1 Introduction: Compositional Semantics for Natural Language

Chapter Outline

1.1 How to study semantics of natural language

This textbook introduces the reader to natural language semantics. We will adopt English as the primary object language or the target of our investigation. However, we will also examine Japanese data when it is substantially different from English and its examination is theoretically crucial.

We shall start with a simple assumption: we use language to exchange information and ideas with others. The speaker utters a sentence to convey some thought to the hearer. When this information is decoded correctly by the hearer, communication is considered successful.

Source: Emma / Adobe Stock Photos.

We wish to characterize this process through formal means. For this purpose, we adopt the basic idea that pieces of language (i.e., linguistic expressions, such as words and phrases) refer to external objects (e.g., people, animals, physical objects). This idea is called **a referential (or denotational) theory of meaning**. Consider (1), for example.

(1) Mary: (Pointing at a dog) That's Fido.

(1) enables the addressee to understand that the dog being pointed at is named *Fido*. A connection is established in the mind of the hearer between the dog in question and the name *Fido*, which is a word in English. This example illustrates that language is frequently employed to establish a connection between a word (in this case, a name) and an entity in the external world. Once the connection is established for the speaker as well as for the hearer, a name can be used without an associated pointing action as in (2).

(2) Fido is a husky.

When someone utters (2), they assume that the addressee knows a specific dog named *Fido* and then claim that the dog in question is a husky, which is a dog breed.[1] The part of a sentence that states something about the subject is called a **predicate**.

For our purposes, the term **predicate** is a cover term for verbs, adjectives, nouns, and prepositions that describe events or states within a sentence.

We have just presented a critical hypothesis according to which the subject of a sentence denotes an external entity, and the predicate denotes a property, which is then assigned to the subject-denoted entity. Although this may seem obvious, it is important to recognize that it is by this process that someone can understand the meaning of a sentence like (2). This process, which involves expressing the meaning of a complex expression by combining the meanings of its parts, illustrates **the principle of compositionality**. This principle can be shown with a syntactic tree, as in (3).

(3)

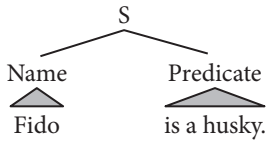

As the syntactic tree in (3) shows, the sentence consists of a Name (*Fido*) and a Predicate (*is a husky*). We shall henceforth ignore the verb *be* and the indefinite article *a* and say that the predicate is simply *husky*.[2] This will be explained more formally in Section 1.3, but the idea is that the verb *be* and the article *a* do not carry relevant semantic information in (3). We assume that these expressions (*Fido* and *husky*) are listed in the mental lexicon (a kind of dictionary in the brain) with their semantic specifications. We then need some mechanism that merges the individual meanings of each part to yield the meaning of the entire sentence.

Based on a referential theory of meaning, we need to specify more explicitly what *husky* denotes. The meaning of *husky* in a referential theory of meaning is the exhaustive list of those individuals that are huskies. The entities that are huskies are concrete entities and not linguistic expressions. Thus, the idea that *husky* denotes these individuals is consistent with the referential theory of meaning. Now that we have an idea of what *Fido* and *husky* mean, we should be able to combine them to obtain the meaning of the entire sentence. In this case, the dog Fido must be among those that are huskies for the sentence to be true. The details of how this procedure works in a formal system are discussed in Section 1.3.

We will use the nouns *reference* and *denotation* interchangeably. Likewise, the verbal expressions *refer (to)* and *denote* are understood to be used in the same way, relating linguistic expressions to external entities serving as their meanings.

1.2 Motivations for adopting a referential theory of meaning

The referential theory of meaning allows us to make semantics more tractable by assuming that meanings are concrete entities in the external world. In this account, meanings are not vague images or thoughts that we have in the brain. This also means that the referential theory of meaning lets us avoid circularity: if we define the meaning of a word in terms of other words, we will never understand the meaning of anything. By making the distinction between linguistic expressions and their meanings explicit, we can avoid this problem because linguistic expressions are said to denote external entities.

In addition, the referential theory of meaning enables us to explain why children can acquire the semantic component of language. Because a child is born with no established vocabulary in any language, we cannot teach them the meaning of a specific expression by using other expressions in the same language. Despite this problem, children are remarkably successful learners of their first language(s). They learn the meaning of each word without relying on dictionary definitions – they learn language directly by using it. It is reasonable to claim that the basic strategy they rely on is analogous to a referential theory of meaning, at least for the first few years of their life. For example, a child's caretaker might see the full moon and say, "Look at the moon! It is beautiful." This type of utterance by adults is likely to establish a firm link in the mind of the child between the word *moon* and the yellowish celestial object in the night sky. The same idea can explain how a child learns the meaning of a common noun, such as *dog*. An adult might point at a dog and say, "That's a cute little dog," or "Watch out. That dog is big and scary." The child might collect such specimens of dogs and form a concept from such examples. Essentially, this is the referential theory of meaning in action. Thus, constructing a semantic theory that is based on the referential theory of meaning has this added advantage and motivation. By constructing a semantic theory based on the idea that words we use refer to the external world, we may be able to shed light on why the acquisition of the semantic component of language is possible. Although a more advanced stage of language acquisition may involve words conveying abstract concepts (such as "*idea*," "*theory*," and "*honesty*"), these abstract concepts may originate from metaphorical extensions of more concrete instances of meaning. Many language acquisition researchers, such as

Gentner (1982) and Gillette et al. (1999), observe that children acquire words referring to concrete entities before they tackle abstract ones.[3]

1.3 Compositional semantics: first steps

Let us now illustrate how a **compositional semantic theory** works based on the idea that the meaning of a word is an entity or entities in the external world. First, we shall introduce **the double bracket notation** ⟦⟧, which indicates the meaning of the expression placed within a pair of double brackets. Regarding (2), which is repeated here as (4), we first show that each word has a meaning different from these words themselves. The interpretation of (4) is determined compositionally from the meanings of the parts as shown in (5).

(4) Fido is a husky.

(5) a. ⟦Fido⟧ = 

b. ⟦husky⟧ = 

Source: 5a: Alef Morais / Pexels; 5b (left to right): Alex Sever / Pexels, Alef Morais / Pexels, Joerg Hartmann / Pexels, Helena Lopes / Pexels and Julissa Helmuth / Pexels.

(5a) shows that the name *Fido* refers to the dog depicted by a picture on the right side of the equation symbol =. (5b) says that the common noun *husky* denotes (i.e., means) the collection of all the husky dogs whose pictures are shown here.

We are now ready to combine these two informal meanings. This process corresponds to the comprehension of what (4) means when we hear someone utter this sentence. Consider (4), for example. Our semantic theory assumes that the meaning of a sentence is either *true* or *false*. More colloquially, we say that (4) is true or that it is false. We can say that it is true when we find a state of affairs in the actual world that is correctly described by (4). Pretheoretically, we might say that (4) is true when the dog Fido is a specimen of the dog breed husky. In our theory, we state that the denotation of *Fido*, the dog pictured to the right of the equation symbol in (5a), is included in the denotation of the common noun *husky*, which consists of the dogs pictured to the right of the equation symbol in (5b). Tentatively, we can propose the semantic rule of "predication" that unites the meanings of the subject name and the predicate in (6). This rule includes the expression *if and only if*, abbreviated as *iff*, which occurs frequently in semantic discussions and is worth explaining. When you see A iff B, this means that A is satisfied when B is, and vice versa.[4]

(6) 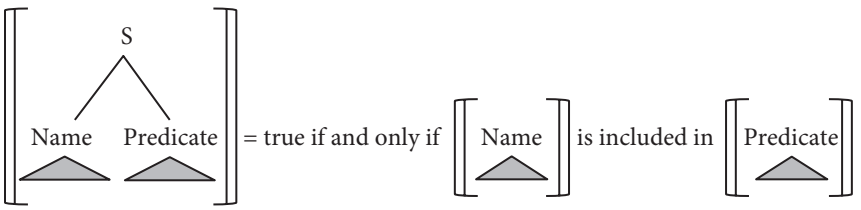

By looking at (5a) and (5b), we find that the dog Fido is indeed included in the meaning (= the denotation) of the common noun *husky*. Thus, we can conclude that (4) is true.

Let us now consider (7).

(7) Max is asleep.

For example, if we find a situation depicted by the picture in (8) and we know that the name of this creature is Max, then we can say that (7) is true.

(8) 

Source: FU / Pexels.

On the other hand, if you see the scene in (9) involving the same dog, you will conclude that (7) is false.

(9) 

Source: Rob Annen / Pexels.

In other words, when we say that (7) is true, we mean that there is a piece of reality that it correctly describes. By contrast, when we say that (7) is false, there is no state of affairs that corresponds to what the sentence describes. Even though the concepts of truth and falsity are not directly observable, we use information about the external world to make these truth-value judgments. Thus, even though we assume that the meaning of a sentence is a truth value, which is admittedly an abstract entity, we are working within a referential semantic theory.

Technically speaking, we assume that the adjective *asleep* also denotes concrete entities, just like the noun *husky*. In the case of *asleep*, its denotation is those entities that are (currently) asleep. This is illustrated by (10). (7) is true when you find the dog Max in the meaning of *asleep* here. Assuming that the denotation of *asleep* is accurate, this is tantamount to checking whether Max has the required property for the truth of (7). Needless to say, the list of those who are asleep changes over time. Thus, (10) should be understood to be the list of those who are asleep now. When we expand the coverage of our investigation to tense and modality in Chapter 8 and onwards, we will need to modify this simple idea about meaning and the use of $[\![\,]\!]$.

(10) $[\![\text{asleep}]\!]$ = 

In more general terms, the meaning of a sentence (independently of uttering it) can be thought of as its **truth conditions**. The truth conditions of a sentence are defined as the conditions that need to be satisfied for the sentence to be true. The reader should recognize that the truth conditions for a sentence do not have to correspond to one specific situation that is mutually exclusive with others. Consider (11).

(11) Mary is happy.

Let us assume that the speaker and the hearer know Mary. The predicate *happy* intuitively describes a positive feeling on the part of someone. Here's one possibility: Mary is having dinner with her friends. She is having a pleasant conversation and is smiling. She is visibly happy. In another possible scenario, she is alone at home watching a movie with a big-screen TV. She is absorbed in the story and is very attentive. She is not talking or smiling but is very happy internally. Note that these are very different situations involving Mary. However, in both situations, (11) is true because its truth conditions are satisfied. This demonstrates that the truth conditions of a sentence can be satisfied in various ways and cannot always be determined solely through visual means. However, we can show that the sentence meaning is firmly linked to a piece of reality in the following sense: if (11) is true, then this means that there is an actual state of affairs that shows that its truth conditions are satisfied.

The reader should also note that the meanings of two distinct sentences can only be distinguished if we consider their truth conditions, because truth values alone are not sufficient to assign different meanings to different sentences. Consider (12a, b).

(12) a. The Earth revolves around the Sun.
b. The Morning Star is the Evening Star.

Given what we know about the Solar System, both (12a) and (12b) are true. However, they are not synonymous. Even if the Earth revolves around the Sun, it is possible that the Morning Star is not the same entity as the Evening Star. Thus, it is possible to imagine a situation where (12a) is true, but (12b) is not. Analogously, it is possible for (12b) to be true without (12a) being true. Thus, we should assume that the intuitive meaning of (12a) or (12b) corresponds to their truth conditions rather than their truth values. Details of this idea will be discussed in the subsequent chapters, especially in Chapters 8 and 9.

Exercise 1.1

State or describe the truth conditions of (i) through (iii).

i. Mary is in New York.
ii. The room is empty.
iii. Jamie has a blue car.

1.4 Our official syntax-semantics interface system in a nutshell

The analysis of (4) provided in Section 1.3 assumes that natural language sentences are interpreted directly: each expression is assumed to have a meaning. This method was adopted to show the fundamentals of referential semantic theory and truth conditions. However, the actual system we adopt in the rest of the book is more involved.

In this section, we present our overall syntax-semantics interface system, enabling the formal material in Chapter 2 to be understood within the context of natural language semantics. We will adopt a logical language that resembles Montague's (1973) Intensional Logic as our intermediate language. That is, each expression in natural language (English or Japanese in our case) is translated into the logical language first. Each logical expression in the translation is then interpreted as an "external" non-linguistic entity, which is specified in set-theoretic terms.

We shall present a simple example sentence in English and show how it is analyzed in our syntax-semantics interface system. Let us examine (2), repeated here as (13), with its provisional syntactic structure.

(13)

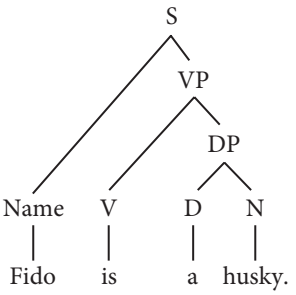

As shown in the tree, the predicate VP is syntactically complex; it consists of the verb *be* and a DP (determiner phrase) consisting of the indefinite article *a* and the common noun *husky*. Despite this complex structure of the predicate, the predicate is just *husky*.[5] Thus, ignoring the verb *be* and the article *a*, (13) is translated into our logical language as in (14).

(14)

1.	Fido ⇒ **Fido**	name
2.	is a husky ⇒ **Husky**	predicate
3.	Fido is a husky. ⇒ **Husky**(**Fido**)	sentence

The arrow ⇒ is used to indicate how a natural language expression is translated into a logical expression of various types. Line 1 shows that the name *Fido* translates into **Fido** in the logical language. There is no surprise here. However, to distinguish between natural language expressions and logical expressions, all logical expressions are bolded, as shown in (14). In Line 2, *husky* is the only part of the VP that is semantically relevant, and it is translated as **Husky**. The first letter of a predicate is capitalized. As shown in Line 3, in this logical language, the predicate is written first, followed by its argument, which is enclosed within parentheses.

We then interpret the sentence given in Line 3 of (14). The predicate **Husky** denotes the collection of all huskies in the entire world, and **Fido** denotes the dog Fido. This is very similar to what we said in (5), except that our official system applies the double bracket notation ⟦⟧ to logical expressions, rather than to natural language expressions. (15) shows how the actual semantic interpretation is obtained from Line 3 of (14).

To be more precise, **Husky** denotes the characteristic function of the set of huskies, as explained below.

(15) ⟦**Husky**(**Fido**)⟧ = ⟦**Husky**⟧(⟦**Fido**⟧)

Let me describe how to read (15). ⟦**Husky**(**Fido**)⟧ is the meaning or denotation of the sentence **Husky**(**Fido**), and this is obtained by calculating ⟦**Husky**⟧(⟦**Fido**⟧). This is written with the so-called algebraic notation, which is exemplified by (16).

(16)

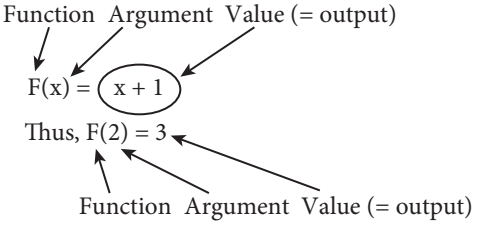

The first line of (16) defines the function F by showing that for any number x, F(x) equals x + 1. F(x) shows that F is the function, and x is its argument. It reads, "F applies to x and yields x + 1." The second line shows a concrete example of how F is used: F applies to 2 in this case. Since 2+1 = 3, the output is 3. We will use the same notation, but the entities involved are not numbers but semantic entities such as people, animals, tables, etc. We need ⟦⟧ to distinguish between expressions and their denotations.

Given the algebraic notation of functions, ⟦**Husky**⟧(⟦**Fido**⟧) is understood to mean that ⟦**Husky**⟧ is the function and ⟦**Fido**⟧ is its argument. ⟦**Husky**⟧ is the so-called characteristic function of the set of huskies.[6] Note that we have shown what Fido looks like in (5a), but let us repeat the info here in (17).

(17) ⟦**Fido**⟧ = 

Source: Alef Morais / Pexels.

The function ⟦**Husky**⟧ is represented graphically in (18).

(18) The meaning of the predicate **Husky** in our logical language

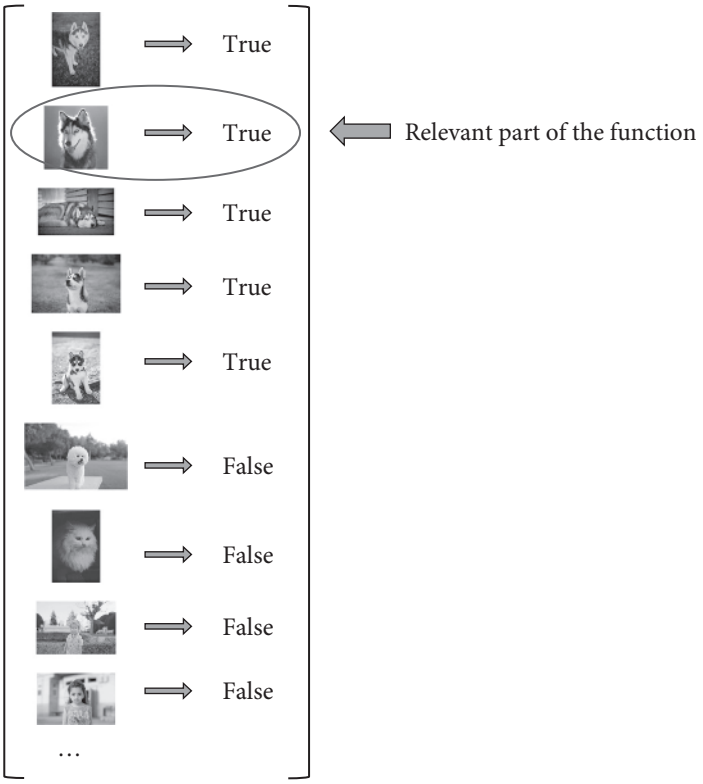

Source: (top to bottom): Alex Sever / Pexels, Alef Morais / Pexels, Joerg Hartmann / Pexels, Helena Lopes / Pexels, Julissa Helmuth / Pexels, David Brown / Pexels, Dionys Bell / Pexels, Cosmin Turbatu / Pexels and Memo Axndv / Pexels.

(18) is a partial representation of ⟦**Husky**⟧. In this representation, the left-side column lists possible inputs, and the right-side column provides the matching outputs (or "values") of this function. Note that there are only two possible values shown on the right-hand side column: True and False. Specifically, the characteristic function (18) is used to classify inputs into two different categories: those that are huskies and those that are not. The output True means that the input is a husky; the output False means that the input is not a husky. Thus, (13), repeated here as (19), requires that we look at the second line in (18), which is circled and shows that the dog Fido is mapped to True. This means that the sentence (13) is true.

(19) Translation: Fido is a husky ⇒ **Husky(Fido)**

Semantic interpretation: ⟦**Husky(Fido)**⟧ = ⟦**Husky**⟧(⟦**Fido**⟧) = True

Function Argument Value (= output)

We have shown that in our system, the predicate Husky itself denotes the characteristic function of the set of huskies. **Husky** can be paraphrased as **λx[Husky(x)]** in our logical language. This representation employs the λ-operator and is analyzed as in (20).

(20) **λx** **[Husky(x)]**

Input Output is True if ⟦**Husky**(x)⟧ = True when x is instantiated by the input

Output is False if ⟦**Husky**(x)⟧ = False when x is instantiated by the input

This will be explained in Chapter 2 in more detail. **λx** shows that x represents a possible input, and the output is determined by whether the input object makes [**Husky**(x)] true. That is, if x is a husky, the sentence **Husky(x)** is true; if not, it is false. In other words, it is semantically equivalent to **Husky**.

This completes the discussion of *Fido is a husky*. The purpose of this discussion is to show very briefly how the whole syntax-semantics interface system works in this textbook. This should make the discussion of Set Theory and Intensional Logic in Chapter 2 more meaningful for the reader. The details of what we discussed here will be examined in the subsequent chapters.

Exercise 1.2

Let us ensure that the reader understands how the lambda expression is utilized in our semantic system. Let us assume that we have three entities {Fido, Max, Leo}, and Leo is a poodle, but Fido and Max are not. ⟦λx[**Poodle**(x)]⟧ is the function F that maps any individual x to True iff x is a poodle.

Describe the function F imitating (18). It should look like this:

$$\begin{bmatrix} \text{Fido} \rightarrow \ \ldots \\ \text{Max} \rightarrow \ \ldots \\ \text{Leo} \rightarrow \ \ldots \end{bmatrix}$$

1.5 Some important syntactic concepts

Some important syntactic concepts are introduced in this section. We will begin with the configurational notions of **dominance**, **direct dominance, and c-command**. We will explain these notions as we simultaneously compare and contrast the structures of the English sentence in (21) and the Japanese sentence in (22). For practical purposes, these sentences are synonymous with each other. In the syntactic trees in (21a, b), NP stands for Noun Phrase, and VP stands for Verb

Phrase. S stands for Sentence, and V stands for Verb. Generally, for any lexical category X, XP is a phrase that contains X.

(21) a.

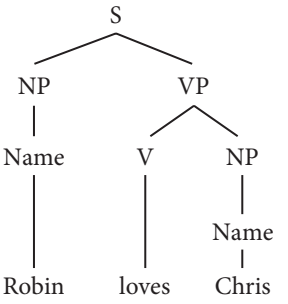

b.

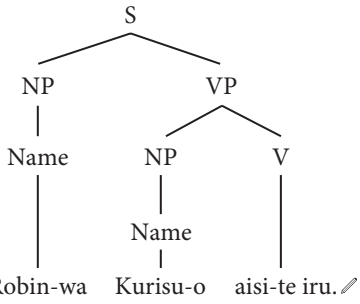

Robin-TOP Chris-ACC love.PRES

'Robin loves Chris.'

Japanese examples are given in a phonemic transliteration, while glosses reflect phonetic form where relevant.

In a syntactic structure, when a constituent A is higher than another constituent B, and one can draw a continuous downward line from A to B, A is said to **dominate** B. When there is no intervening node between A and B, A is said to **directly dominate** B. We often invoke the term **mother node** to describe this situation: A is the mother node for B. For example, in (21a) and (21b), S directly dominates the subject NP; S also dominates the object NP, but not directly.

We can now define the concept of **c-command** as follows: a constituent A is said to c-command B if and only if there is a constituent C that dominates B and directly dominates A. For example, in (21a, b), the subject NP c-commands the object NP, but not vice versa. This is indeed how Chomsky defines the subject NP and the object NP in his early work (Chomsky 1965): the subject is the NP that is directly dominated by the S node, and the object is the NP directly dominated by the VP node. These definitions apply equally well to English and Japanese. The concept of c-command is a way of encoding relative syntactic hierarchy, and it will be useful when we discuss quantifier scope in Chapter 7.

Note also that A does not dominate B, nor does B dominate A.

Exercise 1.3

In the following configuration, which nodes does A directly dominate? Are there any other nodes that A dominates (but not directly)?

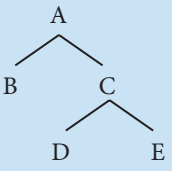

1.6 How we use the Japanese data

This book approaches the semantics of natural language by drawing on two typologically distinct languages: English and Japanese. This approach has many important implications. One is that our semantic theory should be able to deal with natural languages in general, not just English. This is an obvious point, but this perspective should be put into practice. That is, we should examine data from a variety of typologically disparate languages and construct a theory of natural language semantics that takes into consideration a diverse set of data from them. Examining two typologically distinct languages, such as English and Japanese, is a good starting point, and this textbook does just that. I hope that this will invite the reader to reflect on their native language (especially if it is not English) and start working on its semantics. To present the essential semantic concepts and techniques as efficiently as possible, we will discuss Japanese data only when Japanese presents some crucial data that leads to a theoretical point that cannot be captured by English data alone. If English and Japanese data does not present essential differences regarding a particular issue, we discuss the English data alone, assuming that its analysis applies equally well to Japanese.

1.7 Typological differences between English and Japanese

As mentioned in the previous section, we will use Japanese data when differences between English and Japanese are crucial for explaining theoretical claims. Needless to say, there are many differences between English and Japanese that we cannot address in this introductory textbook. We merely disregard them for the sake of efficiency. We do not assume that the reader has any knowledge of Japanese. Therefore, we need to introduce some essential characteristics of this language. Let's examine some key characteristics of Japanese.

The first significant difference between English and Japanese involves headedness. The expression that defines the syntactic category of the entire phrase is called the **head** of the phrase. For example, a VP has a verb as its head, and a PP has a preposition as its head. An English phrase generally has its head at the beginning of the phrase, whereas a Japanese phrase generally has its head at the end. Hence, English is a head-initial language, whereas Japanese is a head-final language. One straightforward example of this difference in headedness is illustrated by English PPs (prepositional phrases) and Japanese PPs (postpositional phrases). This is exemplified in (22a, b).

(22) a. Mary sent a book to Tokyo.
b. Mearii-wa hon-o Tookyoo-ni okut-ta.
Mary-TOP book-ACC Tokyo-to send-PAST

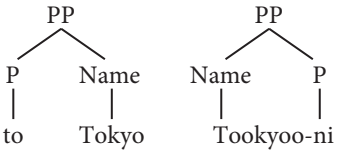

As can be seen in (22a), the preposition *to* occurs before the noun *Tokyo* in English. (22b) shows the Japanese equivalent where the postposition *ni* "to" occurs after the noun *Tookyoo*. The difference in headedness is not a major difference for semantic purposes (at least for simple examples) since our overall theory of semantics does not require a particular order when two expressions are combined to yield the meaning of a larger structure. This point will be discussed in more detail in the rest of this textbook.

Keen readers might see that the difference in headedness is closely related to the word order difference between English and Japanese. If we assume that a sentence contains a VP as a constituent,

English should have a VP with its head (a verb) as the first element, and Japanese should have a VP with its head occurring as the final element. This was shown in (21a, b). The VP portions in (21a, b) are given in (23a, b).

(23) a.

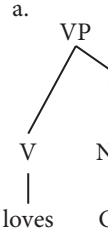

b.

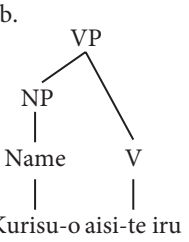

Note that in English the word order within the VP shown in (23a) is VO (verb – object), whereas in Japanese the word order within the VP shown in (23b) is OV (object – verb). Together with the subject that is external to the VP, the entire sentence is SVO in English and SOV in Japanese.

Second, there is a significant difference between English and Japanese in their nominal structure. Japanese has no morphological distinction between count and mass nouns. That is, nouns require neither plural marking nor determiners. In other words, all nouns can occur as bare nouns. Compare (24a) and (24b).

(24) a. * Mary has dog.
b. Hanako-wa inu-o kat-te iru.
Hanako-TOP dog-ACC keep-PROG.PRES
"Hanako has (literally, keeps) a dog"

Note that when *a* is missing from the object of have as in (24a), the sentence is ungrammatical. By contrast, the Japanese sentence (24b) is fully grammatical. It has no marker on the noun for number, nor does it have an indefinite or definite marker. It is not clear whether Hanako has just one dog or multiple dogs. The fact that Japanese lacks definiteness markers and obligatory number marking means that the Japanese nominal structure has fewer requirements than the English nominal structure. English features morphological distinctions between count and mass nouns. When a count noun occurs in the singular form, it must co-occur with a determiner or some other delimiter. When a count noun indicates plurality, it must be accompanied by a marker, such as *-s*. By contrast, Japanese common nouns can occur as bare nouns without determiner-like expressions, even when the intended interpretation is singular and/or definite.[7] This difference between English and Japanese has many significant semantic consequences, which will be discussed in Chapter 13.

Let us touch on a couple of additional differences between English and Japanese. Japanese has internally-headed relative clauses (IHRCs), which resemble nominalized clauses but can refer to external entities, such as apples and oranges. This is impossible in English. Here is an example.

(25) Hanako-wa [Ziroo-ga ringo-o mui-ta]-no-o tabe-ta.
Hanako-TOP Jiro-NOM apple-ACC peel-PAST-NMZ-ACC eat-PAST
"Hanako ate an apple that Jiro peeled."
[Literal] "Hanako ate (the fact) that Jiro peeled an apple."

This will be discussed in detail in Chapter 6.

English and Japanese also differ from each other concerning modal expressions. English has modal auxiliaries like *must* and *may*, but Japanese uses more complex structures.

For example, to express the idea that one is obliged to do something, Japanese employs double negation as shown in (26).

(26) Ika-na-kereba, nara-nai.
go-NEG-if be.good-NEG.PRES
"(I) have to go."
[Lit.] "If (I) do not go, it won't be good."

This topic will be discussed in Chapter 8.

Exercise 1.4

Consider the Malagasy sentence in (i).

i. Mamaky boky ny mpianatra.
reads book the student
"The student reads the book."

Based on the explanation of different patterns of word order (e.g., SOV, SVO, etc.) in the preceding section, what is likely to be the basic word order in this language?

1.8 Semantics and Pragmatics

In this section, we shall discuss some possible issues with the bivalent (two-valued) semantic theory we assume for the rest of the book. Before proceeding, however, I should make clear what a bivalent system is. Recall that a sentence is either true or false depending on whether it satisfies the truth conditions for it. For example, (27) is either true or false at any particular given time, such as the time of writing this sentence.

(27) It is raining (now) in New York.

If it is raining in New York now, it is true. If it is not, then it is false. The hearer may not have access to the necessary information to know whether it is true. However, the speaker has access to the truth conditions of (27) if they speak English. They know that if the conditions hold, (27) is true. Otherwise, it is false. In this way, it is arguable that any given sentence has only two values to choose from if the context is fixed. This system is bivalent regarding sentence meaning.

This system may encounter problems when examining more complex natural language examples. Consider the sentence in (28).

(28) Mary knows that the Sun revolves around the Earth.

Assume, naturally, that the Earth revolves around the Sun, not the other way around. Given this understanding, (28) is not true. However, would it be correct to say that it is false? This does not seem to be the case since it would be odd for someone to react to (28) by saying, "No, she doesn't." Our intuition is that the speaker of (28) has a more fundamental misunderstanding about the relation between the Sun and the Earth. It may be more appropriate to respond by saying, "What do you mean? The Sun does not revolve around the Earth. It's the other way around." In addition, if

(28) was false, then negating the matrix verb should make the sentence true. However, this is intuitively not the case, as shown by the fact that (29) is not true either.

(29) Mary does not know that the Sun revolves around the Earth.

It is arguable that (29) lacks a truth value because it also seems to take for granted that the Sun revolves around the Earth. This is an instance of **presupposition failure.** This issue arises from the idea that some expressions (*know*, in this case) carry presuppositions. The verb *know* presupposes that its complement clause (i.e., the embedded clause) is true. Thus, since the Earth revolves around the Sun, (30) is coherent and is true.

(30) Liz knows that the Earth revolves around the Sun.

However, when we deal with a sentence that has the verb *know* in the main predicate position and a false complement clause, we face the problem of presupposition failure as exemplified by (28). Since (28) is a grammatical sentence, we would need to assign a truth value to it in a bivalent system. However, it would feel wrong to assign it a truth value: it is definitely not true, but it is also not false. The problems associated with presuppositions are significant. However, we will not discuss this issue in the rest of this textbook to keep our book manageable for someone new to semantics.

Another topic that we will not discuss in this textbook is conversational implicature. Consider the conversational exchange between Bill and Chris in (31).

(31) Bill: Can you give me a ride to the airport tomorrow morning?
Chris: I have a job interview early in the morning. Sorry.

In (31), Bill asks whether Chris can give him a ride to the airport tomorrow morning. Chris does not answer this question directly. Instead, he says that he has a job interview early in the morning (tomorrow). This implies that the answer is no. According to Grice (1975), this can be explained by a rule of conversation that requires what one says to be relevant. Since what Chris says must be appropriate to the question Bill has just asked, it must be an indirect answer to the question. Chris says that he has a job interview tomorrow morning; this suggests that he will not be available then. Therefore, Chris's response is an indirect negative answer to the question. This type of reasoning goes beyond what truth-conditional semantic theory can address. We will not discuss Grice's theory of conversational maxims.

The above discussion demonstrates that English (or any other natural language) has numerous meaning-related properties that this textbook will not address. The primary purpose of this textbook is to teach the basic principles of the standard truth-conditional semantic theory, and we will maintain its consistency and simplicity. The ultimate goal of this introductory text is to make sure that a beginning learner understands the basic data and how to approach semantic theory.

Exercise 1.5

In each sentence, specify its presupposition.

i. Jamie regrets having become a plumber.
ii. Robin stopped smoking.
iii. Alex took the math exam again.

Exercise 1.6

Specify a plausible conversational implicature of B's response to A's question.

A: Is Mary home?
B: Well, her car is not in the carport.

Notes

1 The Japanese sentence that corresponds to (2) is (i).

i. Faidoo-wa hasukii-da.
Fido-TOP husky-be.PRES
"Fido is a husky."

We do not deal with the Japanese case separately here since there is no significant difference between (2) and (i).

2 The verb *be*, even though it is obligatory in a sentence like (2), does not appear to have any semantic role to play here and is ignored in our syntax-semantics interface system. The copular verb *be* actually carries tense information, and may also convey some substantial semantic information in other cases, but such subtleties are ignored here for the sake of simplicity of presentation. In some languages such as Russian and Japanese, adjectives can occur with no accompanying verb. Consider (i) (Russian). Japanese examples will be discussed in Chapter 3.

i. Ya schastliv.
I happy
"I am happy."

The indefinite article that precedes the common noun husky is also understood to be semantically empty.

3 I thank Naja Ferjan Ramirez (personal communication) for bringing these references to my attention.

4 The expression "is included in" is vague, but this will be made precise in a fully formal version of our semantic theory.

5 More details will be provided in Chapter 3 as to how these two expressions can be "ignored" in the semantics.

6 The set-theoretic concept of function is also explained in more detail in Chapter 2. Here, we are providing an example so that the reader can understand how the concept is used in the actual example.

7 Japanese has optional plural markers for some nouns. For example, for children (plural), one could use *kodomo-tati* "child-plural." However, this marking is not obligatory and is not simply a plural marker, either. For example, this morpheme can also occur with a name as in *Hanako-tati*, which means "the group of people including Hanako."

References

Chomsky, N. (1965). *Aspects of The Theory of Syntax*. Cambridge, MA: MIT Press.

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In: S.A. Kuczaj, ed. *Language Development: Volume 2. Language, thought and culture*. Hillsdale, NJ: Erlbaum, pp. 301–334.

Gillette, J., Gleitman, H., Gleitman, L., and Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, 73(2), pp. 135–176.

Grice, H.P. (1975). Logic and conversation. In: P. Cole and J.L. Morgan, eds. *Syntax and Semantics, Volume 3: Speech Acts*. New York: Academic Press, pp. 41–58.

Montague, R. (1973). The proper treatment of quantification in ordinary English. In: J. Hintikka, J.M.E. Moravcsik, and P. Suppes, eds. *Approaches to Natural Language*. Dordrecht: Reidel, pp. 221–242.

Other Introductory Semantics Textbooks

Altshuler, D., Parsons, T., and Schwarzschild, R. (2019). *A Course in Semantics*. Cambridge, MA: MIT Press.

Chierchia, G. and McConnell-Ginet, S. (2000). *Meaning and Grammar: An Introduction to Semantics*. Cambridge, MA: MIT Press.

Dowty, D.R., Wall, R.E., and Peters, S. (1981). *Introduction to Montague Semantics*. Dordrecht: D. Reidel.

Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Malden: Blackwell Publishers.

Kearns, K. (2011). *Semantics*. 2nd ed. London: Red Globe Press.

Larson, R.K. and Segal, G. (1995). *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: MIT Press.

Larson, R.K. (2022). *Semantics as Science*. Cambridge, MA: MIT Press.

de Swart, H. (1998). *Introduction to Natural Language Semantics*. Stanford, CA: CSLI Publications.

Zimmermann, T.E. and Sternefeld, W. (2013). *Introduction to Semantics*. Berlin: Walter de Gruyter.