623: Semantics and Pragmatics

Spring 2019

James N. Collins

UNIVERSITY OF HAWAI'I AT MĀNOA

First release, January 2018

1. Introduction

1.1 An overview of phenomena

1.1.1 Entailments

If we accept the truth of one statement, what other statements do we necessarily also accept?

(1) a. Sandy smokes.
b. Sandy smokes cigars.

(2) a. A student smokes.
b. A student smokes cigars.
c. A Swedish student smokes.

(3) a. Every student smokes.
b. Every student smokes cigars.
c. Every Swedish student smokes.

(4) a. No student smokes.
b. No student smokes cigars.
c. No Swedish student smokes.

(5) a. Exactly three students smoke.
b. Exactly three students smoke cigars.
c. Exactly three Swedish students smoke.

1.1.2 Quantifiers

How do we reason about quantificational statements and their entailments?

(6) Every class doesn't have prerequisites.

(7) Sam proofread the document, and he didn't see a typo.

(8) Each member of the Dissertation Reading Committee must personally sign one copy of the signature page. (2017 Guide to Preparing Dissertations for e-Submission, Stanford)

(9) Larry King: Do you have a theory about death?
Trixie Mattel: I think it happens to most of us. (*Larry King Now*, 2018)

(10) Trump says collusion isn't a crime. He's right. It's actually many crimes.
(Washington Post, August 1, 2018)

1.1.3 Implicatures

What meanings emerge only in the course of a conversation via reasoning about what the speaker is trying to convey?

(11) This information is included in the opening titles of "Anaconda," and as the words rolled across the screen I heard a chuckle in the theater. It came from me. (Roger Ebert, 1997)

(12) Rogan: I'm going to pretend half of that applause is for me.
Fox: Some of it is for you.
Rogan: Like 30 percent?
Fox: *Some* of it is for you. (Michael J. Fox and Seth Rogan at the Oscars, 2017)

(13) Honest or dishonest? The following contexts differ only in the price of movie tickets:

a. Kyle and Ellen want to see a movie. Kyle has $20 in his pocket. Tickets cost $8 each.
b. Kyle and Ellen want to see a movie. Kyle has $20 in his pocket. Tickets cost $10 each.

Kyle: *I have $8.*

1.1.4 Domain restriction

How do we avoid over- or under-generalization in quantificational statements?

(14)

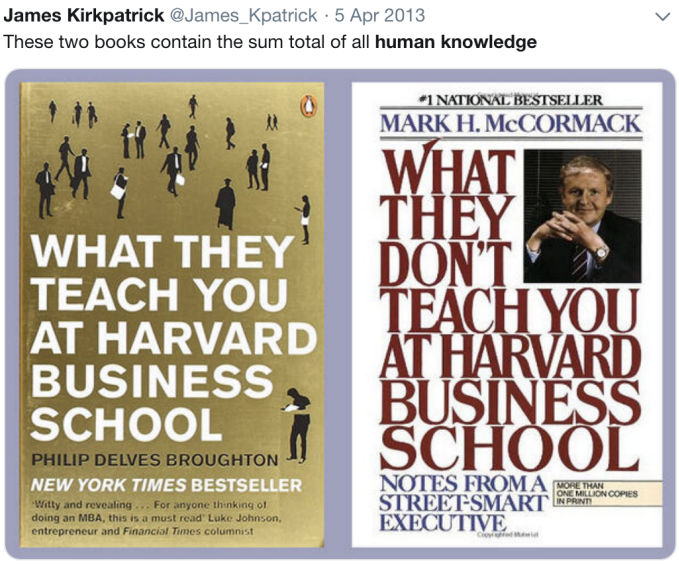

(15) It will not be my intention to do anything to benefit any American. (Ben Carson, 2017)

(16) Supreme Court: *Small v. United States* (April 26, 2005)

a. *Background*: 18 U. S. C. §922(g)(1) forbids "any person [...] convicted in any court [...] of a crime punishable by imprisonment for a term exceeding one year [...] to [...] possess [...] any firearm."

b. *Background*: "Petitioner Small was convicted in a Japanese Court of trying to smuggle firearms and ammunition into that country. He served five years in prison and then returned to the United States, where he bought a gun" and was subsequently arrested.
c. *Small's argument*: any court ranges only over U. S. courts.
d. *Majority opinion*: The meaning of *any* is context dependent. For example, "*any person* may or may not mean to include persons outside the jurisdiction of the state." and "In ordinary life, a speaker who says, *I'll see any film*, may or may not mean to include films shown in another city."
e. *Dissenting opinion* (Thomas, Scalia, Kennedy): The decision "institutes the troubling rule that *any* does not really mean *any*, but may mean '*some subset of any*' ".

1.1.5 Presupposition

Which expressions seem to signal the prior assumptions of the speaker?

(17) Colin Jost: It would be really interesting if [OJ Simpson] killed again...
Michael Che: What do you mean "killed again"? (Late Night with Seth Meyers, Aug 2017)

(18) Mark Halperin: Who did you vote for in the primaries?
Man: Jeb Bush.
Mark Halperin: Oh, you're the one. (MSNBC, Mar 2017)

(19) Still searching for your favorite 90s playlists? (Spotify ad, circa early 2017)

(20) It's [Kate Dickinson]'s birthday today. Let her know you're thinking about her!
(Facebook birthday notification, circa 2017)

(21) Elaine: Is Lippman getting rid of me? It's OK I won't say anything.
Secretary: I don't know anything.
Elaine: Ah, you don't know anything. You see, "*I don't know anything*", means there's something to know. If you really didn't know anything you would have said "*You're crazy.*"
(*Seinfeld*, 1991)

1.1.6 Framing, connotation

What meanings emerge from choices between subtly different expressions?

(22) Lisa: I'm not entirely *uncute*. [...]
Beth: [...] You're pretty. You're very pretty in fact. But *cute*? I don't think so.
Lisa: Well I wasn't aware there was a difference.
Beth: Well of course there's a difference. *Pretty* means pretty. *Cute* means pretty but short and/or hyperactive – like me.
Lisa: Uh huh. What is *beautiful*?
Beth: *Beautiful* means pretty and tall.
Lisa: *Gorgeous*?
Beth: Pretty with great hair.
Lisa: *Striking*?
Beth: Pretty with a big nose. [...]
Lisa: *Exotic*?
Beth: Ugly (*NewsRadio*, 1997)

(23) A New York rabbi said: 'Einstein is unquestionably a great scientist, but his religious views are diametrically opposed to Judaism.' '*But*?' '*But*?' Why not '*and*'?
(Richard Dawkins, *The God Delusion*)

(24) John McCain: It was an energy bill on the floor of the Senate – loaded down with goodies, billions for the oil companies – and it was sponsored by Bush and Cheney. You know who voted for it? You might never know. *That one* [pointing to Obama].
(U.S. Presidential Debate, 2008)

(25) Cecily Strong: I remember our deal. If I guess wrong, I get to 'yum yum' garbage.
Kyle Bennett: ... you "*get to*"? (*Saturday Night Live*, 2014)

1.2 What is meaning?

What does it mean to understand an utterance?

1.2.1 Artificial understanding

Turing 1950: Turing replaced "Can machines think?", which he regarded as "too meaningless to deserve discussion" (p. 442), with the question whether an interrogator could be tricked into thinking that a machine was a human **using only conversation** (no visuals, no demands for physical performance, etc.).

Some objections anticipated by Turing:

> "Thinking is a function of man's immortal soul. God has given an immortal soul to every man and woman, but not to any-other animal or to machines. Hence no animal or machine can think."
>
> "There are a number of results of mathematical logic which can be used to show that there are limitations to the powers of discrete-state machines."
>
> The machine must have a rich, human-like cognitive life and we must be able to verify that.

Searle's Chinese Room argument (Searle 1980)

> Imagine yourself in a room containing a basketful of symbols from a language L that you don't understand, along with a rule book (written in English) for matching symbols in L with other symbols in L. People outside the room pass you strings of symbols in L, you follow your rules, and pass them back symbols in L. The rule book is so good that the symbols you pass back are indistinguishable from the replies of a native speaker of L. You would pass the Turing test, but (Searle says) **no one would say you understand**.

1.2.2 What does it take to interact?

How far away are we from human-like conversational agents?

> Dave Bowman: Open the pod doors, HAL.
>
> HAL: I'm sorry, Dave, I'm afraid I can't do that.
>
> Bowman: What are you talking about, HAL?
>
> HAL: I know that you and Frank were planning to disconnect me, and I'm afraid that's something I cannot allow to happen. (from *2001: A Space Odyssey*, 1968)

> Colbert: Write the show.
>
> Siri: What would you like to search for?
>
> Colbert: I don't want to search for anything! I want to write the show!
>
> Siri: Searching the Web for "search for anything. I want to write the shuffle."
>
> [...]
>
> Colbert: For the love of God, the cameras are on, give me something?
>
> Siri: What kind of place are you looking for? Camera stores or churches?
>
> (from *The Colbert Report*, 2011)

What difficulties are overcome in the first dialogue? What problems come up in the second dialogue?

Logic-driven approaches: transform an expression into a system of symbols a machine can interpret.

> User: Find a block which is taller than the one you are holding and put it into the box.
>
> SHRDLU: OK.
>
> User: How many blocks are not in the box?
>
> SHRDLU: FOUR OF THEM.
>
> User: Is at least one of them narrower than the one which I told you to pick up?
>
> SHRDLU: YES, THE RED CUBE.

What capabilities does SHRDLU need in order to satisfactorily respond to the instructions?

What capabilities does Chat-80 need in order to give the right answers?

> User: Which country bordering the Mediterranean borders a country that is bordered by a country whose population exceeds the population of India?
>
> User: What is the total area of countries south of the Equator and not in Australasia?

1.2.3 Mapping language into true/false statements

Under this logic-driven approach, the machine must be able to verify statements as true or false given some set of facts.

- Block#3 is taller than Block#1.
- Block#7 is not in the box.
- There are three blocks narrower than Block #1.
- Greece borders the Mediterranean.
- Russia has a smaller population than India.
- Canada is south of the Equator.
- Chile is in Australasia.

Natural Language Semantics most commonly take the logic-driven approach to analyzing meaning.

- Semanticists analyze natural language expressions by assigning them a **translation** into some well-defined symbolic system.
- The symbolic system should be free of ambiguities and it should always be totally clear how the semanticist intends the symbols to be interpreted.

A quote from semanticist Lucas Champollion:

> "As a formal semanticist, I try to translate every word to a math expression so that when you put together all the words in a sentence, you get a larger math expression that is true if the sentence is true, and false if the sentence is false. I do this in order to find the precise meanings of certain words and the way these meanings come together in sentences."

Why not just analyze an expression by translating it into another language? Or paraphrase?

Translation: the job of a formal semanticist, illustrated

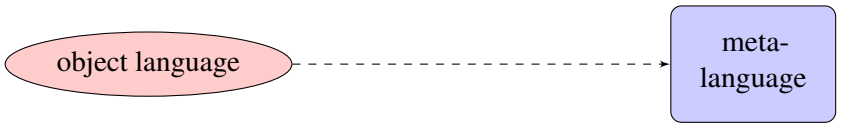

A theory mapping expressions to a well-defined metalanguage is a **theory of linguistic meaning**.

1.3 What is context?

When humans interpret language, we are also always verifying true-false statements. But we don't verify them in relation to a narrow set of facts (like the positions of blocks, or layout of countries). It's a lot more complicated:

> User: Where is **Star Wars** playing in **Honolulu**?
>
> Bot: Star Wars is playing at Kahala Mall.
>
> User: When is **it** playing **there**?
>
> Bot: It's playing at 2pm, 5pm, and 8pm.
>
> User: OK. I'd like 1 **adult** and 2 **children** for **the first show**. How much would **that** cost?

What kind of capabilities does this Bot need?

The Winograd schema (see Levesque 2014):

> The town councillors refused to give the angry demonstrators a permit because they **feared/advocated** violence.
>
> - Who feared/advocated violence?
>
> The trophy would not fit in the brown suitcase because it was so **small/big**.
>
> - What was small/big?

In interpreting language, in order to verify the truth of a statement, we need access to:

- identities of the interlocutors
- the time and place of the conversation
- world information (what time is the movie playing? are councillors or demonstrators violent? do big things fit in small things?)
- information about the discourse (what people/things/places are under discussion?)
- what are the goals of the interlocutor? what is most relevant? informative?
- ...

Context (a very rough definition): the total of prior knowledge assumed by the interlocutors.

1.4 What is a model?

In interpreting an utterance as true/false, it can only be true or false relative to some set of facts, or a context. Therefore, as semanticists, we need a theory of how statements are verified as true or false. Of course, it's difficult to **represent** any sum total of facts known by a human. Therefore, semanticists use models.

Model: a structured representation of facts.

Models are designed by the semanticist to provide a way to check the truth of any statement. They can be laughably simple or inordinately complex. There's only one requirement: *they must provide interpretations for any statement in the metalanguage*.

Interpretation: a semanticist analyses an expression in the object language by *giving it a translation in the metalanguage*, and then checks the interpretation of the translation *relative to a set of background facts* in the model.

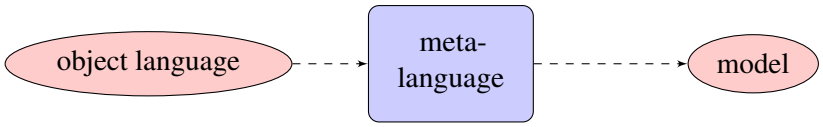

This approach is therefore usually referred to as **model-theoretic semantics**.

1.5 Further reading

- The logic-driven, model-based approach to semantics is very inspired by the work of Richard Montague (and is often called Montagovian semantics). Montague 1973 is the classic, but Gamut 1991a and Gamut 1991b are far more accessible introductions.
- Lewis 1970 is another classic on similar topics, which is less technical than Montague's work, but engages with bigger questions on what linguistic semantics *should* be.
- Probably the most influential book on semantics written from the perspective of generative linguistics is the Heim and Kratzer 1998 textbook, which provides an insightful guide in how to construct an analysis into syntax-semantics interface phenomena.
- Partee, Meulen, and Wall 1990 is a rich resource on the formal underpinnings of natural language semantics (and syntax).
- Readings about formal languages and model theory for this course will mainly come from Benthem et al. 2016, but also see Benthem 1991, Carpenter 1997, Cann, Kempson, and Gregoromichelaki 2009, and Winter 2016 for valuable resources on logic as applied to linguistics.

Bibliography

Benthem, Johan van. 1991. “Language in action”. *Journal of Philosophical Logic* 20 (3): 225–263.

Benthem, Johan van, et al. 2016. “Logic in Action”. Available at www.logicinaction.org.

Cann, Ronnie, Ruth Kempson, and Eleni Gregoromichelaki. 2009. *Semantics: An Introduction to Meaning in Language*. Cambridge: Cambridge University Press.

Carpenter, Bob. 1997. *Type-Logical Semantics*. Cambridge, MA: MIT Press.

Gamut, L. T. F. 1991a. *Logic, Language, and Meaning Vol. 1: Introduction to Logic*. Chicago, IL: University of Chicago Press.

— . 1991b. *Logic, Language, and Meaning Vol. 2: Intensional Logic and Logical Grammar*. Chicago, IL: University of Chicago Press.

Heim, Irene, and Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Levesque, Hector J. 2014. “On our best behaviour”. *Artificial Intelligence* 212:27–35.

Lewis, David. 1970. “General semantics”. *Synthese* 22 (1–2): 18–67.

Montague, Richard. 1973. “The proper treatment of quantification in ordinary English”. In *Formal Philosophy*, edited by Richmond H. Thomason, 247–270. New Haven: Yale University Press.

Partee, Barbara B. H., Alice G. ter Meulen, and Robert Wall. 1990. *Mathematical Methods in Linguistics*. Dordrecht: Springer.

Searle, John R. 1980. “Minds, brains, and programs”. *Behavioral and Brain Sciences* 3 (3): 417–457.

Turing, Alan M. 1950. “Computing machinery and intelligence”. *Mind* 59 (236): 433–460.

Winter, Yoad. 2016. *Elements of Formal Semantics*. Edinburgh: Edinburgh University Press.