24.904
Language Acquisition

Class 6: Words and Their Meanings

Challenges of word learning

1. ~~Finding words in the speech stream~~

2. Associating words with meaning [today+]

Learning word meanings

- What does it mean to *learn* a *word's meaning*?

Word meanings

- What kind of objects are word meanings?
- Let’s try: What is the meaning of *chair?*
  - *an item of furniture*

Word meanings

- Let’s try: What is the meaning of *chair?*

  - Collins Pocket English Dictionary: *A seat with a back and four legs, for one person to sit on*

Word meanings

- Let's try: What is the meaning of *chair?*
  - Collins Pocket English Dictionary: *A seat with a back and four legs, for one person to sit on*
  - The Oxford English Dictionary (2nd edition, 1989): *A seat for one person (always implying more or less comfort and ease); now the common name for the movable four-legged seat with a rest for the back, which constitutes, in many forms of rudeness or elegance, an ordinary article of household furniture, and is also used in gardens or wherever it is usual to sit*

Word meanings

- Let’s try: What is the meaning of *chair?*
  - an architect proposes a dining room in which the chairs are sculpted from the stone that forms the floor. are these *chairs*?
  - in society *x*, all chairs are systematically designed for two people and have 7 legs. are these *chairs*?

discussion adapted from Elbourne 2005

Word meanings

- What kind of objects are word meanings?
  - *extensions* vs. *intensions* (reference vs. sense; particulars vs. concepts, etc.)
    - extension of chair: the set of all chairs
    - intension of chair: the set of all possible chairs
- Whatever the meaning of *chair*, it lets us talk not only about actual chairs we encounter in the world, but also about chairs that are merely possible

Learning

Photo by Pedro Alcocer via Wikimedia. License: CC BY.

Jerry Fodor

- You are shown cards, some of which are *miv* and some *non-miv*.
- You are rewarded when you correctly identify the *miv* cards and punished when you fail

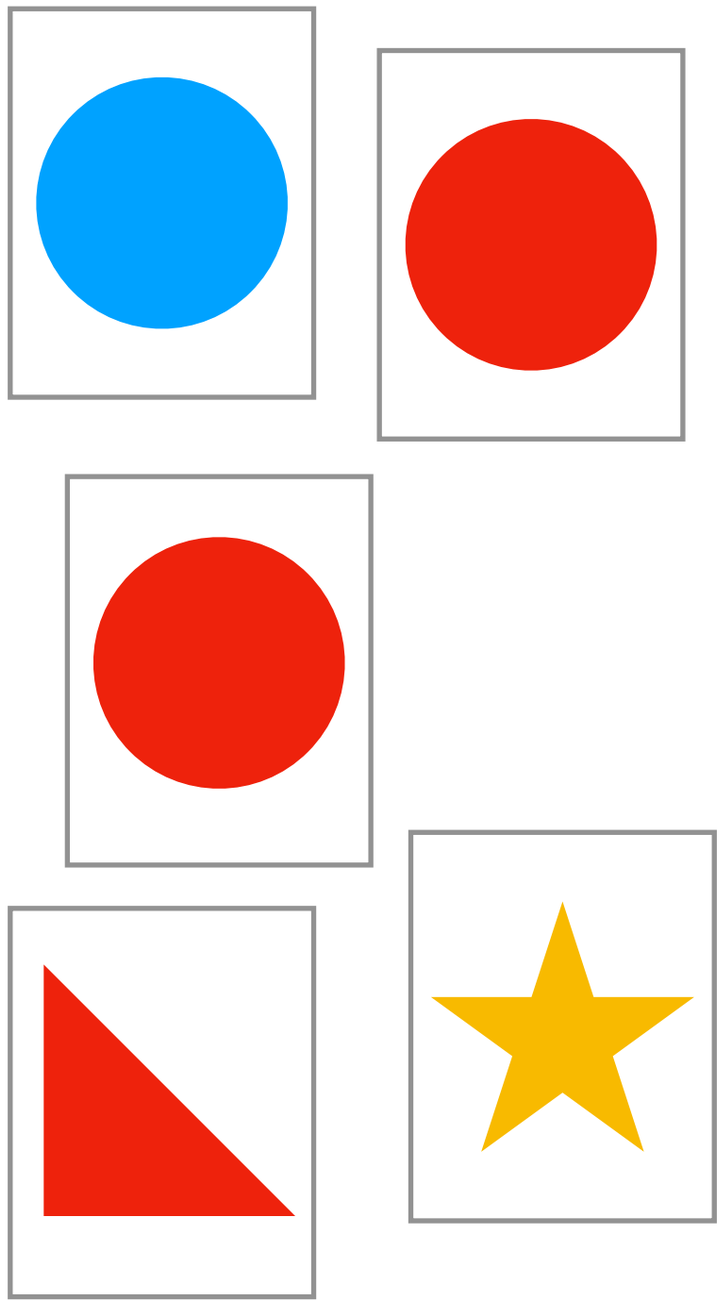

Learning

- **Learning as hypothesis-testing**
  - What you do, according to classical learning theories: consider a hypothesis of the form “X is *miv* iff X is …”
  - The data is then used to assess the truth of the hypotheses with various values of “…”.
    - So if miv means “red and round” then the data will tend to confirm “X is miv iff X is red and round” and disconfirm everything else.

Learning

- Given a hypothesis space, the data (environmental input) can be used to assign a number (a probability) of how well that hypothesis fits the data.

- But where do the hypotheses that are tested come from? i.e. where do the fillers of "…" come from?

- Inductive theories of learning **presuppose the hypothesis space**

The problem of identifying word meanings

A simple story?

“It looks simple. A 14-month-old toddles after the family dog, smacking it whenever she gets close. The dog wearily moves under the table. ‘Dog,’ the child’s mother tells her. ‘You’re chasing the dog. That’s the dog.’ The child stops, points a pudgy hand the dog, and shrieks ‘Daw!’ The mother smiles: ‘Yes, dog.’” (Bloom, 2000)

Defining our problem space

- Not “how children learn words”, but “how children learn words that go with medium sized everyday objects”
- Nothing we will talk about in this class will tell us how the child learns the meanings of “coincidence” or “forget” or “most” or “even”

Learning from observation

- The question:
  - At least some of the time, you probably do learn word meanings by observing the world
  - How does an infant mind do that? What ‘start-up package’ (biases, learning strategies) is it equipped with?
    - infants do not have the store of experience and world knowledge as adults!

Poverty of the Stimulus

- If word meanings are intensions/concepts, we have a POS problem

- How do you go from observing a perceptual slice of the world (a referent associated with the word) to something qualitatively different?

Poverty of the Stimulus

- If, on the other hand, word meanings are equivalence classes of entities grounded in perceptual similarity, the POS problem does not arise

- The task instead is to identify which perceptual features are **relevantly** similar (not all are) and where this knowledge comes from

The gavagai problem

- No one-to-one correspondence between words and perceptual features. The same slice of perceptual experience is compatible with many word meanings (and thus many words)

Quine 1960

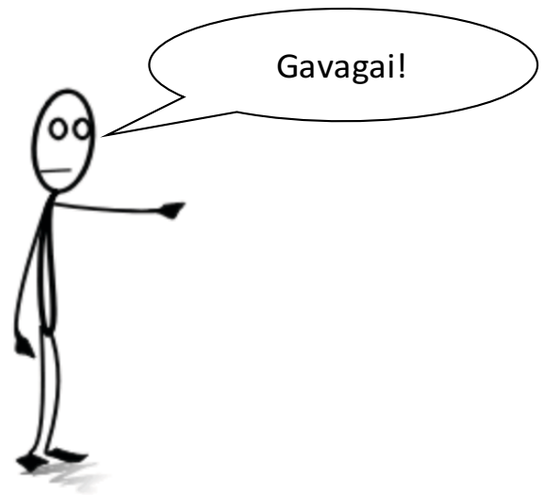

- How do you identify the right level of granularity at which to generalize?

Photo courtesy of USDA NRCS Montana on Flickr. Image is in the public domain.

The gavagai problem

- Quine (1960): the problem of referential uncertainty
  - [[gavagai]] =
    - Rabbit?
    - Rabbit ear?
    - Dinner?
    - Something cute?
    - An animal?
    - A thing that hops?
    - Disconnected rabbit parts?
    - A rabbit but only to the year 2021, then carrots?

The problem of referential uncertainty

- The same problem that the child faces

- Even if something is explicitly labeled in their input (“Look! There’s a doggy!”), how does the child know what specifically that word refers to?

- An infinite number of hypotheses about word meaning is possible. **Input under-specifies the word’s meaning.**

- Yet children figure it out. They have mapped meanings onto some words as early as 6mos of age (Bergelson and Swingley 2012)

Learning biases?

- Innate biases that help constrain hypothesis space
  - mutual exclusivity, newness, ...
  - basic level categories, whole objects as opposed to parts, properties or relations, ...
  - shape (for artifacts), color, texture and shape (for animals), ...

How bad is the referential uncertainty problem really?

- Perhaps Quine's problem is only really pernicious if learning has to happen from a single instance.

- If infants can accrue statistical evidence across multiple referential situations, they might be able to home in on the right target

cross-situational learning

Utterance situation 1: “ball!”

cross-situational learning

Utterance situation 2: “ball!”

cross-situational learning

Utterance situation 3: “ball!”

cross-situational learning

Utterance situation 4: “ball!”

One solution: cross-situational learning

- Only one hypothesis — that "ball" refers to 🔴 — is consistent with all of these situations.

i) Is this true?

ii) If so, can babies reason like this?

		Objects				
		b	***c***	***d***	***s***	***t***
Situation	1	1	1			
	2	1		1	1	
	3	1			1	
	4	1				1

Smith & Yu 2007

- Preferential Looking Paradigm w/ 12-to-14-mo olds
- Training:
  - 6 novel words (*bosa, gasser, manu, colat, kaki, regli*) each associated with a distinct brightly colored shape
  - 30 trials, with 2 objects named with 2 words

manu
colat

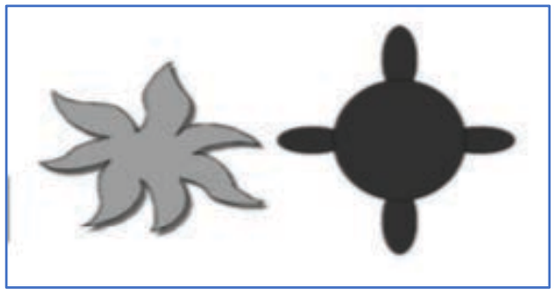

kaki
colat

Smith & Yu 2007

- Test:
  - Two objects; one word. Which object do they prefer to look at?
  - 12 trials (2x per target word)

colat

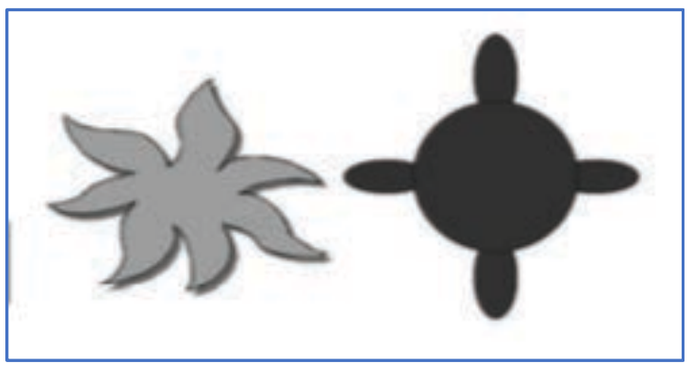

Courtesy Elsevier, Inc., https://www.sciencedirect.com. Used with permission.

Smith & Yu 2007

- Results:
  - Infants preferentially looked at the target
  - Looking behavior varied by word: some word-referent associations better mapped than others

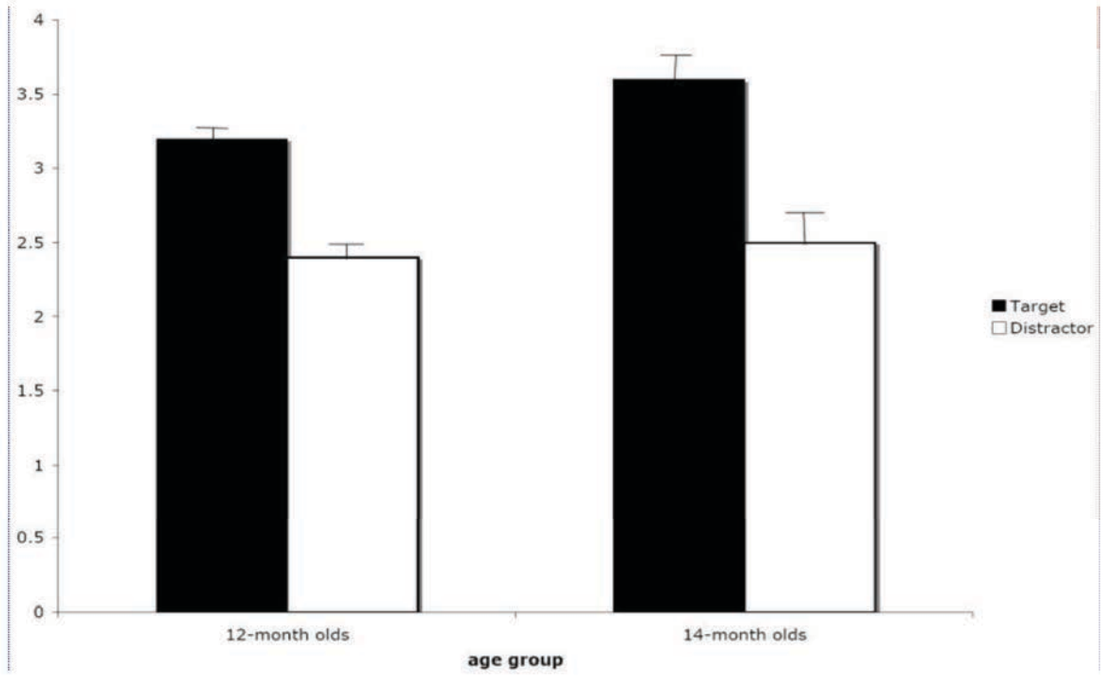

Mean looking time to target and distracter per 8 sec test trial (and standard error of the mean) for younger and older infants.

Smith & Yu 2007

"In sum, these results tell us that cross-situational statistical learning is in the repertoire of young word learners. Despite the ambiguity of word-referent mappings on any individual training trial, infants clearly **accumulate information across trials and use that information to determine the underlying mappings**."

but…

- But how you “determine the underlying mappings” this way?

- Especially if we take seriously our conjecture from earlier that word meanings are not sets of referents

- One possibility: cross-situational learning = accruing evidence for/against a hypothesis across situations (Tenenbaum & Xu 2007)

but… one shot learning

- Carey and Bartlett (1978): “fast-mapping”
  - 3-year-olds were presented with two objects, one blue, one olive-green, and asked: “bring me the chromium one, not the blue one”
  - All of the children retrieved the olive tray, correctly inferring that the
    experimenter intended chromium to refer to this new color.
  - When tested a week later on their comprehension of the word, over half of the children remembered something about its meaning, either that it named olive or that it named a color that resembled olive.

but… what if people just don’t do this?

- What if people can’t actually tabulate word-referent co-occurrence statistics over multiple learning instances in the intended way?

next time

- Propose-but-verify models

- no class Tuesday!