The Study of Meaning

An Introduction to Semantics and Pragmatics

Nathan Klinedinst

Yasutada Sudo

University College London

Chapter 3

Vagueness

In the previous chapter, we used Set Theory to analyse the meanings of nouns, verbs and adjectives in model-theoretic terms. For instance, the extension of the noun 'door' is the set of all doors in a given (possibly unreal) situation, and likewise, the extension of the adjective 'rectangular' is the set of all rectangular entities in the situation.

However, for certain words such as 'big', this strategy might not work very well. This is because in many real life situations, it does not seem completely clear what should count as big and what should count as not big. Certainly, there can be things that are obviously big, and there can be things that are certainly not big. But more often than not there are also mid-sized things that fall between these two extremes. Should such things be members of the extension of 'big', or not?

This is the issue of *vagueness*, which is the main topic of this chapter. It is in fact one of the oldest issues in natural language semantics, dating back to Ancient Greece, and many different theories have since been proposed. Consequently, a thorough review of the literature on vagueness would easily require more than a chapter, perhaps a book on its own. Instead, we will focus on one formal approach to vagueness, namely the *fuzzy set approach*, and examine its theoretical predictions against empirical data, referring the interested reader simply to the articles and books are mentioned in the Further Reading section at the end of the chapter. This will provide us with a concrete example of how to evaluate a formal theory of meaning by identifying its predictions, and assessing them with empirical data.

3.1 Vagueness in Natural Language

Vagueness is ubiquitous in natural language. It is often illustrated with *gradable adjectives* like 'big', presumably because their meanings are obviously vague, but as we will see later in this section, expressions of other categories also give rise to vagueness.

3.1.1 Gradable adjectives

Gradable adjectives are those adjectives that are compatible with expressions like 'very', 'too', 'enough', and can appear in comparative and superlative constructions. These expressions and constructions intuitively have meanings that have to do with extents to which a given property holds. The examples below demonstrate that 'cheap' and 'smart' are gradable adjectives.

(3.1) a. That laptop is too cheap.
b. That laptop is cheaper than this tablet.
c. That laptop is the cheapest.

(3.2) a. She is too smart.
b. She is smarter than he is.
c. She is the smartest.

While many adjectives in English are gradable, not all are. For instance, 'mandatory' is not, as all of the following are infelicitous. As mentioned before, infelicity is standardly indicated in linguistics by '#'.

(3.3) a. #This course is too mandatory.
b. #This course is more mandatory than that one.
c. #This course is the most mandatory.

However, one important caveat here is that one might be able to construe 'mandatory' as a gradable adjective in certain conversational contexts and/or in certain sentences. For instance, (3.4) sounds acceptable.

(3.4) Vaccination should be more mandatory than it is now.

Such cases might force us to say that 'mandatory' has both gradable and non-gradable uses. Similar remarks apply to the adjective 'prime'. Specifically, sentences like (3.5) suggest that 'prime' qualifies as a non-gradable adjectives.

(3.5) a. #She is a very prime/too prime a minister.
b. #She is the more prime minister.
c. #She is the most prime minister.

But again, one can find sentences and/or contexts in which 'prime' exhibits properties of gradable adjectives, as shown in (3.6), which is acceptable.

(3.6) This is a very prime location.

In this case one might argue that 'prime' has different meanings in 'prime minister' and 'prime location', and the former is a non-gradable adjective and the latter is a gradable adjective. This is certainly an analytical possibility, but it should also be acknowledged that it is not obvious if this adjective is ambiguous. If one were to postulate such an ambiguity in meaning, one should seek for independent evidence for the ambiguity in question (See Ch. **??** for more discussion on ambiguity).

It is not the purpose of this textbook to settle all these empirical matters, but there are several things to learn from these observations.

Firstly, despite some of the complications we've just seen, it is clear that adjectives like 'cheap' and 'smart' are gradable adjectives. The unclear cases have to do with alleged non-gradable adjectives like 'mandatory'.

Secondly, for such unclear cases, we can at least say that they have both non-gradable and gradable *uses*, while leaving open whether or not these uses are to be accounted for in terms of lexical ambiguity. But one should still wonder if there is any adjective that is only usable as a non-gradable adjective. Adjectives that are used in technical ways such as ones that pertain to mathematical concepts like 'square', 'prime', 'even', etc. are normally non-gradable, but even these ones allow for gradable uses in some (non-technical) cases. We are actually not aware of a convincing case of an adjective that only has a non-gradable use, but we encourage you to try to find one. Generally, constructing theoretically relevant linguistic examples is neither easy nor trivial, and is a skill that you need to develop through practice, in order to become a good linguist.

Finally, the fact that many cases of what look like non-gradable adjectives can also be used as gradable adjectives might be showing us something fundamental about natural language. That is, the distinction between gradable and non-gradable adjectives seems to be a quite flexible one, and it is unlikely that a given adjective is simply lexically specified to be either once and for all. This gives rise to a number of further questions that one could investigate. For one, the above data involving 'mandatory' and 'prime' could be taken as suggesting that whether they are used as non-gradable or gradable adjectives is determined by surrounding linguistic material. Alternatively, what matters might be something about the conversational context these sentences are used in and is non-linguistic in nature. Or, maybe both of these two factors matter. We will not try to decide which of these is the best hypothesis, but we would like to stress the general lesson to be learned here: Empirical data in natural language is often messy and one's theory might not clearly match observations, but one should not see such data as trouble or an issue to be avoided. Rather, it often tells us something deep about the nature of natural language, and provides us with opportunities to direct our research in further directions.

3.1.2 Relative adjectives and context-sensitivity

As we saw above, gradable adjectives can be modified by expressions like 'very', a comparative marker, superlative marker, and so on, but it is also true that they don't need to co-occur with such expressions, at least in English.[1] This is

[1] It has been observed that in Mandarin Chinese, a sentence like (3.7) requires 'hěn' in front of the adjective. Since this is absent when other modifiers are available, it is not part of the

demonstrated in (3.7). The unmodified form of the adjective is often called its *positive form*.

(3.7) This is cheap.

Gradable adjectives like ‘cheap’ are called *relative adjectives*, because in their positive form, their interpretation is relative to some contextually determined standard. Concretely, in the case of (3.7), the sentence is true when the price of the referent of the subject ‘this’ is below some threshold price, a contextually determined standard. Due to the inherent vagueness, which we will discuss in detail below, we cannot be very precise about the standard and say something like ‘If the price is below £12.53, then it is cheap’. But crucially, the standard seems to vary across contexts, and is not fixed by the meaning of the three words used in the sentence in (3.7). Rather, it seems that it is determined every time this sentence is used, by various factors of the conversational context.

One of the factors that matters for determining the standard for (3.7) is who the speaker is. For example, if the speaker is a billionaire, they could probably truthfully utter the sentence in (3.7), while talking about, say, a hotel room in London that costs £300 a night. But the authors of this book personally do not consider this to be cheap, so we cannot truthfully speak of the same hotel room in the same way. So, for the billionaire, the standard is somewhere high, maybe around £500 a night (or maybe higher?), while for us it is much lower, probably around £50, and this is not determined by the semantics of the expressions used in (3.7).

Another such contextual factor that affects the standard is what object we are talking about. We would say a hotel room that costs £20 a night is cheap, but if we are talking about a hamburger that costs £20, it is definitely not cheap at all.

Thus, in order to evaluate the truth and falsity of a sentence like (3.7) that contains a relative adjective in its positive form, one first needs to fix the standard by inspecting relevant aspects of the conversational context against which the sentence is to be understood. In this sense, relative adjectives are *context-sensitive*.

Relative adjectives are abundant in English. Other examples of relative adjectives include, ‘fast’, ‘smart’ ‘big’, ‘tall’, ‘heavy’, etc. It should b easy to verify the context-sensitivity of these adjectives in the same way as above. For example, for ‘It is fast’, it clearly matters what the subject of the sentence ‘it’ refers to. Specifically, the standard is very different, when ‘it’ refers to a car and when ‘it’ refers to a snail. Similarly, we can demonstrate that the standard can vary depending on who the speaker is. Suppose that ‘it’ refers to a computer from 1990 that was a high-end, expensive model when it was released. From our current perspective, such a computer is not fast at all, but for people in 1990, it was. We encourage you to think of such example contexts that indicate the

gradable adjective. This suggests that positive forms of adjectives cannot be used on their own in Mandarin Chinese, at least in simple constructions like (3.7).

context-sensitivity of the other relative adjectives listed above as well.

Unlike relative adjectives, other adjectives like 'bent' have more or less fixed standards. Firstly, it is easy to demonstrate that *bent* is a gradable adjective.

(3.8) a. The rod is very bent.
b. This rod is more bent than that one.
c. This rod is the most bent.

But unlike the gradable adjectives we have seen so far, the standard of the positive form of *bent* in a sentence like (3.9) does not seem to vary across contexts.

(3.9) It is bent.

That is, (3.9) is true as long as the referent of 'it' is not straight, and it does not seem to matter what the referent of 'it' is and who the speaker is. Rather, as soon as the referent of 'it' is not straight, it is bent. Such adjectives with clear and fixed standards are called *absolute adjectives.*

However, one might wonder if 'bent' is always an absolute adjective. That is, there is maybe a sense in which the standard of 'bent' depends on the context. For instance, if we are talking about a ruler used to draw straight lines, any slight curve should make it bent, but if we are talking about the handle bar of a bike, we are likely to be more tolerant, and regard a slightly bent handle bar as straight. Moreover, two people could disagree about whether a twig is bent or not, perhaps depending on whether they deem it to be useful for whatever purpose they have in mind. This fact suggests that speakers could have in mind different standards for *bent.*

Thus, it is not clear if we can categorise gradable adjectives completely clearly into relative adjectives and absolute adjectives. In other words, we again find ourselves dealing with linguistic data that might not be as clear as we might wish. However, to repeat the lesson from above, this is an opportunity for us to be curious about what is going on, instead of feeling pessimistic about our theory. One analytical possibility here is that the contrast between relative and absolute adjectives is a flexible one, and an absolute adjective like 'bent' can sometimes be used as a relative adjective. According to this idea, then, the situation is reminiscent of the case of gradable vs. non-gradable adjectives: Some adjectives have non-gradable uses but they can also be used as gradable adjectives in some contexts. However, there is another possibility in this case. We might as well get rid of the categorical distinction between relative and absolute adjectives altogether, and assume that every gradable adjective is in fact a relative adjective, and the difference between adjectives like 'cheap' and adjectives like 'bent' is merely a matter of degree. That is, 'cheap' happens to have a standard that is, for some reason yet to be understood, more easily and robustly affected by various contextual factors than the standard for 'bent'. As classification of adjectives is not the main topic of this chapter, we will leave this question open here. Please see the Further Reading section at the end of this chapter, in case you want to read more about this and related topics.

3.1.3 Context-sensitivity and vagueness

We have just seen that relative adjectives like 'cheap' and 'fast' are context-sensitive in the sense that their standards vary from context to context. But there is more than mere context-sensitivity in their semantics, namely, their standards are never precise. For 'cheap', for example, we can never find a context where we know for sure that everything under some particular price, say £10.00, is cheap and everything above that is not cheap. Likewise, we don't seem to use 'fast' in any context to talk about something moving at a higher velocity than some specific speed. The standards of relative adjectives always seem to be unnamable.

This point can be made even clearer, when we consider the fact that context-sensitivity could be removed, or at least reduced to a significant degree, by certain linguistic means. Take the example sentence in (3.10).

(3.10) This boy is tall.

This is context-sensitive and vague, similarly to the other examples we discussed above. That is, whether (3.10) is true or false depends on certain non-linguist factors, such as who else we are considering, among others. Specifically, in a context where we are talking about five-year-old children at some kindergarten, for example, the standard of who counts as tall is relative to this group of children, so the sentence is true of a boy who is, say, 120 cm tall. On the other hand, if we are looking for a tall person who can replace the lightbulb of a ceiling light, the same sentence about the same boy can well be judged as false.

We can remove, or at least reduce, this type of context-sensitivity by using a *for*-phrase as in (3.11).

(3.11) This boy is tall for a five-year-old British boy.

The semantic contribution of the *for*-phrase in this example is to fix the standard to the one that is relative to five-year-old British boys, so the sentence in (3.11) is suitable for the former context mentioned above, but not for the latter one. Nonetheless, however, the standard is still vague in the sense that we are unable to name exactly where the threshold for being tall vs. being not tall is.

What we have just observed, namely, that we can tamper with context-sensitivity without affecting vagueness itself, suggests that the context-sensitivity of the standard and its vagueness are two different things. In fact, there are adjectives in English that have context-sensitive standards but are not vague. For instance, in order to determine whether the sentence in (3.12) is true of false, we need to know two things, namely, who the referent of 'he' is and what properties that person needs to have to count as 'qualified'. Note that neither of these are determined by the meanings of the linguistic expressions, but contextually, and in this sense the meaning of this sentence is context-sensitive.

(3.12) He is qualified.

Crucially, however, the standard of this adjective can be non-vague, at least in some contexts, and in that case we can know exactly what is needed for the sentence to be true. Concretely, suppose that the sentence is used in a context of job application and that all that is required for one to be qualified for the job in question is that one have a college degree. This standard is not vague. Thus, this is a case where the standard is context-sensitive but not vague.

It should be noticed that this example also involves a non-gradable use of the adjective 'qualified'. As with other examples of non-gradable adjectives, the same adjective 'qualified' can sometimes be used in a gradable way. For instance, (3.13) is acceptable.

(3.13) She is more qualified than he is.

Importantly, when used as a gradable adjective, 'qualified' seems to have a standard that is vague, and is also context-sensitive.

Thus, we have seen examples where the standard is both vague and context-sensitive, as well as one case, namely, (3.12), where the standard is context-sensitive but not vague. A natural question that arises at this point is, are there also cases where the standard is vague but not context-sensitive? It seems to us that there are no such cases. If this is correct, then it might be that vagueness is deeply rooted in context-sensitivity one way or another, although not all cases of context-sensitivity are vagueness. Different theories of vagueness in fact make different claims about how vagueness relates to context-sensitivity. We will not be able to discuss these different views in detail here, but the interested reader is referred to the Further Reading section at the end of this chapter.

3.1.4 The Sorites Paradox

In the previous subsection, we stated that the vagueness of relative adjectives like 'tall' remains even when its context-sensitivity is removed. This is intuitively so, but what exactly do we mean by this? Or in other words, how do we know when we have vagueness?

The standard diagnostic for showing that a given expression is vague is to construct a version of the so-called *Sorites Paradox* (or the Paradox of the Heap). The word 'sorites' is related to the Greek noun *sorós*, which means 'heap' or 'pile', and the paradox is named so, because the original form of it, attributed to the Greek philosopher Eubulides of Miletus, was about heaps of sand. We will come back to this original form of the paradox below, but let us focus here on examples with gradable adjectives.

A version of the Sorites Paradox for 'tall' goes as follows. Let us say we find a man who is 200 cm tall. Everyone would say he is tall. Now suppose a hypothetical situation in which he was 1 mm shorter. Our intuition tells us that in that case he would still be clearly tall. Generally, a small difference like 1 mm will not make someone who is clearly tall not tall. A similar observation also holds for a short person. If we find a man who is 140 cm tall, then he is clearly short. Even if he were 1 mm taller than he actually is, he would still be short. Again, a 1 mm difference makes no difference as to whether someone counts as tall or not.

Now, the paradox is constructed as follows. The man who is 200 cm tall is clearly tall, and we have just convinced ourselves that shrinkage by 1 mm would not change that, so someone who is 199.9 cm tall is also tall. But note that this reasoning applies recursively. That is, due to the same assumption that a difference by 1 mm would not make a difference, someone who is 199.8 cm should be tall as well. This is true, so there is no problem yet, but a paradox will arise we can keep doing this, which forces us to conclude that someone who is 199.7 cm tall is tall, that someone who is 199.6 cm tall is tall, and so on, because at some point, we will have to conclude that someone who is 142.6 cm tall is tall, that someone who is 142.5 cm tall is tall, etc. These conclusions are clearly absurd, hence the paradox.

The puzzle that underlies this paradox is that while it is true that a small difference like 1 mm makes no difference with respect to who is tall and who is not, a big difference should make a difference, so someone who is 200 cm tall is tall, but someone who is 140 cm tall is not tall. However, in the reasoning above, we only considered two people whose heights differ by 1 mm at a time, and we could not detect where exactly we had the first non-tall person. Yet, it must be that at some point the people under consideration should not count as tall anymore. The puzzle of vagueness comes from this fact that the cut-off point cannot be identified in precise terms, although we know that it exists somewhere and we cross it as we go down the scale of height.

We call expressions vague when they give rise to a paradox like the one we have just seen. We can easily demonstrate that the expression 'tall for a five-year-old British boy' is also vague in the same sense. That is, we start with a five-year-old British boy that is clearly tall. Then we make the assumption that 1 mm difference makes no difference. From this, a counter-intuitive conclusion follows that a five-year-old British boy who is 90 cm tall is tall for a five-year-old British boy.

One can similarly construct such paradoxes for other gradable adjectives. This is left as an exercise for this chapter.

To summarise, the Sorites Paradox shows that vagueness is inherent in the meaning of expressions like 'tall', and they simply cannot be used with respect to a *crisp* standard in any context. That is, if we take two people whose height only differ slightly, say, by 1 mm, they do not differ with respect to this predicate. That is, if one of them is tall, so is the other one, and if one of them is not tall, neither is the other one. On top of this, if one of them is an unclear case of tall, then the other one as well. However, the assumption that 1 mm difference makes no difference for 'tall' gives rise to a paradox that short people should also count as tall. This issue stems from the fact that the boundary between tall and not tall is inherently not sharp.

3.1.5 Vague expressions of other categories

We have so far only discussed the vagueness of gradable adjectives, but vagueness is rampant in natural language and one can easily demonstrate that expressions of various parts of speech show vagueness.

Recall at this point that the original form of the Sorites Paradox uses the noun 'heap', whereby demonstrating that this noun has vague meaning. It goes as follows. 1 million grains of sand clearly forms a heap. If we remove one grain from it, we will have 999,999 grains of sand, and that is still clearly a heap. Removing another grain will still keep it a heap. So the small difference of one grain does not turn a heap into a non-heap. But then, if we continue this reasoning, we will end up concluding that 3 grains of sand forms a heap, which is counter-intuitive. This paradox demonstrates that the noun 'heap' counts as a vague expression.

Similarly, verbs can be vague. For example, consider 'ran'. According to the analysis we proposed in the previous chapter, its extension is the set of all people who ran, but what counts as running? If someone is moving with their feet at a reasonably high speed, they are definitely running, but there is no clear cut-off point from which the movement ceases to be running and starts to count as walking, and here too, we can construct a Sorites paradox. If someone is moving with their feed at 15 km/h, they are definitely running. If they were moving at a speed 0.1 km/hr slower than this, they would be still running. Then we would have to conclude that if they were moving at 1 km/h, they would be running.

In an exercise at the end of this chapter you will see more examples of nouns and verbs that are vague, including relational nouns and transitive verbs.

3.2 The Fuzzy Set Approach to Vagueness

In the previous chapter, we used sets as our main formal tool to talk about extensions, but they are not very well suited for capturing vague meaning. This is because by definition, a set completely demarcates the world into things that are in the set and things that are not in the set. Nothing can be both in the set and not in the set, and nothing can be neither in the set nor not in the set. This is just how Set Theory works.

Note that context-sensitivity is not necessarily a problem for Set Theory. What counts as cheap changes depending on various contextual factors, but one could simply say that different models, which represent different states of affairs, assign different sets as the extension of 'cheap'.

On the other hand, vagueness poses a fundamental issue. If the extension of 'cheap' is a set, then there should be a very clear boundary between things that are cheap and things that are not cheap. Also, in that case, it should not make much sense to say that something is both cheap and not cheap, given that nothing can be both inside and outside the same set, but it seems that statements like 'It's cheap and not cheap' actually sound sensible.

For the rest of this chapter, we will introduce *Fuzzy Set Theory*, which is an extension of Set Theory that allows for sets that have blurred boundaries, and discuss if it gives us a reasonable theory of vagueness.

3.2.1 Fuzzy Set Theory

Recall that in Set Theory a set is determined by what is in the set and a statement of the form '$x \in S$' is simply objectively determined to be either true or false, for any x. There is nothing in between.

In Fuzzy Set Theory, membership is not always black and white, and can be gray. In fact, different shades of gray are allowed. For instance, one can express something like 'x is a member of S' is a little truer than 'y is a member of S', but neither of them is completely true or completely false.

The way this is formally implemented is by representing a set with a *membership function*. A membership function assigns a real number between 0 and 1 (including 0 and 1) to anything in the world, be it an entity or a set. The set of real numbers between 0 and 1 (inclusive) is often denoted by $[0, 1]$.

By way of illustration, let us consider some arbitrary membership function. Let us call it m. Suppose that m assigns 0 to London, and 1 to Tokyo. This means that London is not a member of the set in question, while Tokyo is. Now suppose that m assigns 0.64 to Moscow. We will discuss how this is to be interpreted later, but roughly, it means something like Moscow is not entirely in the set but more in the set than not, because it's closer to 1 than to 0. Suppose also that m assigns 0.23 to Istanbul. Then this means that Istanbul is a bit in the set but less so than Moscow. These statements can be written as in (3.14), using the notation '$m(x)$', which stands for the value that m assigns to x and is read 'm applied to x'.

(3.14) $m(\text{London}) = 0 \quad m(\text{Tokyo}) = 1$
$m(\text{Moscow}) = 0.64 \quad m(\text{Istanbul}) = 0.23$

By assumption m assigns a value to all entities and sets in the universe, and by doing so it *characterises* a fuzzy set. That is, a fuzzy set is a set whose membership is just as according to the membership function that characterises it. Different membership functions characterise different fuzzy sets. For example, the following membership function, k, characterises a different set from m.

(3.15) $k(\text{London}) = 0.15 \quad k(\text{Tokyo}) = 0.3$
$k(\text{Moscow}) = 0.2092 \quad k(\text{Istanbul}) = 0.8$

In Fuzzy Set Theory, we can no longer simply list the members of a set, because membership is fuzzy and comes with different degrees. As such, sets in Fuzzy Set Theory, or *fuzzy sets*, are a little more abstract than sets in standard Set Theory. Yet, it is important to notice that Fuzzy Set Theory is capable of representing the same sets as standard Set Theory. Specifically, when the membership function always returns 0 or 1, that amounts to the same thing as a set in the classical sense. As a concrete example, consider the set of prime numbers. We consider this set to have a clear boundary, or to be a *crisp* set, so is a kind of set that can be described in standard Set Theory. Concretely, it can be defined as (3.16).

(3.16) $\{ x \mid x \text{ is a prime number} \} = \{ 2, 3, 5, 7, 11, \dots \}$

In Fuzzy Set Theory, the same set can be characterised by the following membership function, p.

(3.17) For any prime number x, $p(x) = 1$, and for everything else y, $p(y) = 0$.

Thus, p never assigns an intermediate value, so it is totally clear which things belong to the set and which don't, and we can regard p as representing the same thing as (3.16). It should be obvious that any set from standard Set Theory can be paired with such a crisp membership function that always returns 0 or 1. This means that Fuzzy Set Theory can represent everything Set Theory can, but it can also represent things that Set Theory cannot, namely sets with fuzzy boundaries. In this sense, Fuzzy Set Theory is a proper extension of Set Theory.

3.2.2 Describing vague extensions with fuzzy sets

The membership function of a fuzzy set is just a formal object that assigns a real number between 0 and 1 to each thing in the universe and set that can be formed out of them and other sets, but how we understand such numbers is up to us. To make this point concrete, suppose $m(\text{Toronto}) = 0.88$. We could interpret this as meaning that we are 88% confident that Toronto is a member of the fuzzy set that m characterises, for example. But this is not the only possible interpretation. It could be understood as saying that 88% of Toronto is in the set in question, or even that 88% of the people we consulted in our department at UCL on 15 December 2021 said Toronto is in the set.

This point is worth stressing. Fuzzy sets are nothing more than formal objects, and how we use them in our formal analysis is up to us researchers who apply it to an empirical phenomenon of interest. We can in fact say the same thing about any formal tools we make use of in linguistics and other empirical sciences.

Here, we wish to use fuzzy sets to capture vagueness of natural language semantics, so we would like to understand these numbers to represent something about this linguistic phenomenon. One possibility is to assume that these num-

Figure 3.1: The visual stimulus used by Alxatib & Pelletier 2011: p. 307.

bers represent how confident we are in saying that a given object is in the set. For example, $m(\text{Toronto}) = 0.88$ is interpreted as we are 88% confident that Toronto is in the set that m characterises. This interpretation is called the *epistemic interpretation*, because it has to do with our knowledge (the word 'epistemic' originally comes from '*epistēmē*', which means 'knowledge' in Classical Greek).

By adopting the epistemic interpretation of fuzzy sets, we are essentially committed to the epistemic view of vagueness, according to which vagueness stems from our lack of knowledge. To understand this view, let us consider a concrete example situation depicted in Figure 3.1, which is the visual stimulus used in the questionnaire study reported in Alxatib & Pelletier 2011. The second man from the left is intuitively neither clearly tall nor clearly not tall. According to the epistemic view of vagueness, this intuition reflects our lack of knowledge about the word 'tall'. More specifically, this view assumes that the objective reality is actually categorical and there is no vagueness, so the objective fact is that either this man is tall or he is not tall, but since we do not know which is the case, we hesitate to say he is tall or he is not tall. On the other hand, for the third man from the left, we are confident that he is tall, and for the first man from the left, we can confidently say that he is not tall.

Conversely, according to this view of vagueness, if we knew everything about everything, there would not be vagueness. Recall, however, that vagueness seems to be somehow inherent to words like 'tall', and there does not seem to be a way to use this word in a totally crisp way in any context. To explain this under the epistemic view of vagueness, one would have to assume that such inherently vague words are words for which we simply have no way of knowing for sure how their extensions are delineated, even though the objective reality is completely categorical and these words do have completely clear objective standard.

One might feel uncomfortable with this assumption, because this view is committed to the existence of things that we cannot know, no matter how hard we try. Some such people who are not happy with this view have proposed other ways of looking at vagueness.

For instance, one could assume that vague expressions actually have vague meaning, and therefore it make perfect sense to say things like 'The second man from the left is 0.62 tall'. According to this view, these numbers are not about our knowledge, but something more fundamental about the ontology of vague categories. Then, one could be 100% confident about the fact that the second man in Figure 3.1 is 0.62 tall, for example. An important consequence of adopting this view is that you have to end up admitting the existence of degrees of truth and falsity, because the statement 'The second man from the left is tall', on this view, is neither completely true nor completely false, but something like 0.62 true. This amounts to postulating infinitely many truth-values, so let us call this interpretation of fuzzy sets the *many-valued interpretation*.

The epistemic and many-valued interpretations are examples of different ways of understanding fuzzy set membership and of understanding where vagueness ultimately comes from. We only sampled two views here but there are in fact many more. This introductory textbook obviously cannot settle this debate, or review arguments that have been raised for or against different views. We will therefore leave this issue open, and simply adopt the many-valued interpretation for the rest of the chapter, as it is the interpretation that is most often associated with the fuzzy set approach to vagueness. If you wish to read more about different views of vagueness, you can find some recommended readings in the Further Reading section at the end of this chapter.

3.2.3 Degrees of vagueness

A particularly nice feature of the fuzzy set approach to vagueness is that it is capable of describing different degrees of vagueness. That is, there can be two words, both of which are vague but one of which is less vague than the other. For example, consider two adjectives, 'tall' and 'royal'. As we have seen above, 'tall' is a prototypical case of a word with vague meaning. Note in particular that clear cases are extreme cases, namely, either very tall people, who are definitely tall, or very short people, who are definitely not tall. As a matter of fact, the actual height distribution is known to be very close to the so-called normal distribution, which looks like a bell curve, and the most common height is the average height and extreme cases are very rare. So the majority of people actually are borderline, non-clear cases of 'tall'.

By contrast, the adjective 'royal' has many clear cases. All those people in the British royal family are royal, for example. Similarly people like the authors of this textbook who have nothing to do with royal families of any country or area (at least in recent history), we are clearly not royal. But there are some unclear cases. For example, there are people in France and Russia who are remotely related to former royal families but are commoners now. You might or might not call such people royal. Either way, such unclear cases are certainly much rarer than clear cases of royal and non-royal people.

Thus, with respect to the actual facts, 'tall' is more vague than 'royal' is in the

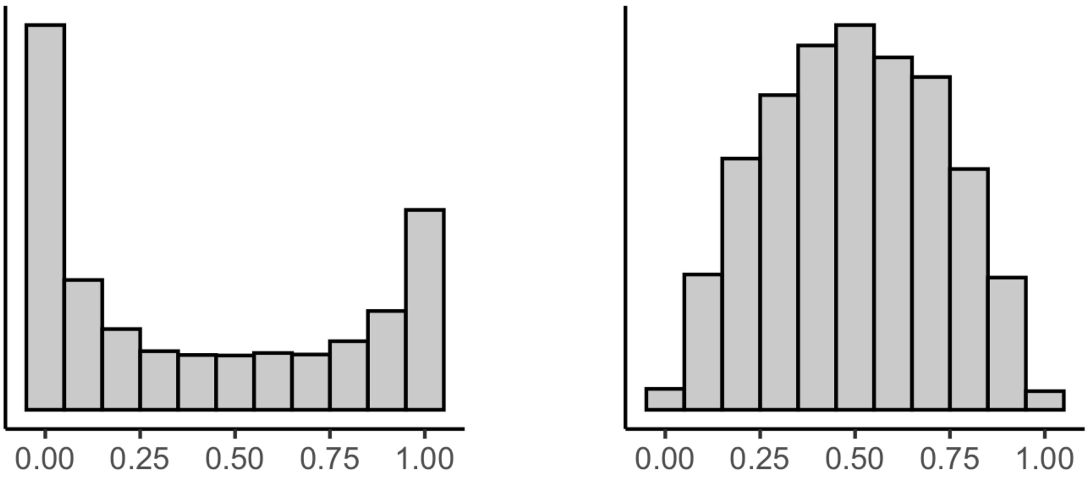

Figure 3.2: Two example membership functions. For each of them, the x-axis represents the output of the function and ranges from 0 to 1, and y-axis represents how many entities the membership function assigns which value. The membership function on the left represents the less vague extension and the function on the right represents the more vague extension relative to a hypothetical model.

sense that more things are in the gray zone for the former than for the latter. The fuzzy set approach to vagueness allows us to nicely capture this difference. For the vaguer expression 'tall', we assign as its extension a membership function that tends to return a value in the middle. For the crisper expression 'royal', on the other hand, we can assign it as its extension a membership function that tends to return extreme values, 0 or 1, but still returns values in the middle for some things in the model. We can visualize this difference as in Figure 3.2.

3.3 Contradictions with Vagueness

Fuzzy sets are a very intuitive way of describing vague meanings, and as such very attractive as the formal foundation for a theory of vagueness. However, there are some potential issues with this approach. In this section, we will focus on the issue of *contradiction*.

Recall that a sentence that is always false, such as (3.18), is called a contradiction.

(3.18) There is something but there is nothing.

In this conjunctive sentence, the second sentence, or *conjunct*, is the negation of the first, so the two conjuncts say opposite things. Conjoining them amounts to saying that both these conjuncts are true, so the resulting meaning is contradictory.

Interestingly, however, when a vague predicate is involved, a sentence of the same form does not necessarily have a contradictory meaning. For instance, consider (3.19).

(3.19) Alice is tall and not tall.

This sentence can actually be used to mean that the referent of 'Alice' is a borderline case of 'tall'. Sentences of this kind are sometimes called *borderline contradictions*, although they are not contradictions in the standard sense.

A theory of vagueness needs to be able to account for why borderline contradictions have sensible meaning, and in this section, we will discuss whether the fuzzy set approach to vagueness can achieve it. Note that sentences like (3.19) are syntactically complex, involving 'not' and 'and', and we haven't talked about their semantics yet. In order to evaluate predictions of the fuzzy set approach to vagueness for sentences like (3.19), we need to augment it with *compositional analyses* of 'not' and 'and' first.

3.3.1 Compositionality

As already mentioned in Ch. **??**, there are reasons to believe that natural language semantics is a *compositional* system. What this means is that in natural languages, meanings can be composed to give rise to other meanings in a systematic way.

It is an important fact about natural language that it contains infinitely many grammatical expressions and each of them has some meaning, as we discussed in detail in Ch. **??**. Compositionality is a key to understanding this fact. If natural language were not compositional, the meanings of all expressions would have to be simply memorized, but that would be simply impossible if there were infinitely many of them to remember.

Another key observation is that while meanings of syntactically simple expressions like 'car' and 'brown' are largely arbitrary in the sense that there's no principled reason why their meanings are conveyed by these particular words, other than socio-historical reasons, the meanings of syntactically complex expressions are systematic and predictable. For example, the meaning of 'brown car' is very systematically related to the meaning of 'red car' and as well as to the meaning of 'brown dog', because the same words are involved. This might seem like a truism, but in theory, there could be a language where 'brown car' means what it means in English while 'red car' means something completely different, e.g. what 'tasty coffee' means in English. This high degree of systematicity suggests that there must be a general, systematic way of composing the meaning of an adjective and the meaning of a noun.

We do not have enough space in this book to develop a full compositional semantic theory that can deal with every sentence and construction in English, which would certainly require more than a single book and a lot more research than has been done so far, but we can still home in on the relevant corner of the English grammar and develop a compositional semantic theory for it. What we will focus on will be sentences of the form 'X is Y', where X is some referring term like a proper name and Y is an adjective in its positive form, as well as more complex sentences formed from them by using 'not' and 'and'. Let us discuss them in turn.

3.3.2 Sentences with positive adjectives

The fuzzy set approach to vagueness assigns a fuzzy set as the extension of an adjective, and a fuzzy set is characterised by some membership function that maps everything in the universe to some real number in $[0, 1]$. As we discussed above, there are different ways of understanding these numbers, but whatever interpretation one adopts, for a sentence like 'This is big', it makes sense to regard the number that the membership function for 'big' assigns to the extension of 'this' as the extension of the sentence.

Let us be a bit more concrete by assuming an example model P. Let us assume that in P, 'tall' is associated with the membership function t, and the extension of the proper name 'Alice', $[\![\text{Alice}]\!]^P$, is some entity a. Now suppose that $t(a) = 0.6$. Since we adopt the many-valued interpretation, we understand it as meaning that it's 60% true that a is 'tall', and it makes sense to regard the extension of the sentence 'Alice is tall' with respect to this model P as this number, 0.6. This analysis is summarised in (3.20).

(3.20) a. $[\![\text{Alice}]\!]^P = a$
b. $[\![\text{tall}]\!]^P = t$
c. $[\![\text{Alice is tall}]\!]^P = t(a) = 0.6$

This analysis is generalizable to all sentences of the same form, as in (3.21).

(3.21) For any model M, if $[\![X]\!]^M$ is an entity and $[\![Y]\!]^M$ is a membership function,

$$[\![X \text{ is } Y]\!]^M = [\![Y]\!]^M([\![X]\!]^M).$$

Recall at this point that the fuzzy set approach can represent crisp sets, sets of the kind that the standard version of Set Theory can represent, in terms of membership functions that always return 0 and 1 and never an intermediate value. Using such membership functions, we can account for non-vague adjectives as well. For example, consider the sentence 'This number is prime' with respect some arbitrary model N. For the sake of argument, let us assume that 'this number' refers to the number 14 in this model, which we simply regard as an abstract entity, and 'prime' has its extension a crisp membership function, p, which assigns 1 to every prime number and 0 to everything else. Then, applying (3.21) to this sentence, we get (3.22c).

(3.22) a. $[\![\text{this number}]\!]^N = 14$
b. $[\![\text{prime}]\!]^N = p$

c. $[\![\text{This number is prime}]\!]^N = p(14) = 0$

Both (3.20) and (3.22) are instances of (3.21), so they demonstrate that this analysis of sentences of the form 'X is Y' can deal with both vague and crisp cases. Furthermore it is a compositional semantic analysis, because it specifies how the meaning of the parts of this syntactically complex expression 'X is Y' is derived in terms of the meanings of its parts, X and Y.[2] We will account for the meanings of 'not' and 'and' below by postulating similar interpretation rules that embody compositionality.

3.3.3 Negation

Let us now discuss how negation 'not' is to be accounted for. To have a compositional analysis for it, we would like to state the meaning of 'X is not Y' in terms of 'X is Y', the latter of which we have just analysed.

It is perhaps easier to start with a crisp case. Let us take the above example, 'This number is prime', where 'this number' refers to the number 14. This sentence is clearly false, and we saw in (3.22) that our analysis accounts for it, as the extension of this sentence is 0. When negation is added to it, as in 'This number is not prime', the sentence becomes clearly true. This suggests that negation turns a false sentence to a true sentence.

This is indeed the standard analysis of negation (we will come back to this in Chapter **??**). but we also need to deal with fuzzy sets. For instance, if the extension of 'Alice is tall' is 0.6 in some model, what should be the value for 'Alice is not tall' in that model? According to the many-valued interpretation we are assuming here, the extension of 'Alice is tall' being 0.6 means that it is 60% true that the referent of 'Alice' counts as tall in this model, so it makes to say that it is 40% true that she does not count as tall. Then, the extension of 'Alice is not tall' in this model should be 0.4.

More generally, we can state the compositional semantic analysis of 'not' in the fuzzy set approach as follows.

(3.23) For any model M, if $[\![X]\!]^M$ is an entity and $[\![Y]\!]^M$ is a membership function,

$$[\![X \text{ is not } Y]\!]^M = 1 - [\![X \text{ is } Y]\!]^M.$$

Using this rule, the above example can be analysed as (3.24). We call the relevant model A.

[2] We simply ignored 'is' in our analysis here, but this is only for the sake of simplicity, and obviously wrong. That is, 'is' carries some meaning that has to do with tense, and that needs to be accounted for in a full analysis of this grammatical construction. However, the semantics of tense is far beyond the scope of this textbook and we cannot discuss it here. Nonetheless, it is fair to say that our simple analysis still captures certain core aspects of the semantics of the construction in question, and factoring in tense will actually not invalidate it.

(3.24) a. $[\![\text{Alice is tall}]\!]^A = 0.6$
b. $[\![\text{Alice is not tall}]\!]^A = 1 - [\![\text{Alice is tall}]\!]^A = 1 - 0.6 = 0.4$

Notice that this analysis accounts for crisp cases in the same way. If a given sentence of the form 'X is Y' is completely true, its extension will be 1, and the rule predicts that the negation of it, 'X is not Y', will have $1 - 1 = 0$ as its extension. If it is completely false, its extension will be 0 and the extension of its negation will be $1 - 0 = 1$.

3.3.4 Conjunction

Turning now to conjunction, 'and', let us again start with a crisp case. Let us suppose that the referent of 'this number' is 13 and the referent of 'that number' is 14 in the current conversational context, which we model with a model C. Then, 'This number is prime' is true while 'That number is prime' is false with respect to this model.

(3.25) a. $[\![\text{this number is prime}]\!]^C = 1$
b. $[\![\text{that number is prime}]\!]^C = 0$

What if we conjoin these two sentences with 'and', as in (3.26)? Clearly, it is false, so we want its extension to be 0.

(3.26) This number is prime and that number is prime.

More generally, when two crisp sentences are conjoined, the whole conjunctive sentence is true if both of the component sentences are true, and false if at least one of them is false.

This much should be easy, but what if vague predicates are involved? For example, take a sentence containing a vague predicate, say 'Alice is tall', and assume its extension happens to be 0.6 in model B, which is to say that 'Alice' refers to someone who is not clearly tall and not clearly short. Let us conjoin this sentence with another sentence that is crisp in B, say 'Alice is French'. Suppose that Alice is a very clear case of a French person and therefore, this second sentence is completely true, which is to say that its extension is 1 with respect to B. These assumptions are summarised in (3.27).

(3.27) a. $[\![\text{Alice is tall}]\!]^B = 0.6$
b. $[\![\text{Alice is French}]\!]^B = 1$

In this scenario, how should we judge the conjunction of these two sentences, given in (3.28)?

(3.28) Alice is tall and Alice is French.

Intuitively, this sentence should not be completely true, although the second conjunct is completely true. Rather, (3.28) sounds as vague as the first conjunct. Incidentally, the order of conjuncts does not seem to matter. So (3.29) receives a similar borderline judgment as (3.28).

(3.29) Alice is French and Alice is tall.

What if one of the conjuncts is borderline as in the above examples but the other one is completely false? Consider the following sentence in the same scenario as above, where the referent of 'Alice' is borderline tall and unquestionably French and hence not Australian.

(3.30) Alice is tall and Alice is Australian.

This sentence sounds completely false, and intuitively this is because of the false conjunct. Again the order of the conjuncts does not seem to matter, so (3.31) is judged as simply false.

(3.31) Alice is Australian and Alice is tall.

Therefore, for sentences like (3.28) and (3.29), crisp falsity dominates the meaning, while crisp truth is in a way ignored. From these observations we can hypothesise that the extension of a conjoined sentence is whichever is the smaller of the extensions of the conjuncts. This analysis can be stated as follows.

(3.32) For any model M, and for any two declarative sentences X and Y,

$$[\![X \text{ and } Y]\!]^M = \min([\![X]\!]^M, [\![Y]\!]^M).$$

The $\min$-function is the function that takes a sequence of numbers and returns the smallest among them. Thus, when one of the conjuncts is false, which is to say that its extension is 0, the extension of the entire conjunction will be 0 as well. If the extension of one of the conjuncts is greater than 0 but less than 1, then the extension of the entire conjunction will never be 1.

Let us see what this analysis predicts for a conjunction where both conjuncts contain vague predicates, as in (3.33).

(3.33) Alice is tall and Alice is rich.

Suppose that Alice is borderline tall such that the extension of 'Alice is tall' is 0.6. In a model where Alice is a billionaire and is undoubtedly rich, then (3.33) sounds borderline true, so it makes sense to assign the same extension to (3.33) as 'Alice is tall'. In a model where Alice is not rich at all and hence the second conjunct is completely false, the entire conjunctive sentence is simply false.

These cases are parallel to the above examples where one of the conjuncts is crisp, and the analysis makes the same predictions as before. Now, consider the following model: The extension of 'Alice is tall' is 0.6 and the extension of 'Alice is rich' is 0.4. According to the general rule in (3.32), the extension of (3.33) in this model should be 0.4. It might not be entirely obvious if this is the right prediction, but at least it does not seem completely off, so let us say that the theory is working well so far.

3.3.5 Borderline contradictions

We now have compositional analyses of 'not' and 'and'. They together make predictions about sentences like the following.

(3.34) This number is prime and it is not prime.

Suppose that the referent of 'this number' is a prime number, say 13, and the pronoun 'it' in the second conjunct refers to the same number. Then, intuitively (3.34) is false in that case, and this is as predicted by our analysis. Specifically, the fist conjunct is true, so its extension is 1. But then since the second conjunct is its negation, so its extension will be $1 - 1 = 0$. The conjunction then takes the smaller of these, so the extension of (3.34) will be 0.

What if 'this number' refers to a non-prime number, e.g. 14? The prediction stays the same, because this time the extension of the first conjunct will be 0, which renders the entire conjunction false. Therefore, (3.34) will be always false no matter what the referent of 'this number' is, which means that it is a contradiction. This is as we want it to be.

Let us now examine the prediction for borderline contradictions such as (3.35).

(3.35) Alice is tall and Alice is not tall.

Suppose that the referent of 'Alice' is of moderate height, so the extension of 'Alice is tall' is somewhere in the middle of 0 and 1, say 0.6. This number will be the extension of the first conjunct here. What about the second conjunct? Since it's the negation of the first conjunct, its extension should be $1 - 0.6 = 0.4$. The extension of the entire sentence in (3.35) will then be the smaller of these two, 0.4.

In a different model where the extension of 'Alice is tall' is 0.3, the extension of the entire conjunction in (3.35) will be 0.3, because in this case, the extension of the second conjunct will be 0.7 and the smaller one of these is 0.3.

A moment's reflection reveals that the extension of (3.35) can never be 1. This is because the extensions of the two conjuncts are systematically related. That is, when the extension of one of them, whichever it is, is 1, then the other one needs to be 0, which will make the extension of the entire conjunction 0. If one of them is 0.2, then the other one needs to be 0.8, and the extension of the entire conjunction will be the smaller of these two, namely 0.2. Therefore, the largest

	True	False	Can't tell
#2 is tall	46.1%	44.7%	9.8%
#2 is not tall	25.0%	67.1%	7.9%
#2 is tall and not all	44.7%	40.8%	14.5%

Figure 3.3: A summary of the results for the four statements about the second man from the left in Fig. 3.1 reported in Alxatib & Pelletier 2011.

value that the extension of (3.35) can ever be is 0.5, and that is the case when the extensions of both conjuncts are 0.5.

This is a good prediction, because as a matter of fact, one could use the sentence in (3.35) to say that Alice is a borderline case of tall. So overall, the fuzzy set approach to vagueness accounts for simple sentences of the form 'X is Y', their negations, and the conjunctions of such sentences.

3.3.6 Experimental data

Having identified the predictions of the fuzzy set approach to vagueness, let us now discuss some empirical data. Above, we relied on our native intuitions, but arguably evaluating gradient theoretical predictions against introspective judgments is far from straightforward.

Recently, a number of linguistic phenomena in the area of the study of meaning have been investigated using experimental methods, and vagueness is one of them. Here we will review experimental results reported in Alxatib & Pelletier 2011 in order to see if the fuzzy set approach to vagueness can account for them.

Alxatib & Pelletier (2011) asked undergraduate students at Simon Fraser University to evaluate several statements about each of the men depicted in Figure 3.1. The statements of particular interest here are those about the second man from the left, whose height is $5'11'$, falling exactly in the middle of these five suspects. Table 3.3 summarises the results of Alxatib & Pelletier's questionnaire for three sentence about this man. Note that the participants could give one of three answers, 'True', 'False' and 'Can't tell'.[3]

Firstly, the positive sentence '#2 is tall' was judged to be true 46.1% of the time and to be false 44.7% of the time, indicating that the participants of the survey considered this man to be a borderline case of 'tall', as intended.

Importantly, this does not necessarily mean that the fuzzy set approach should assign 0.461 as the extension of this sentence, because it is not guaranteed that the truth-value should correspond directly to the proportion of times people should choose 'True' in this questionnaire. Ideally we would like to have an ex-

[3]The participants also gave judgments about one more sentence about the same man, as well as the corresponding four sentences about all the other men in the picture, but these results are omitted here for the sake of simplicity.

plicit bridging assumption that predicts what the results should look like based on the extension predicted by the theory, but we do not have a precise one *a priori*. It will be certainly reasonable to assume that the theoretically predicted extension E of a given sentence should positively correlate with the proportion P of 'True' answers in the questionnaire, it does not need to be a simple equation that $P = E$. It could be $P = 0.9 \times E$ or $P = E^2 + 0.05$, for example. However, we can still triangulate from the results of the other sentences what the bridging assumption should be like.

Consider, for example, the negative version of this sentence, '#2 is not tall'. Here, more people said it was false than it was true. Let us see whether this is compatible with the fuzzy set approach. According to the fuzzy set approach, the extension of this sentence should be 1 minus the extension of '#2 is tall'. If the extension of '#2 is tall' is 0.461, then the extension of '#2 is not tall' will be predicted to be $1 - 0.461 = 0.339$. If we assumed the bridging assumption that these numbers should directly correspond to the rate of true answers in this experiment, the predictions would not match the data, because although the positive sentence was judged to be true 46.1% of the time, the negative sentence was only judged to be true 25.0% of the time, rather than 33.9%.

This suggests that we need a more complicated bridging assumption. For example, we could assume that the proportion of true answers in this experiment is 0.711 times the extension of the sentence. In that case, we can assign the following extensions to make the results compatible with the fuzzy set approach. For expository purposes, we will refer to the model that represents the state of affairs represented in Figure 3.1 by L.

(3.36) a. $[\![\text{\#2 is tall}]\!]^L = 0.648$
b. $[\![\text{\#2 is not tall}]\!]^L = 1 - 0.648 = 0.352$

Bu multiplying these extensions by 0.711, we get 0.461 and 0.250, which match the results. This demonstrates that there are ways to make the results for the first two sentences compatible with the fuzzy set approach.

To explain the full results reported in Alxatib & Pelletier 2011, however, we would need a more complex bridging assumption. This is because the one we considered just now would predict that no sentence should be judged to be true more than 71.1% of the time in this experiment, but Alxatib & Pelletier observed that almost everyone judged '#3 is tall' to be true, where #3 is the tallest of the five and his height is indicated to be $6'6''$, for example.

Instead of putting more effort into finding a better bridging assumption, let us discuss the third sentence at this point, which poses a more fundamental issue for the fuzzy set approach to vagueness. As we remarked above, the extension of a sentence of this form is at most 0.5, and when the two conjuncts have different extensions, it will be identical to whichever is smaller. Since the negative sentence was judged to be true less often in this experiment, we can assume that the negative sentence has a smaller extension. Then, the prediction of the fuzzy set approach is that the conjunctive sentence '#2 is tall and not tall' should have the same extension as the negative sentence '#2 is not tall'. To the extent that the bridging assumption is about the extension and the extension alone, we expect these two sentences to be judged to be true to the same extent. However, this is not what we observe in the results. Rather, the conjunctive sentence is judged to be true more often than the negative sentence.

This is an issue but this issue alone does not constitute enough evidence to dismiss the fuzzy set approach, because there's a theoretical possibility that a more complex bridging assumption might be able to explain the data based on the extensions that the fuzzy set approach assigns to these sentences. However, one thing that is clear is that the bridging assumption must refer to something other than the extensions as well, and we will leave open what such an assumption might say.

One important lesson of this discussion is that linguistic data such as truth-value judgments might not directly correspond to the theoretical predictions, which makes evaluation of a theory against empirical data a little complicated, and conclusions not as firm as one might wish. However, this is how linguistics and other empirical sciences make progress. More often than not, a single set of data is not sufficient to reject a theory. Rather, in order to reject a theory, we often need to evaluate it with data about different aspects of its predictions, and also compare it to different theories. If the theory requires a bridging assumption that is more implausible than another theory, for example, we could say that the data favors the latter, and by accumulating such arguments, we can make the former theory look less and less plausible, ultimately to a point where we can say that we no longer need to consider it.

3.4 Chapter summary

In this chapter we discussed the issue of vagueness, which is a long-standing issue in natural language semantics. Below are the key terms and concepts introduced in this chapter.

- Vagueness is often brought up in the context of the semantics of adjectives. Adjectives are divided into *gradable adjectives* and *non-gradable adjectives*, and gradable adjectives are further divided into *relative adjectives* and *absolute adjectives*.
- Relative adjectives are prime examples of vague expressions, but vague expressions can be found in other parts of speech as well.
- A vague expression gives rise to the *Sorites Paradox*.
- In order to model vague meanings in formal terms, we introduced *Fuzzy Set Theory*, an extension of Set Theory and discussed its potential issues against quantitative data from a recent study.

As already remarked, vagueness is a theoretically very rich area, and the fuzzy set approach to vagueness is just one approach among many. Since we do not have enough space here to give a comprehensive overview of the topic or even discuss all the controversial aspects of the fuzzy set approach. For this reason, we will not make a theoretical commitment here, and try to avoid vagueness as much as possible in the following chapters. Yet, we hope that you found the topic interesting, and would like to encourage you to check out the further reading section to learn more about it.

3.5 Further reading

If you find vagueness interesting, we recommend:

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

This book is a highly accessible introduction to the topic of vagueness. Part I contains a number of very interesting, original examples of vagueness, and fun to read. Part II is a good overview of different theoretical approaches and ideas that have been proposed for vagueness, including the fuzzy set approach. The author furthermore discusses applications of theories of vagueness in computer science and elsewhere in Part III, which is stimulating.

If you would like shorter overviews, there are some written by philosophers:

- J. Robert G. Willaims. 2012. Vagueness. In Gillian Russell & Delia Graff Fara (eds.), *The Routledge companion to philosophy of language*, 143–152. London: Routledge
- Roy Sorensen. 2018. Vagueness. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Summer 2018. Metaphysics Research Lab, Stanford University

There are also overview articles written by linguists:

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

We would also like to mention one article that is specifically about the fuzzy set approach to vagueness.

- Uli Sauerland. 2011. Vagueness in language: The case against fuzzy logic revisited. In Petr Cintula et al. (eds.), *Understanding vagueness: logical, philosophical and linguistic perspectives*, vol. 36, 185–198. Rickmansworth: College Publications

Unlike the articles above this is not meant to be an overview article, and you might find its technical aspects somewhat challenging, but having read the present chapter, you should be able to follow it. It also presents another experiment, and discusses its results with respect to the fuzzy set approach to vagueness.

Exercises

Q1. Using original examples, discuss whether each of the following adjectives should be classified as a gradable or non-gradable adjective, as well as whether it has both gradable and non-gradable uses.

i) far	iv) interesting	vii) recursive
ii) purple	v) empty	viii) amphibious
iii) Canadian	vi) chemical	ix) huge

Q2. Discuss whether each of the following gradable adjectives should be classified as a relative or absolute adjective, and in case it has both relative and absolute uses, raise original examples that illustrate the two uses.

i) full	iv) early	vii) cold
ii) intelligent	v) possible	viii) successful
iii) young	vi) dry	ix) green

Q3. Construct Sorites paradoxes for the following vague expressions.

i) heavy	iv) sing	vii) boy	x) slowly
ii) funny	v) talk	viii) friend	xi) easily
iii) pink	vi) love	ix) chair	xii) happily

Q4. Discuss whether the following nouns and verbs are vague or not. They might have vague uses and non-vague uses at the same time. Note that for relational nouns and transitive verbs, vagueness can be about both or either one of their arguments.

i) student	iv) passport	vii) speak
ii) language	v) dance	viii) meet
iii) coffee	vi) disappear	ix) cancel

Q5. An adjective X is said to be a *contrary* of Y, when for any suitable argument, X and Y cannot both apply to it at the same time, and it maybe that neither of them applies. For instance, 'rich' and 'poor' are contraries, because they cannot both apply to the same person at the same time. If someone is rich, they are not poor and if someone is poor, they are not rich. Furthermore, there can be someone who is neither rich nor poor. Come up with three other pairs of adjectives that are contraries of each other.

Q6. An adjective X is said to be a *contradictory* of Y, when for any suitable argument, either X or Y must apply but they cannot both apply at the same time. For instance, 'dead' and 'alive' are contradictories of each other, because anything that can be dead or alive must be either dead or alive but cannot be both. Come up with one more pair of adjectives that are contradictories of each other.

Q7. Suppose that C is a model where the following is the case.

- $[\![\text{Daniel is employed}]\!]^C = 1$
- $[\![\text{Daniel is British}]\!]^C = 0$
- $[\![\text{Daniel is tall}]\!]^C = 0.73$
- $[\![\text{Daniel is bald}]\!]^C = 0.22$

Using the analysis of negation in (3.23), compute the extensions of the negations of the above sentences.

i) $[\![\text{Daniel is not employed}]\!]^C =$
ii) $[\![\text{Daniel is not British}]\!]^C =$
iii) $[\![\text{Daniel is not tall}]\!]^C =$
iv) $[\![\text{Daniel is not bald}]\!]^C =$

Q8. In the same model as above, C, compute the predicted extensions of the following sentences using the analysis of conjunction given in (3.32).

i) $[\![\text{Daniel is employed and Daniel is tall}]\!]^C =$
ii) $[\![\text{Daniel is British and Daniel is tall}]\!]^C =$
iii) $[\![\text{Daniel is not tall and Daniel is bald}]\!]^C =$
iv) $[\![\text{Daniel is not British and Daniel is tall}]\!]^C =$
v) $[\![\text{Daniel is not bald and Daniel is bald}]\!]^C =$

Q9. To account for sentences of the form '(either) X or Y', let us assume the following rule:

(3.37) For any model M, and for any two declarative sentences X and Y,

$$[\![\text{(either) } X \text{ or } Y]\!]^M = \max([\![X]\!]^M, [\![Y]\!]^M).$$

The max-function takes a sequence of numbers and returns the largest one among them. Using this analysis, compute the extensions of the following sentences with respect to the same model C as above.

i) $[\![\text{Either Daniel is employed or Daniel is tall}]\!]^C =$
ii) $[\![\text{Either Daniel is British or Daniel is tall}]\!]^C =$
iii) $[\![\text{Either Daniel is not tall or Daniel is bald}]\!]^C =$
iv) $[\![\text{Either Daniel is employed or Daniel is not British}]\!]^C =$
v) $[\![\text{Either Daniel is not bald or Daniel is bald}]\!]^C =$

Q10. Here are two *incorrect* analyses of 'and' in the fuzzy set approach.

i) For any model M, and for any two declarative sentences X and Y,

$$[\![X \text{ and } Y]\!]^M = [\![X]\!]^M + [\![Y]\!]^M.$$

ii) For any model M, and for any two declarative sentences X and Y,

$$[\![X \text{ and } Y]\!]^M = [\![X]\!]^M \times [\![Y]\!]^M.$$

For each of these analyses, explain what it is incorrect by identifying examples sentences for which it makes predictions that do not match intuitions.

Bibliography

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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