis Jeffrey 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 vagueness*. 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](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, Lluís 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*, Summer 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-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.