

Over the course of this class we have seen lots of examples of how things can be ambiguous (having multiple possible meanings) in terms of their morphology (e.g., *unlockable* can mean un-lockable or unlock-able) and in terms of their syntax (e.g., different constituency structures for a phrase like 认识邻居的小孩 or a sentence like *Veena said it rained yesterday*).

But what about semantics? Can things be ambiguous because of their semantics?

Of course words can have multiple meanings, which can create ambiguity. An example I like (I noticed this many years ago, but it became especially relevant since summer 2019) is the word for "police", 差人. The history of this word is that the 差 comes from, e.g., 出差的差 (because historically, the Hong Kong police recruited many foreign police officers). But nowadays I like to interpret the word as if it comes from 差劲的差, and thus 差人 looks like it means "crappy people". (Of course these two 差 have different pronunciations, but when written down they look the same.)

But I'm asking about a different kind of semantics. Rather than ambiguity because words can have different meanings, what I want to ask is, can sentences have different meanings because of different *semantic structures*? Let's consider some examples here.

Neg raising

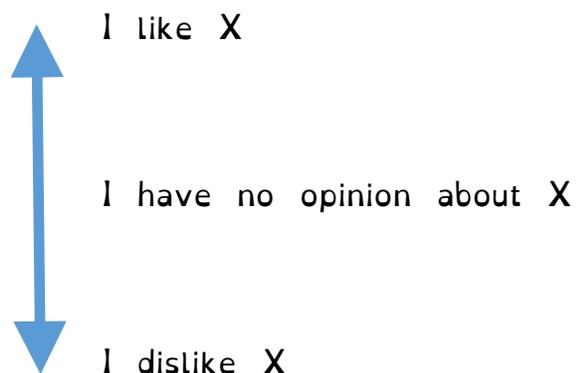
Look at the headline "印度政府"不希望"官員參加達賴喇嘛紀念活動" from <https://www.voachinese.com/a/4280572.html>. How would you translate

this in English? (You may want to try it yourself before you read on.)

I don't know about you, but I would say something like "The Indian government hopes officials won't participate in the activity commemorating the Dalai Lama". I would *not* say "The Indian government doesn't hope officials will participate in the activity...". Why this difference?

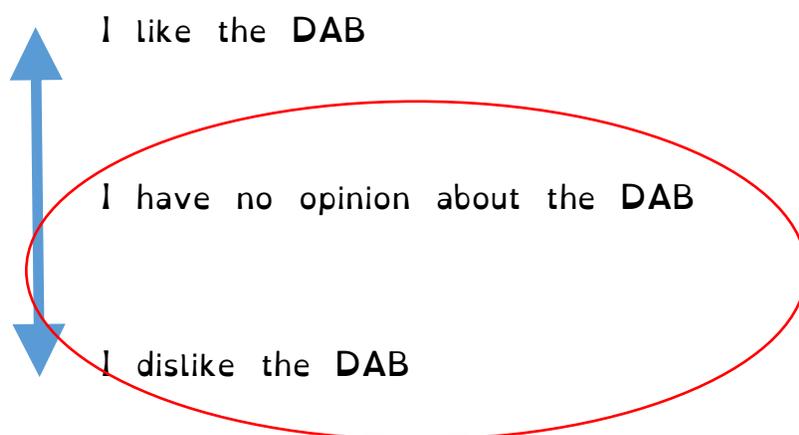
This has to do with something called "neg raising". ("Neg" is short for "negative".) Think of a sentence like "I don't think he'll come to the party" or "I don't like the DAB". Do these things mean what they literally say?

Imagine that "liking" can be on a scale, anywhere from really hating something to really loving something. Like this:



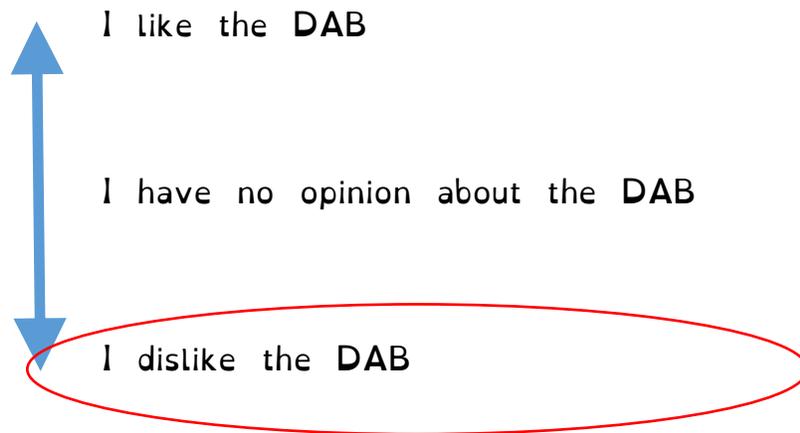
In this example, if something is near the top of the likability scale, we might say "I like it". If something is near the bottom of the likability scale, we might say "I dislike it". And if something is in the middle, we don't care; maybe we don't know enough to say we like it or hate it, or maybe we just really don't care about it. For me, there are some things in the world I really like, some things I really hate, and some things I just don't really have any feeling about.

So, literally speaking, if I say "I don't like the DAB" that would mean it is not in that high likability range. Maybe I really dislike it, or maybe I just don't have an opinion either way; the important thing is, I can't say I really *like* it. In other words, if I say "I don't like the DAB", the meaning should be somewhere in the red circle shown below.



However, "I don't like the DAB" would not normally be interpreted like this. Usually when someone says "I don't like X", people understand that to mean the person specifically *dislikes* X. In other

words, it is understood to have a meaning within the red circle below:



In terms of literal semantics, "I don't like the DAB" could be paraphrased as "It is not the case that [I like the DAB]". But the way people interpret it is instead more like "It is the case that [I dislike the DAB]". And in fact, the latter is probably what the person really meant to say; we would rarely say e.g. "I don't like ice cream" to mean that we feel neutral about ice cream. That's why this is called "neg raising" or "negative raising". The idea is that in the real meaning of the sentence, there is something "negative" in the more inner proposition of the sentence ("I dislike the DAB"), and when we go to actually say it we "raise" that negative to the more outer proposition of the sentence ("It is not the case that..."), which is how a speaker like me can end up saying a sentence like "I don't like the DAB" when I actually mean is "I dislike the DAB".

The same thing applies for a sentence like "I don't think John is coming to the meeting". It usually is really meant to mean "I think John is *not* coming to the meeting". You can think through how this happens because of neg raising.

The interesting thing is that there seem to be some fairly arbitrary and language-specific constraints on what kinds of verbs or predicates can allow neg raising. In English, I can say "I don't want to have ice cream" when I mean "I want to not have ice cream", and "I don't think he's a nice guy" when I mean "I think he's not a nice guy", etc. But it does not seem very acceptable to say "I don't hope you'll do that" when I actually mean "I hope you won't do that"; "I don't hope you'll do that" seems ungrammatical to me. On the other hand, in Chinese, 不希望 seems pretty normal, and to me a sentence like 我不希望你那样做 seems acceptable as a way to express that you hope someone doesn't do something. What do you think?

Negative strengthening

A somewhat similar phenomenon is called "negative strengthening". Imagine I say "Carrie Lam is not a genius". (If you are a fan of Carrie Lam, maybe you can replace her name with Donald Trump to make this example clearer.) Again, like in the "liking" example, imagine there is a scale of smart-ness: someone might be a genius, they might be pretty smart, they might be average, they might be a little dumb, or they might be really dumb. If someone says "Carrie

Lam is not a genius", then literally that means she's not a genius, but she might be anywhere else on the scale --- maybe she's not a genius but she's pretty smart. Do you think this is what a person saying "Carrie Lam is not a genius" means?

In reality they probably meant not only is she not a genius, but she's not even smart; they probably mean to say she's pretty dumb. Likewise if you say about someone "He's not Einstein" --- you probably don't just mean that he's not as smart as Einstein,¹ but you probably actually mean he's even pretty dumb. If you say "he's not handsome", you probably don't just mean that he's not super-attractive, but in fact that he's kind of ugly. This can work with bad descriptions, too; if you say about someone "She's no idiot", you probably don't just mean that she's not an idiot, but in fact that she's pretty smart.

Negative strengthening is an interesting phenomenon and it has some weird interactions with other semantic and pragmatic phenomena; see, e.g., this recent paper by Nicole Gotzner:

<https://www.frontiersin.org/articles/10.3389/fpsyg.2018.01659/full>. Does negative strengthening work in Chinese in a similar way as it does in English, or are there differences? I have no idea, I haven't really thought about it! Can you think of any examples?

¹ To be fair, even Einstein is not Einstein. Sure, the theory of general relativity that he promoted has been hugely influential up to this day, and definitely he has made big contributions to science. But apparently a lot of the work attributed to him was actually done by his wife, and he published much of it without giving her credit.

De Morgan's Laws

In math and logic we often talk about *operators* like AND, OR, and NOT. These are also important operators in semantics. Weird things often happen when they are put together.

In math and logic, there are rules, called De Morgan's Laws, about what happens when you combine NOT with the other operators. Let's see what some of them are.

First, "not (x and y)" is supposed to mean the same thing as "not x or not y". In other words, imagine if I told you to go around the room to find something that is not "red and square". You could find something that's not square (like a red ball) or something that's not red (like a blue book). For something to be "red and square" it needs to be both; so for something to *not* be red and square, all it needs to be is not red or not square.

Next, "not (x or y)" is supposed to mean the same thing as "not x and not y". In other words, imagine that I told you to go around the room and find something that's not "red or square". I wouldn't be satisfied if you found a red ball; I'd say "sure that's not square, but it's still red; I want something that's neither red nor square". I also wouldn't be satisfied if you found a blue square book, for the same reason. I'd only be satisfied if you found something that's not red AND not square.

So there we have it, that's De Morgan's laws:

$\text{not}(X \text{ and } Y) == \text{not}(X) \text{ or } \text{not}(Y)$

$\text{not}(X \text{ or } Y) == \text{not}(X) \text{ and } \text{not}(Y)$

Do these work the same way in language? In other words, is language "logical"? People often act like language should be logical; for example, in colloquial English people sometimes say "I didn't see nobody" to mean they didn't see anybody, and teachers (following prescriptive grammar) tell them that is wrong because, logically, "I didn't see nobody" means you saw somebody. But is it right to expect language to follow the rules of logic? Let's see how De Morgan's laws work in language.



Imagine that I said *I guess John didn't bring beer or wine to the party*. From the above picture (showing three men: one brought wine, one brought beer, and one brought nothing), which guy do you think I'm talking about? Is John the one who brought wine but not beer, the one who brought beer but not wine, or the one who brought nothing?

In English, this would be interpreted as the guy who brought nothing. In other words, in English, *John didn't bring beer or wine* means John didn't bring beer AND John didn't bring wine. Or, to express that more abstractly, not(bring-beer or bring-wine) means not(bring-beer) and not(bring-wine). Does that fit with De Morgan's laws?

(yes, it does seem to.)

What about in Chinese? If you translate this sentence to 我猜張三沒有帶啤酒或者紅酒, do you interpret it the same way I did in English? Or do you get a different interpretation? Can you explain how you interpret this sentence in Chinese, and whether or not it's different from English? (There's no wrong answer here; based on research I have read on this and based on my previous experience discussing this example in previous classes, different Chinese speakers seem to have different interpretations.)

Scope ambiguity

Scope ambiguity is, for me, the most complicated topic out of any we have in this class. Long story short, the issue is that sentences like *Every kid climbed a tree* have multiple interpretations (it could mean that there is one tree, and every kid climbed that tree; or it could mean that every kid climbed his or her own tree, e.g. there were ten kids and ten trees, so kid 1 climbed tree 1, kid 2 climbed tree 2, kid 3 climbed tree 3, etc.).² Furthermore, there appear to be some differences in when scope ambiguity occurs in English versus Chinese. Look at the following English sentences and their possible interpretations:

- *Every plate was broken by someone.*
 - One terrible guest came, and that person broke every plate.
 - Each guest broke their own plate. Like we had a party and at the end of dinner said "Ok, when I count to three, each of you pick up your plate and smash it!"
- *A plate was broken by every guest.*
 - Each guest broke their own plate (see the scenario above).
 - There was one huge plate, and all the guests worked together to break it. (Like we said "ok, everybody grab one side of the plate. When I count to three, we all lift together, so we can pick it up and smash it!")
- *Every guest broke a plate.*
 - Each guest broke their own plate.

² Technically, there are two issues at play here: scope ambiguity, and distributive-collective ambiguity. For the purpose of this subject, we will ignore the details of this distinction. However, if you study advanced semantics, this issue will become important; for a given sentence that has two scope readings, one of the scope readings may have distributive-collective ambiguity and one may not.

- There was one huge plate, and all the guests worked together to break it.
- *Someone broke every plate.*
 - One terrible guest came, and that person broke all the plates.
 - Each plate got broken, but each by a different person.

Here I have manipulated whether the sentence is passive or active, and whether the sentence has "every" in the subject and "a"/"some" in the object, or "a"/"some" in the subject and "every" in the object. But every sentence has multiple possible interpretations. Of course, some interpretations may be more likely than others; for "Every plate was broken by a wild person", it sounds to me more like one person did all the breaking. But if I change some words of the sentence but keep the same structure, I could get a sentence like "Every city was damaged by a bomb"; now it sounds much more likely that different bombs damaged different cities (probably because we know that once a bomb damages one city, the bomb is destroyed, so it can't continue to go on and damage other cities, unless its explosion is so big it covers the entire world).

But what about Chinese? If you translate all these sentences to Chinese, are they all still ambiguous in the same way, or not?

(For a little more info about these kinds of scope ambiguities, see this paper by Chia-Ying Chu:

<http://www.lingref.com/cpp/galana/5/paper3077.pdf>.)