Introduction to Computational Semantics
Lecture notes

Simone Teufel and Weiwei Sun

Michaelmas 2022

1 Introduction

1.1 What is Semantics?

You may have heard the term semantics in every day language in phrases such as "Now, let's stop talking about this, because we are getting into semantics here." Semantics here is roughly equivalent to "hair-splitting", i.e., arguing excessively about very small, insignificant differences. This is a colloquial meaning of the word "semantics" which goes back to the 1930s.

However, semantics also has a technical definition in linguistics. Semantics roughly means "meaning", but are "semantics" and "meaning exactly the same thing? As is often the case with technical terms, there is some overlap in meaning, but the technical term's definition is to a subset of things. Meaning is a vast phenomenon in the world, which is multi-faceted and includes many things that are not connected to language directly. For instance, if two people who are very familiar with each other exchange what is called a "meaningful glance" when they experience a particular situation or hear a certain story, this form of meaning is very personal and situated in the moment. They might both have experienced something very relevant to the story, they might know exactly how the other person feels about this, the shared meaning could be so complex that one might have trouble describing it even in a very long text.

However, this kind of meaning is not tied to a particular element of language, and it's not generally applicable to every speaker of the language. In linguistics, we call "semantics" the subdiscipline that deals with specialised kind of meaning that all language users share about a particular item of language. If we look at the proto-Indo-European (PIE) roots of the two words, we see that "meaning" is related to thinking and "semantics" is related to seeing. Semantics is something we can observe like a natural scientist, because it's defined as the language-specific parts of meaning that all speakers of the language will tend to agree on. Meaning is something internal, very general and/or very personal to a person's perception of the world, and because linguistics understands itself a science, it does not deal with it. (Please note that you should never use the expression "semantic meaning", because it's a tautology[1].

1.2 De Saussure's observation

De Saussure, a language philosopher living from 1957 to 1913, made the groundbreaking observation that what is needed as precondition for a language is a societal agreement to tie a particular symbol for a particular semantics in the world: whenever the symbol is uttered, competent speakers of the language know that the other person is referring to the exact semantics that the other party had in mind when they uttered the symbol. De Saussure called the symbol the "signifier" (that which gives meaning) and the semantics in the world the "signified". He also made the observation that the mapping between the signifier and the signified needs to be *arbitrary*, i.e.,

[1] Tautologies are semantically odd statements which transmit the same semantics more than once in the same statement. Examples: female mother, other alternative

not tied to the semantics it refers to. For instance, calling a cat a “miao” is not arbitrary, because the mapping is based on the sound a cat makes. Calling a cat a cat makes the magic of language happen, because there is nothing about the string “cat” that directly maps to the meaning of cat. It is not the case that the “c” refers to the head of the cat and the “t” to the tail of the cat, or any such thing.

The signifier can be in any medium that language operates in. This can be sound (phonemes) or written words (graphemes) in any syllabic or idiogrammatic language. The medium can also be a sign language.

The arbitrariness of the sign allows a language to express all that needs to be expressed, because for new or complicated concepts, new signifiers can be invented infinitely, and one is not tied to trying to “draw” or otherwise map semantics to a signifier. This gives language the power to express abstract concepts and actions that can’t be “sharaded out”[2].

Are there languages where the signifier and the signified are not arbitrarily connected? Are sign languages just “playing out” the meaning in a sharades kind of manner? They are not. They are abstract languages, it is just the case that the symbol is transmitted by visual signing instead of by sound. But maybe ideogrammatic languages such as Aegyptian hieroglyphs or Chinese characters are? This is based on the mistaken idea that all aspects of such languages are pictograms, which they are not. While the early development of some concrete concepts indeed goes back to ideograms, it’s important to note that these are very few. The vast majority of characters display arbitrary connection of signifier and signified. If they were not, the language were unable to transmit the myriad of very fine distinctions and semantic facets that a language is able to transmit. In fact, with language we can transmit meaning from one brain to the other in infinite detail.

1.3 Lies, Truth and Misleading

Truth obviously matters in the real world. It can make the difference between life and death. Semantics is the study of how what we say maps onto the truth in the world. (It doesn’t always map, and it isn’t supposed to; that is why language is such a powerful tool. It allows us to directly lie, by saying something that is obviously and to our own knowledge not true. It also allows us to *hide* some knowledge by using psychological effects that our listeners will experience if they hear particular statements. Importantly, we can predict what those psychological effects are. A little story to drive this home: Imagine you are in the situation that you are trying to hide some children in your house from a hostile authority figure who is in search of these children, and who has entered your house and is questioning you. They have the power to search your house and take the children away if they don’t like your answers or believe you are lying. They ask you “How many children are in the house?”, and your answer is ”This one here” (say you have a child with you that they aren’t after, and you point to it). If in fact there are 2 more children hidden in the house, have you lied or not? Your statement is strictly speaking true (please work out why), but you have surely mislead the listener. What you *suggested* rather than stated is surely false, because the common usage of giving any number is that there are *exactly* that many children in the house, not *at least* that many. However, if the hostile authority figure believes you then you have achieved the effect you wanted: you have concealed the truth from them without lying. Immanual Kant was the first philosopher to point out the difference between suggested and stated semantics explicitly. You will hear a lot of about Truth Conditions later on in this course.

1.4 Small difference, big result

Let us demonstrate this with a set of sentences that on the surface look quite similar.

(1) a. Kim promised Jo to do the dishes.

[2]Sharades is a game where people try to guess concepts and actions by other people playing them out with one’s body

b. Kim wanted to do the dishes.

c. Jo failed to do the dishes.

d. Kim persuaded Jo to do the dishes.

e. Kim managed to do the dishes.

The sentences look similar on the surface, but when it comes to truth, there are large differences in truth conditions. Imagine a dialogue system that is confronted with these statements and that needs to tailor its response depending on the following questions:

1. Does a washing up action take place or not?
2. Who is doing the washing up?
3. Do we learn anything else about the washing up?

Exercise for the reader: Please answer these questions for each of the sentences above. What we will do in this course about this is the following:

- Lexical semantics is the subdiscipline of defining the difference in meaning between "persuade" and "promise"
- Compositionality is the phenomenon that the resulting semantics of a larger piece of language is a predictable function of the semantics of its pieces
- Syntax-semantics interface looks at phenomena such as "control" - some of these statements have two verbs in them with shared arguments, and the linguistic phenomenon of control explains who the subject of the subordinate clause is
- Pragmatics explains which part of the semantics is due to what we called above "psychological effects" (e.g., presuppositions), which are *predictable* elements of semantics that are transmitted to a listener along with the stated semantics.

Please notice that the differences in meaning look subtle (e.g. to an automatic model of understanding that operates mainly on surface similarity with trained examples), but is immediate and absolute to a human listener, independent of their personal circumstances. It is language alone that transmits this aspect of meaning, which we earlier defined exactly as semantics.

1.5 On the search for a useful meaning representation

A precise characterisation of semantics should be our goal as computational linguists. Language is a phenomenon in the world that has grown into such a complex system that it goes beyond the algorithmic understanding of either machine nor human. All we can do is observe it, like any scientist, describe it, and do experiments with it, leading to deductive conclusions. We can also build simulations. This is what natural sciences do with complex phenomena in the world.

One of the things we will need if we want to talk about semantics is a meta language, in which we can express our observations about semantics. We could use a natural language such ase English. We could try and use a programming language such as Ruby or Scala. We could use maths as our meta-language, with its familiar objects and notations, e.g., a vector or a matrix. We could use a logic formalism such as lambda calculus. Or we could try and express semantics in the form of a formal automata, e.g., a finite state machine.

Descriptions as we find them in human-readable lexica use a natural language. For instance, there are two senses of the word "blue" – the first concens visual effects involving wavelengths and our Netzhaut, whereas the second is metaphorical and refers to a sad or depressed feeling.

What is very common in today's NLP is to use a representation of semantics in the form of a vector with several dimensions. If such a vector is created by a method based on distributional semantics, we often refer to it as a "word embedding". You will hear much more about this representation and how it can be created in L90. Such representations are very well suited to similarity comparisons between words. One of the disadvantages of this represenation that is often mentioned however is its lack of interpretability. For instance, if we look at one particular dimension in the vector (along with its magnitude), in today's vector representations we can no longer assign any meaning to the magnitude. (In simpler, earlier meaning representations of this form, it was possible to directly interpret a dimension as the cooccurrence count between the target word and one of, say 2000, context words (or a function of that cooccurrence count). However, even this information does not give us anything tangible beyond the fact that in a large text, two concepts occurred together a certain number of times.)

Other, language independent vectors are possible, such as the RGB vector that directly expresses the mixture of red, green and blue values corresponding to the intensity of three wavelengths from which all colours can be mixed together. However, what is the meaning of "blue" in that case? The average of all blue hues that a language user of a particular language acknowledges as still being blue?

This would make it impossible for an artist such as van Gogh to have understood the semantics of "blue", because when he was alive, the RGB notation hadn't yet been invented. However, arguably nobody understands colour as well as artists, and in particular van Gogh.

A side effect which is relevant in this example and relatively well researched concerns the cultural dependence of colour perception (effects are particularly strong for green/blue hues). The colour hues that most speakers of West-European languages (Romance and germanic) would call "blue" are distinguished in two colours in Russian, goluboy and siniy. A Russian speaker perceives them as distinct from each other as they would either of these from green. The opposite phenomenon was observed in Japanese, where before contact with the West in 1865, "green" as a linguistic colour did not exist. The hues nowadays seen as "green" by Japanese speakers were clustered either with yellow or with blue; the colour name "midori" for green was only introduced in the Meiji era after 1868. The old meaning of "blue" is visible in old poems where the leafs are called "ao" (blue), and even in the colours of the traffic lights, one of which is "ao" (blue).

1.6 An example of a "precise meaning representation" and its usage

Shapeworld is an example of how precise meaning representations can be used to connect language (and thus system's language ability) to situations in the real world (and thus truth values). It was created by Alexander Kuhnle (supervised by Ann Copestake) and chosen as runner-up for BCS's PhD thesis of the year 2020. (available as technical report UCAM-CL-TR-942.html from the Computer Lab's website). At the center of this representation is a simulated microworld where images involving objects of particular colours sizes and shapes can be randomly created (within fixed parameters set from the outside). In parallel with the images, statements in English about the image are created, and the task is to determine which of these statements are true and which are false.

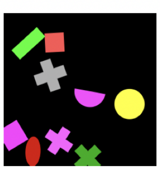

A magenta square is to the right of a green shape.

The lowermost green shape is a cross.

A red shape is the same shape as a green shape.

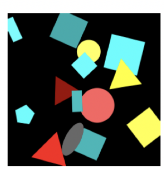

At least half the triangles are red.

More than a third of the shapes are cyan squares.

More than one of the seven cyan shapes is a square.

Here, green statements are true and red statements are false. Shapeworld can generate a huge test and training set of such false and wrong statements. It knows whether a statement is true because the image and the linguistic description are created from the same representation of objects. This representation is called a world model (more about this in the lecture about Truth Values).

{ color: {name: *black*, shade: *0.0*}, noise-stddev: *0.1*, size: *64*, objects:
[{ center: {x: *0.47*, y: *0.28*}, color: {name: *yellow*, shade: *-0.24*},
rotation: *0.06*, shape: {name: *cross*, extent: {x: *0.10*, y: *0.10*}} },
{ center: {x: *0.49*, y: *0.65*}, color: {name: *red*, shade: *0.26*},
rotation: *0.76*, shape: {name: *cross*, extent: {x: *0.08*, y: *0.08*}} },
{ center: {x: *0.15*, y: *0.91*}, color: {name: *yellow*, shade: *-0.16*},
rotation: *0.27*, shape: {name: *pentagon*, extent: {x: *0.09*, y: *0.08*}} },
{ center: {x: *0.80*, y: *0.37*}, color: {name: *red*, shade: *-0.12*},
rotation: *0.53*, shape: {name: *circle*, extent: {x: *0.12*, y: *0.12*}} },
{ center: {x: *0.92*, y: *0.73*}, color: {name: *yellow*, shade: *-0.42*},
rotation: *0.73*, shape: {name: *cross*, extent: {x: *0.09*, y: *0.09*}} }] }

For instance, we see 5 objects here, both as visual representation in the image and as world model descriptions, which give us their properties and positions in a symbolic manner. The world model is an example of a precise meaning representation. What Kuhnle has created around it is methodology dealing with the precise meaning representations.

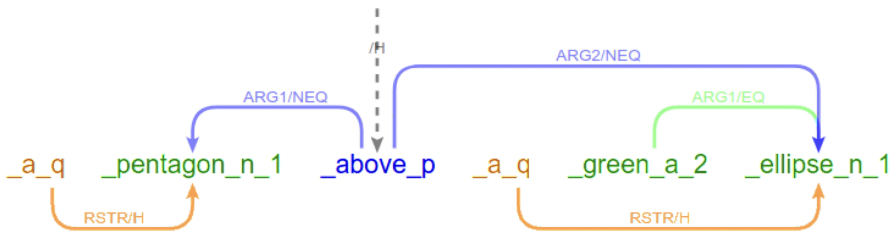


```
[ LTOP: h0 INDEX: e2 RELS: <
    [ _a_q LBL: h4 ARG0:  x3 RSTR: h5 BODY: h6 ]
    [ _pentagon_n_1 LBL: h7 ARG0:  x3 [...]  ]
    [ _above_p LBL: h1 ARG0:  e2 [...]  ARG1:  x3 ARG2:  x8 ]
    [ _a_q LBL: h9 ARG0:  x8 RSTR: h10 BODY: h11 ]
    [ _green_a_2 LBL: h12 ARG0:  e13 [...]  ARG1:  x8 ]
    [ _ellipse_n_1 LBL: h12 ARG0:  x8 [...]  ]
  > HCONS: < h0 ≅ h1, h5 ≅ h7 h10 ≅ h12 > ]
```

"A pentagon is above a green ellipse."

Shapeworld generates natural language statements by translating known relationships, properties and object descriptions into a particular form of *semantic represenation* called DMRS. This representation roughly has the same power as First Order Predicate Logic (FOPL).

We want to create a sentence about something being to the left of something else (or reference R). The algorithm below describes how to create wrong and correct statements of this kind. We first find an instance E' of type E which is indeed to the left of our reference R (E"s X-value is smaller than E's (by more than a minimum distance), and if E' is neither a lot higher or lower than E ("a lot" being defined as the left-right difference between E and E'). If we want to make a false statements, we need to make sure that no duch object e' of class E exists. For each relation it wants to express, this can then be verbalised by an NL generator that creates language based on semantic representations in DMRS. In this case, it results in the statement "a pentagon is above a green ellipse".

Shapeworld is useful as a test suite for visual Question Answering (VQA) because it can capture meaning precisely, even if its world is a micro-world. Traditional datasets for VQA run into problems with vagueness. They typically use real images. For instance, for the image on the left, crowdsourcing had resulted into three questions related to the image.

- What object is shining on the animal?
- What objects is the cat sitting behind?
- How many cats?

Out of the three questions in the image, only the third is objectively answerable as "one" (at least if you mean real cats, not cat-shaped shadows). The others cannot be clearly answered. It is possible that the object shining on the cat is a strong lamp; we cannot see the object on the photo and while it's plausible that it's the sun, this is based on inference and not absolutely true. Concerning the second question, the intended answer might have been "a stack of books"; however, due to perspective and uncertainty about what still counts as "behind", there can be a discussion about whether the cat is also sitting behind the computer mouse type object to the left of the books. In sum, there is doubt whether real-world images can ever be precise. It's also unclear how to create a test set so that it can be kept world-knowledge free, which is a desideratum of the VQA task, as we don't want to measure the system's inference capability, but it's understanding of precisely the scene at hand (and nothing else). The results are therefore not as objective as they should be. (This might be testable in an experiment with many human participants, although this is not a tradition in VQA). In contrast, Shapeworld is precise, guaranteed to be world-knowledge free, and speakers of the language are likely to agree 100% about truth or otherwise of a statement.

1.7 The name of the game in computational semantics (and our position in this course)

What goal does computational semantics set itself? It is to understand a text (or a spoken utterance). In order to be able to accomplish this, we need to have knowledge of many things. The areas of syntax and semantics describe how language locally assembles its semantic units. A key aspect of this is compositional semantics, which is taught in symbolic NLP courses. Lexical semantics concerns the semantics individual words. In this type of semantics, the meaning of individual words are left untreated (also called "atomic"). Much modern NLP research, which has a different approach to lexical semantics (namely the distributional approach), assumes that lexical semantics is a solved problem. In this course, we present a mixture of lexical semantics and compositional semantics, and also a mixture of phenomena and technical solutions to describing, modelling or manipulating these phenomena. The last few lectures are given to pragmantics, the field dealing with meaning in a larger context such as a conversation. One aspect of pragmatics are the semantic components which are left unsaid (often purposefully), but which can nevertheless be calculated (exactly) by a human listener. In pragmatics, there is typically lots of theoretical research but few computational approaches.

2 Word senses

3 Events