Title	Two ways to study meaning : semantics and pragmatics : chapter 1 of the textbook Introduction to semantics for non-native speakers of English
Sub Title	
Author	Tancredi, Christopher
Publisher	慶應義塾大学言語文化研究所
Publication year	2021
Jtitle	慶應義塾大学言語文化研究所紀要 (Reports of the Keio Institute of Cultural and Linguistic Studies). No.52 (2021. 3) ,p.163- 177
JaLC DOI	10.14991/005.00000052-0163
Abstract	
Notes	研究ノート
Genre	Departmental Bulletin Paper
URL	https://koara.lib.keio.ac.jp/xoonips/modules/xoonips/detail.php?koara_id=AN00069467-00000052-0163

Reports of the Keio Institute of Cultural and Linguistic Studies 52 (2021), 163~177

Two Ways to Study Meaning: Semantics and Pragmatics

Chapter 1 of the textbook

Introduction to Semantics for Non-native Speakers of English

Christopher Tancredi

Introduction

This book is an introduction to semantics. But what is semantics? A dictionary will tell you that semantics is the study of meaning. But that only makes things clear if we know what *meaning* is. The words *mean* and *meaning* describe many very different things. Only some of those things are included in the study of semantics. Consider examples.

(A) a. The **meaning** of the word "dog" is "four-legged animal of the canine family".
b. "Physician" **means** the same thing as "doctor".
c. The sentence "The physician is a puppy" **means** that the unique salient doctor is a young four-legged animal of the canine family.

(B) a. By "The physician is a puppy", I **mean** that the doctor is cute and harmless.
b. Those clouds **mean** rain.
c. I understand what she said, but I don't know what she **means**.
d. The true **meaning** of a rule only becomes clear when people break it.

The (A) examples are about the literal meaning of expressions in a language, here English. The (B) examples are about what significance things have. The study of

semantics includes meaning as illustrated in the (A) cases but not that in the (B) cases. Its aim is to describe and explain the literal meanings of expressions in language.

Semantics and Pragmatics

In studying meaning in language, it is important to separate the literal, grammatical meaning of a sentence from what it can be used to communicate. This difference can be seen in (Ac) and (Ba) above. Both are about the sentence "The physician is a puppy". (Ac) focuses on the literal meaning of the sentence. This is a meaning that very few people would want to communicate and that is almost certainly false. (Ba), in contrast, focuses on a message that is easily communicated. That message, however, is not the literal meaning of that sentence. The literal meaning seen in (Ac) we call the *semantic meaning* of the sentence, and the message communicated in (Ba) we call a *pragmatic meaning*.

Though *semantics* and *pragmatics* are both studies of meaning, they look at meaning in different ways.

Semantics: The study of grammatical meaning.
Pragmatics: The study of the meanings people communicate to others using language.

Sometimes it is unclear whether a *meaning* is part of the semantics or part of the pragmatics. One way to decide is to ask whether the meaning changes when the context changes. If so, that part of the meaning is likely pragmatic. If not, it is more likely semantic. Consider an example adapted from Kate Kearns' textbook *Semantics*.

(C) I forgot the paper.

The semantic meaning of the sentence is that there is a salient time in the past at which the speaker of the sentence forgot the uniquely salient paper. Pragmatically, however, the sentence can communicate very different things. Consider (C) uttered by Betty in the following situations:

Situation I:
Betty goes out to buy coffee and the morning newspaper. She comes back with only coffee. (C)
Situation II:
Betty investigates a murder. The victim has a scrap of wallpaper stuck to his shoe. Betty finds many other clues, and comes up with a hypothesis about what happened, but then realizes it doesn't account for the scrap of wallpaper. (C)

In Situation I, we understand *forgot* to mean *forgot to buy*, and we understand *the paper* to mean *the newspaper*. In Situation II, in contrast, we understand *forgot* to mean *forgot to take into consideration*, and *the paper* to mean *the scrap of wallpaper on the victim's shoe*. Since our understanding of what was *forgotten* and of what kind of *paper* was involved depends on the context, this understanding is likely to come from the pragmatics. Both cases involve forgetting, though, as well as a unique salient paper. Since these parts do not change with the change of context, they are likely to come from the semantics.

Determining how the message that is communicated depends on the context is an important part of understanding language. A full theory of linguistics needs to include an explanation of this dependence. The reason for bringing issues of communication up here, however, is to be able to put them aside. They belong mostly to the study of pragmatics, not semantics. We will look at some pragmatic effects in a little more detail in Chapter 2, but for the rest of this chapter and the rest of the book after Chapter 2 we focus on semantics.

Semantics

Sense and Denotation

Recall that we are taking semantics to be the study of grammatical meaning. More precisely, we take semantics to pair expressions of a language with meanings, where an **expression** is *a word*, *phrase or sentence*.

Words:	*dog*, *the*, *flew*
Phrases:	*the US President*, *in Yoyogi Park*
Sentences:	*John walked*, *I'm happy today*

Sense

Before looking at how this pairing works, it is first necessary to ask what kind of a thing meanings are. Two notions are closely associated with meaning. The first is the **sense** of an expression. This is roughly what you know when you know the meaning of an expression. For individual words, *sense* is like what a dictionary gives you, though complex expressions have senses too.

> *dog*: a four-legged animal of the canine family.
>
> *the US President*: The unique individual who leads the administrative branch of the US government and commands the US armed forces.
>
> *John walked*: At some time in the past, an individual named *John* moved by alternatively transferring his weight to one leg (or set of legs) and advancing the other leg (or legs).

According to Gottlob Frege, who invented the concept in the late 19th century, *senses*, like the languages they are part of, are things in the world that are independent of people's understanding. The sense of the word *dog* is not something we can see or touch, but for Frege it is as real as an actual dog. People can grasp a sense, either fully or only partly, but what they have in their minds when grasping a sense is not itself a sense. It is an **idea**. In the 1950s, Noam Chomsky argued that language is a property of individual minds. Public languages like English, for Chomsky, do not exist separately from the speakers who speak it. The properties that a public language has come from the properties of the speakers' individual languages. If we accept Chomsky's view of language, then Fregean *ideas* become the basic building blocks of meaning rather than Fregean *senses*. Below, I will use the term *sense* to mean either a Fregean sense within a public language or a Fregean idea within an individual's language.

Denotation

The second notion associated with meaning is the **denotation** of an expression. This is the thing(s) or situation(s) that an expression picks out, or is true of. For some expressions, like names and definite descriptions (expressions starting with the word

the), what is picked out is often obvious. The name *Tokyo*, for example, picks out the city Tokyo. So does the definite description *the capital of Japan*. In other cases, we find it odd to consider expressions as picking out anything at all. What, for example, does *in Yoyogi Park* pick out? We certainly don't normally use that expression to refer to things. However, we do take it to *apply to*, or to *be true of* certain things, namely of all those things that are in Yoyogi Park. In such cases, we take the *denotation* of the expression to be the set of things that the expression is true of. The case of sentences is even less intuitive than the case of expressions like *in Yoyogi Park*. Sentences, too, have denotations. However, **declarative** sentences -- sentences that make a declaration, or that state something that is either true or false -- have one of only two values as their denotation: either True, if the sentence is true, or False, if the sentence is false. Examples of these three classes of expressions and their denotations are given below.

The capital of Japan:	Tokyo
Tokyo:	Tokyo
in Yoyogi Park:	The objects, people, animals, etc. in Yoyogi Park.
John walked:	True/False (True if John walked, False if John didn't walk).

A third term used to talk about meaning is **reference**. Some people use the word *reference* to mean the same as the word *denotation*. In this book, *reference* is used only for expressions that intuitively refer to things. These expressions are usually names like "Tokyo" and "Shakespeare" or definite descriptions like "the capital of Japan". If an expression has a reference, its reference is also its denotation. Many expressions, however, such as sentences, verbs and preposition phrases, have a denotation but not a reference.

Relating Sense and Denotation

While *sense* is more closely related to how people think about things, *denotation* is more closely related to the things themselves. Of the two, *sense* is closer to what people intuitively think of as meaning. For giving a formal analysis of meaning, however, it turns out that *denotation* is the more useful concept. The two concepts are

connected: *sense* determines *denotation*. That is, knowing the sense of an expression makes it possible to pick out its denotation. Knowing the denotation of an expression, however, does not make it possible to pick out its sense.

To see this connection, consider the following examples:

(D)	expression	denotation
	The capital of Japan	Tokyo
	The city in Japan with the largest area	Tokyo
	The city in Japan with the largest population	Tokyo

The three expressions each pick out, or **denote**, the same thing, namely the city Tokyo. They pick it out in different ways, though. Those different ways can be thought of as different paths to the denotation. Each path is like a separate sense. Choosing any of the three paths described in (D) will get you to the denotation *Tokyo*. Knowing only what the denotation is, however, does not tell you which path was followed to get there.

Compositionality

Having identified sense and denotation as two things we want from meaning, the next question is how to assign the correct sense and denotation to the expressions of a language. For individual words our options are limited. Some words can be defined in terms of other words. *Physician*, for example, can be defined as meaning the same thing as *doctor*. However, this will only give us a sense and denotation for *physician* if we already have a sense and denotation for *doctor*. We obviously could not get these by defining *doctor* as meaning the same thing as *physician*. That would make our definitions circular. We could get them by defining *doctor* using other expressions, but this will only help if these other expressions have an independently identifiable sense and denotation. At some point, the process of defining expressions using other expressions needs to come to an end. At this point, we need an independent way of assigning a sense and denotation to the expressions used in the final definitions. Call these expressions the semantically basic expressions of a language.

Basic Expressions

How do we assign a sense and denotation to a basic expression? It is no good trying

to give a definition to describe the sense. If the definition is accurate and if it actually does give the sense, then this shows that the expression is not a basic expression. This makes it very difficult to identify senses clearly. Usually, the best we can do is give hints about what the sense is. Could we instead give the denotation? For a noun like *cat* we can certainly point to objects that are in the denotation. But we cannot point to the entire denotation. That would require being able to point to every cat in the world throughout history.

If senses and denotations cannot be clearly identified, why do we think they are important parts of meaning? We can get an idea why by considering a simplified case of language learning. Children learn the meanings of at least some basic expressions (nouns like *ball*, *cat* and *table*, for example) by connecting the expressions to particular objects. That connection can be made without involving sense or denotation. The word can at that point be like a name for that particular object. Once a child has several objects connected to one word, however, the word can no longer work like the name of a single object. This pushes the child to find another way that the expression can work. Noticing that all the things called *cat* have properties in common, the child associates the word *cat* directly with these properties. The child then uses these properties to pick out objects. Since many objects have these same properties, the expression then comes to pick out many more objects than the child has yet seen. In this way, a basic expression can change from working like a name of a particular object to working like a word that picks out a set of properties true of many objects. The properties act like the sense of the expression, while the many objects this sense picks out are the denotation of the expression. Of the three processes given above -- naming an object, finding properties shared by several objects, and using those properties to pick out all objects that have them -- only the first is taught. The other two processes are hidden from sight. It is these hidden processes, though, that supply us with an internal sense and a denotation for our basic expressions.

Complex Expressions

In principle we could treat every expression as a basic expression. However, that would mean having to be taught the meaning of every expression when we first hear it. This would make it difficult to explain our ability to create new expressions. Most of the sentences we hear and produce are new to us, and yet we rarely have any difficulty understanding them and using them. We certainly do not need to be taught

their meanings one by one in order to do so. We must have a different way of figuring out those meanings.

To anyone who knows a language, it is obvious what that different way is. We build up the meaning of the sentence by combining, or **composing**, the meanings of the words that make it up. The details of how this works will make up a large part of this book. However, we can show one basic property of this process of **composition**, namely that the meaning of a complex expression depends on more than just the meanings of the words it is made from. We can show this by creating two sentences with different meanings out of the same words. Take the words *the*, *dog*, *cat* and *chased* for example. From these words we can create the following two sentences:

The dog chased the cat.
The cat chased the dog.

The difference in meaning cannot come from the choice of words since these are the same for the two sentences. It rather has to come from how the words are put together.

The study of syntax, or phrase structure, has shown us that what matters in putting words together into sentences is not the order in which the words occur but rather the way they are structured into more complex expressions. In both of the above sentences, the words *the* and *dog* go together to make the complex expression *[the dog]*, and *the* and *cat* combine to make *[the cat]*. In the first sentence, *chased* combines with *[the cat]* to produce the verb phrase *[chased [the cat]]*, and this verb phrase combines with *[the dog]* to produce the full structured sentence *[[the dog] [chased [the cat]]]*. In the second sentence, the expressions combine differently, producing the structure *[[the cat] [chased [the dog]]]*. The semantics then needs to show how this difference in structure affects meaning.

The principle that connects structure and meaning is the **principle of compositionality**.

> **Principle of Compositionality:** When two expressions combine in the syntax to make up a larger expression, their meanings compose to make up a more complex meaning.

When the expressions being combined are words, their meanings have to be given separately, either as basic expressions or as words defined from basic expressions. When the parts themselves are complex, the meanings of the parts have to be calculated first before being composed into the meaning of the whole.

Explaining the Unbounded Nature of Language

The principle of compositionality helps us explain how we understand new expressions. It also helps us explain why there is no upper limit on how many sentences can be generated (built) and understood. This **infinite generative capacity** of language comes from the syntax allowing structures to contain other structures of the same type. This makes syntactic structures **recursive**, and makes it possible in principle for structures to contain other structures that contain other structures, and on and on without limit. The ability to understand these structures comes from the compositionality of semantics. As long as the structures generated by the syntax are well formed, the semantics will assign them meanings based on the meanings of their parts and how they go together.

We can see recursion in simple examples like the following:

> *[The dog barked at [the cat that chased [the rat that ate [the cheese that was left on [the table that was built by [the man that ...]]]]]]*

This is a case of **unbounded recursion**. We could continue repeating this process of embedding one structure inside another in the same way for as long as we want. There are other cases of recursion that are **bounded**. So-called center embedding is one of them, illustrated below. The process that generates the first two sentences cannot be repeated. Doing so generates the third structure, but this is not a sentence of English, or at least not an understandable one.

> *[[The cat [that the dog barked at]] chased the rat]*
> *[[The rat [that the cat chased]] ate the cheese]*
> * *[[The rat [[that the cat [that the dog barked at]] chased]] ate the cheese]*

Types of expressions

There are many different expressions in language, but they can be put into a fairly

small number of groups depending on what type of expression they are. Here we will focus on four such groups. The first group contains **names**. A name is an expression that picks out an object or individual but does not describe it. A name can name a person, like *Shinzo Abe*, a place, like *Tokyo*, or something more abstract such as a group, as with *The Beatles*.

Expressions in the second group, **definite descriptions**, sometimes function like names, as with *The Beatles* or *The Netherlands*, but more often pick out an object or individual as the unique thing having a particular property. *The tallest man living in Paris*, for example, picks out one person, but unlike a name, it does so by picking out whoever has a particular property that no one else has, namely the property of being tallest among men living in Paris.

The third group contains **1-place predicates**. These are expressions that describe objects. Typical examples include adjectives like *blue*, which describes all things that are blue; intransitive verbs like *sleep*, that describes all things that sleep; and nouns like *dog*, that describes all things that are dogs.

The fourth group contains **sentences**, which are expressions that describe situations. The sentence *Shinzo Abe is sleeping* describes situations in which Shinzo Abe is sleeping, while *Tokyo is in Russia* describes situations in which Tokyo is in Russia. Sometimes, as in this second case, the only situations that are described by a sentence are unreal situations.

From Denotations to Intensions

Earlier it was claimed that expressions have a *sense* and a *denotation*. Identifying senses, though, was seen to be difficult. In particular, for basic expressions we had no way of saying what their sense is. This means that if we tried to base our semantics on senses, our semantics would lack a solid foundation. Can we base our semantics on denotations? True, we cannot identify the denotations of many expressions either. However, denotations often include things that we can identify and that all speakers can agree are in the denotation -- individual cats, for example. In this way, denotations give us a much more solid foundation to build a theory of semantics upon.

While denotations play a role in meaning, it is clear that they cannot do everything

that we want from a theory of semantics. Among other things, we want our semantics to distinguish between the meanings of expressions that are intuitively felt to differ. Denotations do a terrible job of that. This is clearest in the case of sentences. All true sentences have the same denotation -- *True* -- and yet intuitively they do not all have the same meaning. Even if we put sentences aside, problems remain. Consider the 1-place predicates *animal with a heart* and *animal with a kidney*. Suppose it turns out that the animals with a heart are exactly the same as the animals with a kidney. Then these two expressions have the same denotation. Intuitively, however, they do not have the same meaning. This once again shows that denotations cannot count as meanings.

Possible Worlds

The problems just seen can be overcome by considering denotations in **possible worlds** in addition to in the actual world. As we will see below, this is enough to distinguish among intuitively distinct sentences that are true in the actual world as well as among intuitively distinct 1-place predicates that happen to have the same denotation at the actual world.

What is a possible world? The easiest way to think of a possible world is as a rearrangement of things in the actual world, with the possibility of adding objects or taking objects away. In the actual world, everything is as it is. My computer is on my dining room table, my reading glasses are to the left of the computer, and behind the computer is a bottle of mayonnaise, waiting to be put back in the fridge. One possible world would be a world in which everything is like it is in the actual world except that my reading glasses are to the right of the computer instead of to the left. In another possible world, my reading glasses are behind the computer and the mayonnaise is to the left of the computer. In still another possible world, there are no reading glasses on the table. Other possible worlds differ more from the actual world: worlds in which people do not exist, worlds containing black holes but no stars, worlds containing nothing but 2 pieces of dust.

To see how possible worlds help us to distinguish expressions, consider again the expressions *animal with a heart* and *animal with a kidney*. You were asked to imagine that these expressions have the same denotation in the actual world. Even if this is so, we can imagine a possible world in which some animal – perhaps a camel, or perhaps

some animal that doesn't exist in the actual world, like a unicorn – has a heart but no kidney. In that world, the denotation of *animal with a heart* will include this animal, but the denotation of *animal with a kidney* will not. If we collect together the denotations of our two expressions at all possible worlds, the two collections will differ. The collections as a whole will then distinguish between expressions intuitively felt to differ in meaning.

Extensions and Intensions

We work this idea out by distinguishing two versions of "meaning". The first, the **extension** of an expression, is the denotation of that expression at a specific world. If the world is not identified, it is taken to be the actual world. The second, the **intension** of an expression, is a set of pairings of a world *w* with the extension of the expression at *w*.

> **Extension** of an expression *E* at world *w*: The denotation of *E* in *w*.
> **Intension** of an expression *E*: The set of pairs <*w*, *e*> such that *w* is a possible world, and *e* is the extension of *E* at *w*.

We illustrate these concepts below with expressions of the four types considered above, based on the following assumptions:

$w_@$ = the actual world
$w_1, w_2, \ldots$ = possible worlds

Consider first names. The name *Donald Trump* names Donald Trump. Who does the name pick out at other worlds? We have to be careful about what we mean here. There are other worlds in which someone else was given the name *Donald Trump*. Suppose that in w_1, the name *Donald Trump* was given to the actual world Barack Obama, and the actual world Donald Trump was named *Fred Flintstone*. From within w_1, we would say that the name *Donald Trump* names Donald Trump, but the person we would point to as Donald Trump would be the person we identify in the actual world as Barack Obama. So which individual do we associate with the name *Donald Trump* at world w_1, the actual world Donald Trump or the actual world Barack Obama?

Since our semantics is meant to explain meanings in our language as we actually

speak it, we take the name *Donald Trump* to pick out the actual world Donald Trump at w_1. That is why I say ***at*** w_1 rather than ***in*** w_1. The fact that this person is called *Fred Flintstone* ***in*** w_1 doesn't matter. The connection between a name and the individual it names is a part of our language. Changing that connection in other possible worlds then counts as changing the language. Since our semantics needs to describe our actual language, not a language we speak in some other possible world, we have to keep that connection the same. The name *Donald Trump*, then, picks out the same individual at every world where that individual exists. The extension and intension of the name are as follows:

Donald Trump:
Extension at $w_@$: Donald Trump
Intension: {<$w_@$, Donald Trump>,
<w_1, Donald Trump>,
<w_2, Donald Trump>, …}

Consider next definite descriptions. In the actual world, the US president in 2020 is Donald Trump. However, we can imagine other possible worlds, such as w_1 in which it is Barack Obama, w_2 in which it is Hilary Clinton, and many others. In this case, the extension of the expression *the US president in 2020* at the actual world is Donald Trump, while the intension is an infinite set of pairs, pairing each possible world with the individual who is the US president in 2020 at that world:

The US president in 2020:
Extension at $w_@$: Donald Trump
Intension: {<$w_@$, Donald Trump>,
<w_1, Barack Obama>,
<w_2, Hilary Clinton>, …}

Notice that though the individuals picked out as the US president in 2020 differ from world to world, we use our actual world understanding of the expression *the US president in 2020* to pick those individuals out.

Consider next the 1-place predicate *happy*. To simplify things, imagine that there are only 3 individuals in the actual world $w_@$ and in the two other worlds w_1 and w_2 that

we use for illustration, namely Donald Trump, Barack Obama and Hilary Clinton. In $w_@$, Donald Trump and Barack Obama are happy, but Hilary Clinton is not. In w_1, all three are happy. In w_2, nobody is happy. In this case, the extension of the word happy at $w_@$ is the set containing Donald Trump and Barack Obama, while the intension is the set pairing worlds with the set of people happy in those worlds:

Happy:
Extension at $w_@$: {Donald Trump, Barack Obama}
Intension: {<$w_@$, {Donald Trump, Barack Obama }>,
<w_1, {Donald Trump, Barack Obama, Hilary Clinton}>,
<w_2, { }>, …}

Notice that the set of individuals paired with w2 has no members. We call this set the **empty set**.

Finally, consider the sentence *The US president in 2020 is happy*. If $w_@$, w_1 and w_2 are as described above for the illustration of definite descriptions and 1-place predicates, then this sentence is true at $w_@$ because in $w_@$, Donald Trump is the US president in 2020 and Donald Trump is happy. It is also true at w_1 because in w_1, Barack Obama is the US president in 2020 and Barack Obama is happy. Finally, it is false at w_2 because in w_2, Hilary Clinton is the US president in 2020 but Hilary Clinton is not happy. This gives us the following as the extension and intension for the sentence:

The US president in 2020 is happy:
Extension at $w_@$: True
Intension: {<$w_@$, True>,
<w_1, True>,
<w_2, False>, …}

Intensions bring us closer to our intuitive notion of meaning than extensions do. However, they are still not as close as senses. For a declarative sentence, the sense determines a truth value by giving us **truth conditions**. These are the conditions that make a sentence true when they are satisfied, and false when they are not satisfied. Intuitively, for the sentence *The US president in 2020 is happy*, the truth conditions involve one individual being US president in 2020 and being happy. While the truth

conditions of the sentence determine the intension, however, the intension does not determine the truth conditions. We cannot recover a specific set of truth conditions from the intension {<$w_{@}$, True>, <w_1, True>, <w_2, False>, …}, for instance. The intension is rather the result of what the truth conditions say about each and every possible world.