Psych 150/ Ling 155: Psychology of Language

Lecture 7

Learning aspects of words

Announcements

Be working on HW2, due on 5/5/15.

Be working on the review questions for acquisition.

Learning aspects of words

Where are they?
(word segmentation)

What exactly do they sound like?
(word form identification)

What do they correspond to?
(word-meaning mapping)

Word segmentation

Word segmentation:
Divide spoken speech into individual words

tuðəkǽsəlbijándðəgáblɪnsíri

Word segmentation

Word segmentation:

Divide spoken speech into individual words

tuðəkǽsəlbijándðəgáblɪnsíɾi

tu ðə kǽsəl bijánd ðə gáblɪn síɾi

to the castle beyond the goblin city

Word segmentation

Word segmentation:

It's harder than you think when you don't know the language!

http://sites.sinauer.com/languageinmind/wa04.01.html

Word segmentation

"One of the infant's earliest tasks, then, is to figure out which strings of sound form these basic units…babies are confronted with speech in which multiple words are sewn seamlessly together, and they have to figure out all on their own where the edges of words are…spoken language does not present convenient breaks in sound to isolate words."

- Sedivy 2014, p.107

Pauses between words don't really happen

Word boundaries are not necessarily evident in the acoustic waveform

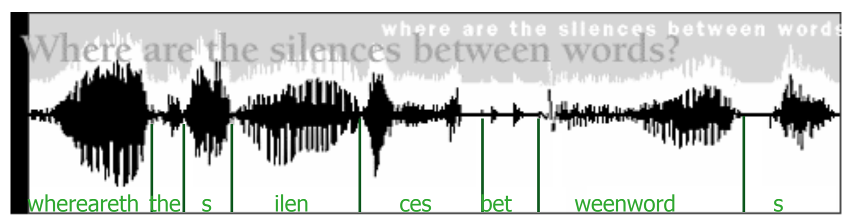

Top-down influence

Chicken Little's repercussions

"Boy, he must think we're pretty stupid to fall for that again."

The sky is falling!

or

This guy is falling!

Top-down influence

[redacted]

The white house is under a tack.

The White House is under attack.

- **Adults** can use top-down information (knowledge of words and the world) to help them with word segmentation.
- What about **infants** who have none or few words in their vocabulary?

Segmentation trajectory

"Babies begin to produce their first words at about a year or so, but they start to identify word breaks at a much younger age than that. In fact, the whole process is under way by 6 or 7 months."

— Sedivy 2014, p.109

How do we know? Use methods like the head-turn preference procedure that capitalize on babies' tendency to devote their attention differently to things that are novel vs. things that are new. (Head-turn preference: Turn their head towards a stimulus they prefer.)

A note about infant preferences: familiarity vs. novelty effects

Sometimes children seem to have a "familiarity preference" where they prefer to look at something similar to what they habituated to. Other times, children seem to have a "novelty preference" where they prefer to look at something different to what they habituated to.

Kidd, Piantadosi, & Aslin (2010, 2012) provide some evidence that this may have to do with the informational content of the test stimulus. There may be a "Goldilocks" effect where children prefer to look at stimuli that are neither to boring nor too surprising, but are instead "just right" for learning, given the child's current knowledge state.

A note about infant preferences: familiarity vs. novelty effects

Practical implications:

Researchers typically look for a difference in infant responses in whichever direction to indicate that infants notice a difference between the test stimuli.

An example of segmentation investigation

Jusczyk & Aslin 1995: testing 7.5-month-old looking preferences

Familiarization stimuli

Her bike had big black wheels. —> həɹbajkhædbɪgblækwilz

The girl rode her big bike. —> ðəgəɹlɹodhəɹbɪgbajk

Her bike could go very fast. —> həɹbajkkʊdgovɛɹifæst

The bell on the bike was really loud. —> ðəbɛlɑnðəbajkwʌzɹililawd

The boy had a new red bike. —> ðəbɔjhædənuɹɛdbajk

Your bike always stays in the garage. —> jɔɹbajkalwezstezɪnðəgəɹɑdʒ

An example of segmentation investigation

Jusczyk & Aslin 1995: testing 7.5-month-old looking preferences

Test stimuli

bike	vs.	dog
bajk		dɑg

7.5-month-old babies look longer at the familiar word.

Implication: They segmented it out of the familiarization stimuli. (However 6-month-olds couldn't yet when they tested them the same way.)

One segmentation cue: Familiar words

Familiar words can help babies break up the speech stream.

How do we segment this artificial language stream?

bankiritubendudifin

Hard to do with no other information!

One segmentation cue: Familiar words

Familiar words can help babies break up the speech stream.

How do we segment this artificial language stream?

ban-kiri-tuben-dudi-fin

What if we recognize the units *kiri* and *dudi*? Much better! (Though still not necessarily perfect — how do we know *tuben* isn't two words, like *tohug* in English?)

One segmentation cue: Familiar words

Familiar words can help babies break up the speech stream.

How do we segment this artificial language stream?

ban-kiri-tuben-dudi-fin

Still, we get two additional words for free, since they're only one syllable: *ban* and *fin*.

Repertoire: *kiri*, *dudi*, *ban*, *fin*

One segmentation cue: Familiar words

Familiar words can help babies break up the speech stream.

How do we segment this artificial language stream?

tamfinatbankirisan

tam-fin-at-ban-kiri-san

This can help us segment new sequences, and add new units to our repertoire.

Repertoire: *kiri, dudi, ban, fin, tam, at, san*

One segmentation cue: Familiar words

Familiar words can help babies break up the speech stream.

Experimental evidence that 6-month-olds can do this in certain situations: Bortfeld (2005) showed that when unfamiliar words are next to words that are really familiar (like *mommy* or the baby's own name), these unfamiliar words can be segmented.

Maggie's bike had big black wheels.

The girl laughed at Mommy's feet.

Another cue: Stress

Infants distinguish between stressed and unstressed syllables, and they learn language-specific biases. English infants prefer words to begin with stress [trochaic: *BUnny*] (Jusczyk et al. 1993, Jusczyk et al. 1999) while French infants prefer words to end with stress [iambic: *année*] (Vihman et al. 1998).

"I went to the castle beyond the goblin city, which was very hard to get to. I saw the goblin king."

Another cue: Stress

Infants distinguish between stressed and unstressed syllables, and they learn language-specific biases. English infants prefer words to begin with stress [trochaic: *BUnny*] (Jusczyk et al. 1993, Jusczyk et al. 1999) while French infants prefer words to end with stress [iambic: *année*] (Vihman et al. 1998).

"I went to the castle beyond the goblin city, which was very hard to get to. I saw the goblin king."

Pretty good strategy for English…

Another cue: Stress

Infants distinguish between stressed and unstressed syllables, and they learn language-specific biases. English infants prefer words to begin with stress [trochaic: *BUnny*] (Jusczyk et al. 1993, Jusczyk et al. 1999) while French infants prefer words to end with stress [iambic: *année*] (Vihman et al. 1998).

"I went to the castle beyond the goblin city, which was very hard to get to. I saw the goblin king."

…though it's not perfect

Figuring out language-specific cues

But how do infants learn these language-specific stress biases? Swingley (2005) suggests that they arise from the initial words infants extract by using statistical cues. This initial set of words is sometimes called a proto-lexicon.

went	castle	goblin	city
very	hard	get	saw
king			

All words in this English proto-lexicon appear to begin with a stressed syllable.

Statistical information available

"…[infants] have other, more flexible and powerful tricks up their sleeves."

—Sedivy 2014, p.115

"…there are measurable statistical regularities that distinguish recurring sound sequences that comprise words from the more accidental sound sequences that occur across word boundaries."

— Saffran, Aslin, & Newport 1996

to the castle beyond the goblin city

Statistical information available

"…[infants] have other, more flexible and powerful tricks up their sleeves."

—Sedivy 2014, p.115

"…there are measurable statistical regularities that distinguish recurring sound sequences that comprise words from the more accidental sound sequences that occur across word boundaries."

— Saffran, Aslin, & Newport 1996

Statistical regularity: *ca + stle* is a common sound sequence

to the castle beyond the goblin city

Statistical information available

"…[infants] have other, more flexible and powerful tricks up their sleeves."

—Sedivy 2014, p.115

"…there are measurable statistical regularities that distinguish recurring sound sequences that comprise words from the more accidental sound sequences that occur across word boundaries."

— Saffran, Aslin, & Newport 1996

No regularity: *stle + be* is an accidental sound sequence

to the castle beyond the goblin city

word boundary

Transitional probability

“Within a language, the transitional probability from one sound to the next will generally be highest when the two sounds follow one another in a word, whereas transitional probabilities spanning a word boundary will be relatively low.” - Saffran, Aslin, & Newport 1996

Forward Transitional Probability = Conditional Probability

ForwTrProb(AB) = Prob(B | A)

Forward transitional probability of sequence AB is the conditional probability of B, given that A has been encountered.

ForwTrProb(“go” ”blin”) = Prob(“blin”$_{curr}$ | “go”$_{curr-1}$)

Read as “the probability of ‘blin’ in the current position, given that ‘go’ has just been encountered in the previous position”

Transitional probability

“Within a language, the transitional probability from one sound to the next will generally be highest when the two sounds follow one another in a word, whereas transitional probabilities spanning a word boundary will be relatively low.” - Saffran, Aslin, & Newport 1996

Forward Transitional Probability = Conditional Probability

ForwTrProb(“go” ”blin”) = Prob(“blin” | “go”)

Example of how to calculate ForwTrProb:
go…
…bble, …bbler, …bbledygook, …blet, …blin
(5 options for what could follow “go”)

TrProb(“go” “blin”) = Prob(“blin”$_{curr}$ | “go”$_{curr-1}$) = 1/5

Transitional probability

“Within a language, the transitional probability from one sound to the next will generally be highest when the two sounds follow one another in a word, whereas transitional probabilities spanning a word boundary will be relatively low.” - Saffran, Aslin, & Newport 1996

Backward Transitional Probability = Conditional Probability

BackTrProb(AB) = Prob(A| B)

Backward transitional probability of sequence AB is the conditional probability of A coming just before, given that B has been encountered.

BackTrProb(“go” ”blin”) = Prob(“go”$_{curr-1}$ | “blin”$_{curr}$)

Read as “the probability of ‘go’ in the previous position, given that ‘blin’ has just been encountered in the current position”

Transitional probability

“Within a language, the transitional probability from one sound to the next will generally be highest when the two sounds follow one another in a word, whereas transitional probabilities spanning a word boundary will be relatively low.” - Saffran, Aslin, & Newport 1996

Backward transitional Probability = Conditional Probability

BackTrProb(“go” ”blin”) = Prob(“go” | “blin”)

Example of how to calculate BackTrProb:
…blin
go…, Du…
(2 options for what could precede“blin”)

BackTrProb(“go” “blin”) = Prob(“go”$_{curr-1}$ | “blin”$_{curr}$) = 1/2

Transitional probability

“Within a language, the transitional probability from one sound to the next will generally be highest when the two sounds follow one another in a word, whereas transitional probabilities spanning a word boundary will be relatively low.” - Saffran, Aslin, & Newport 1996

ForwTrProb
Prob(“stle”$_{curr}$ | ”ca”$_{curr-1}$) = high Why? “ca” is usually followed by “stle”
BackTrProb
Prob(“ca”$_{curr-1}$ | “stle”$_{curr}$) = high Why? “stle” is usually preceded by “ca”

Implication: *castle* likely is a word

to the castle beyond the goblin city

Transitional probability

“Within a language, the transitional probability from one sound to the next will generally be highest when the two sounds follow one another in a word, whereas transitional probabilities spanning a word boundary will be relatively low.” - Saffran, Aslin, & Newport 1996

ForwTrProb
Prob(“be”$_{curr}$ | ”stle”$_{curr-1}$) = low Why? “stle” is not usually followed by “be”
BackTrProb
Prob(“stle”$_{curr-1}$ | “be”$_{curr}$) = low Why? “be” is not usually preceded by “stle”

Implication: *stlebe* likely is not a word (instead, each of these syllables belongs to different words)

to the castle beyond the goblin city

word boundary

8-month-old statistical learning

Saffran, Aslin, & Newport 1996

Testing 8-month-old infants on their ability to use TrProb statistical cues, measuring infant sustained visual fixation preferences.

Stimuli: 2 minutes of repeating four artificial language words made up of three syllables each. Distinguished only by transitional probability between syllables within words vs. syllables across words.

http://whyfiles.org/058language/images/baby_stream.aiff
http://sites.sinauer.com/languageinmind/wa04.03.html

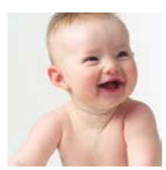

Findings: Eight-month-old infants track and use transitional probability information!

Two-day-old infant sensitivity

Teinonen et al. 2009

Testing 2-day-old infants using EEG, measured during “active” sleep, when the ERPs of newborns are less wildly variable.

Stimuli: 15 minutes of artificial language words made up of ten 3-syllable words.

Findings: Two-day-old infant brains track transitional probability information!

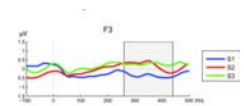

Recap: Word segmentation

Infants are capable of using a variety of cues to help them identify useful units (words) in a fluent speech stream.

The earliest cues they seem to use are statistical cues and (later) familiar words.

Later, they use cues that can be derived from segmented words in their proto-lexicon, like lexical stress patterns.

Word forms

Task:

Map variable acoustic signals to more abstract word forms

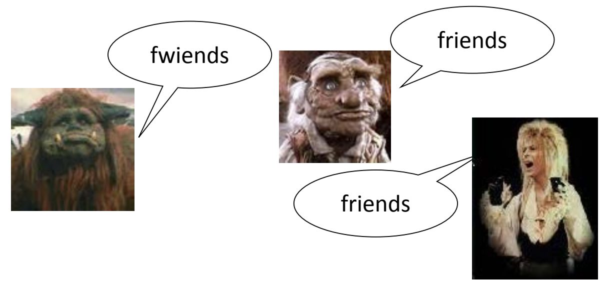

"friends"

Word forms

Word learning: mapping between concept, word, and word's variable acoustic signal

"goblin"

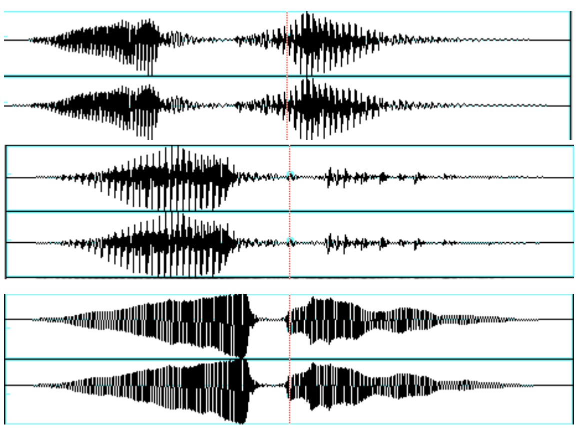

Investigating children's word forms

Basic question:

How much detail do children use to represent word forms?

"Switch" Procedure: measures looking time

…I like the *bih*…look at the *bih*…

Habituation

Test

Same:
look at the *bih!*

Switch:
look at the *dih!*

Stager & Werker 1997

Learning nonsense word forms that are minimal pairs (differ by one phoneme)[‘bih’ vs. ‘dih’], and comparing against words that are not [‘lif’ vs. ‘neem’]

bih…bih…bih…bih…bih…

Habituation

Test

Same:
bih!

Switch:
dih!

Stager & Werker 1997

8-month-olds & 14-month-olds

bih…bih…bih…bih…bih…

Habituation

Test

Same:
bih!

Switch:
dih!

Stager & Werker 1997

8-month-olds & 14-month-olds

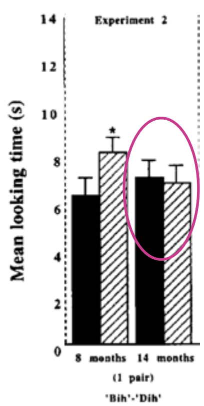

No difference in looking time = 14-month-olds didn’t notice the difference!

Stager & Werker 1997

8-month-olds & 14-month-olds

But 8-month-olds did! They have a difference in looking time. They look longer at the “bih” object when it is labeled “dih” - so they must know “b” and “d” are different.

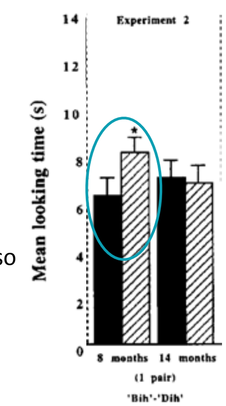

Stager & Werker 1997

14-month-olds

lif…lif…lif…lif…lif…

Habituation

Test

Same:
lif!

Switch:
neem!

Stager & Werker 1997

14-month-olds

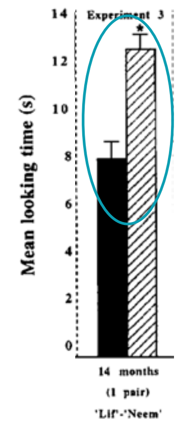

Here, the 14-month-olds look longer at the "lif" object when it's labeled "neem". They notice the difference.

Stager & Werker 1997

14-month-olds

Infants unlikely to associate label with checkerboard pattern (that is, to treat it like a word that has a referent/ meaning)

bih…bih…bih…bih…bih…

Habituation

Test

Same:
bih!

Switch:
dih!

Stager & Werker 1997

14-month-olds

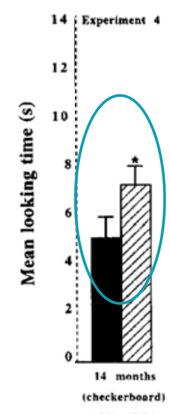

Here, the 14-month-olds look longer at the "bih" "object" when it's labeled "dih". They notice the difference.

Implications

14-month-olds can discriminate the minimally contrasting words (Expt. 4)

…but they fail to notice the minimal change in the sounds when they are paired with objects, i.e., when they are *words with associated meaning* (Expt. 2)

14-month-olds *can* perform the task, when the words are more distinct (Expt. 3)

Therefore, 14-month-olds use more detail to represent sounds than they do to represent words!

One important methodological check

Fennell & Waxman 2010: 14-month-olds can pass this switch task if the communicative purpose of the novel word label is made more salient.

	Stager & Werker 1997	Fennell & Waxman 2010
Habituation	*bih…bih…bih…bih…bih…*	…I like the *bih*…look at the *bih…*
	Issue: Is *bih* a label like "toy"? An exclamation like "wow"? Something else?	Non-issue: *bih* is definitely a label for the object.
Test	*dih!*	Look at the *dih…*
	(This is fine if it means "wow"!)	(This is definitely strange, given the habituation.)

One important methodological check

Fennell & Waxman 2010: 14-month-olds can pass this switch task if the communicative purpose of the novel word label is made more salient.

The communicative intent of the novel word can also be made clear by training items that show familiar objects and labels.

Fennell & Waxman 2010

Kitty!

Car!

Shoe!

Bih!

Non-issue again: *bih* is definitely a label for the object.

Look at the *dih…*

(This is definitely strange, given the habituation.)

One important methodological check

Fennell & Waxman 2010: 14-month-olds can pass this switch task if the communicative purpose of the novel word label is made more salient.

When there's clear intent for the novel word to be a label, 14-month-olds can pass the Switch task just fine.

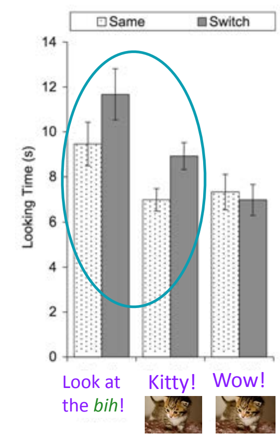

One important methodological check

Fennell & Waxman 2010: 14-month-olds can pass this switch task if the communicative purpose of the novel word label is made more salient.

When it's not clear the novel word is intended as a label (in fact, it seems to be more of an exclamation like "wow"), 14-month-olds look just like they did in the Stager & Werker (1997) experiment.

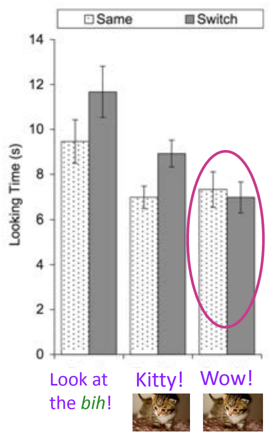

Recap: Implications for word form learning

14-month-olds do store a word's form with enough precision to discriminate between this minimal pair.

— Is this true for all words they know?
— When do they learn to do this?
— What triggers the ability to do this?
(Find out in Acquisition of Language II!)

Word-meaning mapping

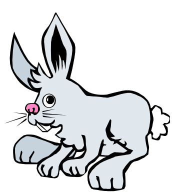

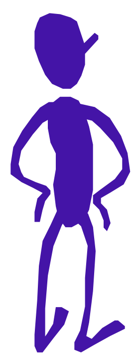

What does "gavagai" mean?

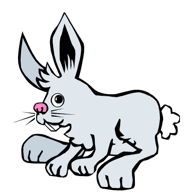

What does “gavagai” mean?

Rabbit?
Mammal?
gray rabbit?
Animal?
Carrot eater?
vegetarian?

Ears?
Long ears?
Is it gray?
Fluffy?
What a cutie!

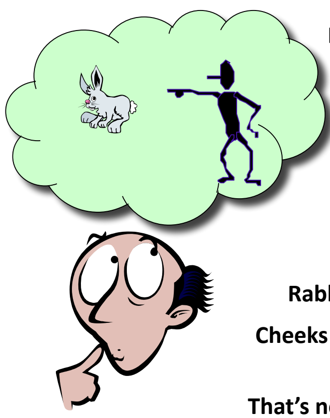

Thumping
Hopping
Scurrying

Stay!
Look!

Meal!
Rabbit only until eaten!
Cheeks and left ear!

That’s not a dog!

Same problem the child faces

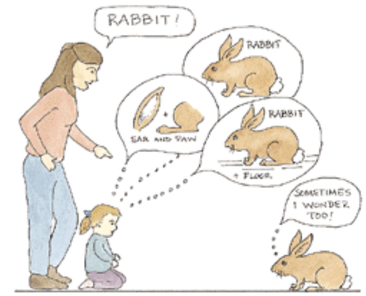

A little more context…

“Look! There’s a goblin!”

The mapping problem

Even if something is explicitly labeled in the input (“Look! There’s a goblin!”), how does the child know what *specifically* that word refers to? (Is it the head? The feet? The staff? The combination of eyes and hands? Attached goblin parts?…)

Quine (1960): An infinite number of hypotheses about word meaning are possible given the input the child has. That is, the input underspecifies the word’s meaning.

So how do children figure it out? Obviously, they do….

Even by 6 to 9 months, infants recognize many familiar words in their language, like body parts and food items (Bergelson & Swingley 2012).

eyes, mouth, hands, …

milk, spoon, juice, cookie, …

Some basic tasks

"I love that *dax*."

Which referent?

dax = lion, tiger, giraffe, tree, lake, crib, grass, flower, leaf, …?

Some basic tasks

"I love that *dax*."

Which referent?

dax = lion, tiger, giraffe, tree, lake, crib, grass, flower, leaf, …?

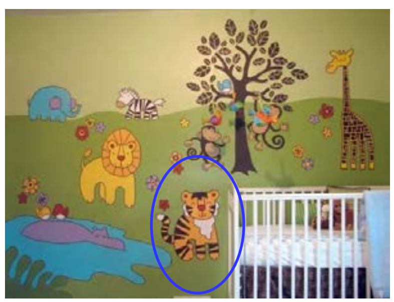

What level of specificity?

Dax = tiger, large predatory cat, animal?

Which referent?

Children begin by making an initial fast mapping between a new word they hear and its likely meaning. They guess, and then modify the guess as more input comes in.

Experimental evidence of fast mapping

(Carey & Bartlett 1978, Dollaghan 1985, Mervis & Bertrand 1994, Medina, Snedecker, Trueswell, & Gleitman 2011)

ball

bear

kitty

[unknown]

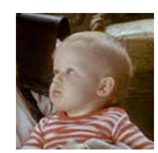

Which referent?

Children begin by making an initial fast mapping between a new word they hear and its likely meaning. They guess, and then modify the guess as more input comes in.

Experimental evidence of fast mapping

(Carey & Bartlett 1978, Dollaghan 1985, Mervis & Bertrand 1994, Medina, Snedecker, Trueswell, & Gleitman 2011)

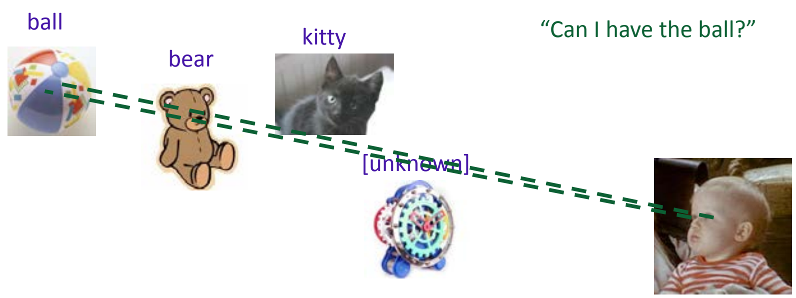

Which referent?

Children begin by making an initial fast mapping between a new word they hear and its likely meaning. They guess, and then modify the guess as more input comes in.

Experimental evidence of fast mapping

(Carey & Bartlett 1978, Dollaghan 1985, Mervis & Bertrand 1994, Medina, Snedecker, Trueswell, & Gleitman 2011)

A slight problem…

"…not all opportunities for word learning are as uncluttered as the experimental settings in which fast-mapping has been demonstrated. In everyday contexts, there are typically many words, many potential referents, limited cues as to which words go with which referents, and rapid attentional shifts among the many entities in the scene."

— Smith & Yu (2008)

One answer: cross-situational learning

Infants accrue statistical evidence across multiple trials that are individually ambiguous but can be disambiguated when the information from the trials is aggregated.

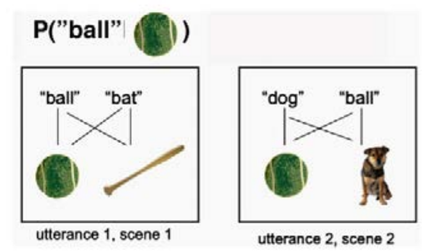

Fig. 1. Associations among words and referents across two individually ambiguous scenes. If a young learner calculates co-occurrences frequencies *across* these two trials, s/he can find the proper mapping of "Ball" to BALL.

Testing cross-situational abilities

Yu & Smith (2007): Adults seem able to do cross-situational learning (in experimental setups).

Smith & Yu (2008): 12- and 14-month-olds can, too.

Something to think about…

The real world isn't necessarily as simple as cross-situational experimental setups (which typically only include two referents for very young children) - often, there will be many potential referents.

(A similar issue to the one fast-mapping has.)

Fig. 1. (*A*) A plausible word learning environment for the word shoe. (*B*) The simulated word-learning environment for shoe found in most cross-situational word-learning experiments.

Something else to think about…

The child's perspective of real world events may make cross-situational learning more feasible, as compared to a neutral third party (the way a photograph represents the world). This is likely because certain things are more salient from a child's perspective due to object foregrounding and degree of clutter in line of sight (Yurovsky, Smith, & Yu 2013).

Fig. 1. (*A*) A plausible word learning environment for the word shoe. (*B*) The simulated word-learning environment for shoe found in most cross-situational word-learning experiments.

What level of specificity?

What does *fep* mean?

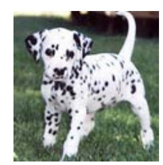

fep

http://sites.sinauer.com/languageinmind/wa05.02.html

What level of specificity?

What does *fep* mean?

fep

dalmatian = subordinate
dog = basic
animal = superordinate

What preference do you have?

Most adults assume the basic level as a default.

What level of specificity?

What about if you saw these examples of *feps*?

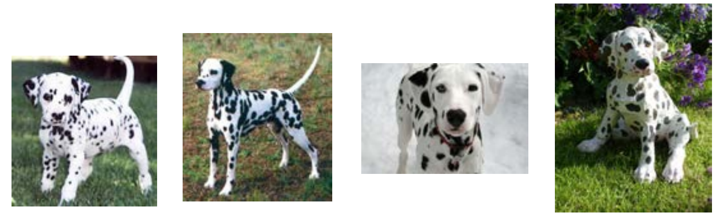

fep *fep* *fep* *fep*

You probably think it's a dalmatian now. Why?
It's a suspicious coincidence: Why is no other animal or other kind of dog a *fep* if *fep* can really label any animal or any kind of dog?

What level of specificity?

Xu & Tenenbaum (2007) found out that 3-year-olds have exactly the same response to suspicious coincidences, and use them to tailor their hypotheses about what a new word means.

fep = …

	dalmatian (subordinate)	*dog* (basic)	*animal* (superordinate)
"This is a fep."			
"This is a fep."			
"This is a fep."			

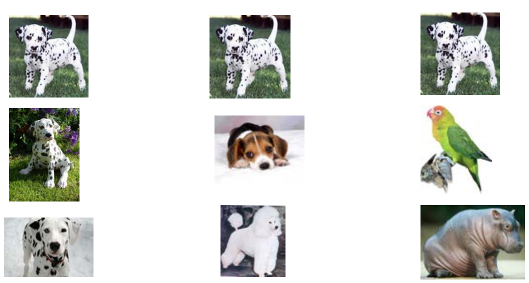

Recap: Word-meaning mapping

The mapping problem can be very hard, especially given all the different ways languages use words to carve up the conceptual space.

However, young children can use different strategies to help guide their hypotheses about what words refer to, including fast mapping, cross-situational learning, and leveraging suspicious coincidences.

You should be able to do up through 38 on the learning review questions and up through 6 on HW2.

Extra Material

8-month-old statistical learning

Saffran, Aslin, & Newport 1996

Testing 8-month-old infants on their ability to use TrProb statistical cues

Measure of infants' response:

Infants control duration of each test trial by their sustained visual fixation on a blinking light.

Idea: If infants have extracted information (based on transitional probabilities) during the habituation trials, then they will have different looking times for the different test stimuli.

Artificial language

Saffran, Aslin, & Newport 1996

4 made-up words with 3 syllables each

Example:

tupiro, golabu, bidaku, padoti

Artificial language

Saffran, Aslin, & Newport 1996

Infants were familiarized with a sequence of these words generated by speech synthesizer for 2 minutes. Speaker's voice was female and the intonation was monotone. There were no acoustic indicators of word boundaries.

Sample monotone speech:

http://whyfiles.org/058language/images/baby_stream.aiff
http://sites.sinauer.com/languageinmind/wa04.03.html

tu pi ro go la bu bi da ku pa do ti go la bu tu pi ro pa do ti…

Artificial language

Saffran, Aslin, & Newport 1996

The only cues to word boundaries were the transitional probabilities between syllables.

Within words, transitional probability of syllables = 1.0

Across word boundaries, transitional probability of syllables = 0.33

tu pi ro go la bu bi da ku pa do ti go la bu tu pi ro pa do ti…

Artificial language

Saffran, Aslin, & Newport 1996

The only cues to word boundaries were the transitional probabilities between syllables.

Within words, transitional probability of syllables = 1.0

Across word boundaries, transitional probability of syllables = 0.33

tu pi ro go la bu bi da ku pa do ti go la bu tu pi ro pa do ti…

Testing infant sensitivity

Saffran, Aslin, & Newport 1996

Test trial:

Each infant presented with repetitions of 1 of 4 words

2 were "real" words

(ex: *tupiro, golabu*)

2 were "part" words whose syllables came from two different words in order

(ex: *pirogo, bubida*)

tu pi ro go la bu bi da ku pa do ti go la bu tu pi ro pa do ti…

Testing infant sensitivity

Saffran, Aslin, & Newport 1996

Results:

Infants listened longer to novel items (part-words)

(6.77 seconds for real words, 7.60 seconds for part-words)

Implication: Infants noticed the difference between real words and part-words from the artificial language after only 2 minutes of listening time! They are sensitive to the transitional probability information.

Smith & Yu (2008): Experiment

Infants were trained on six novel words obeying phonotactic probabilities of English:
bosa, gasser, manu, colat, kaki, regli

These words were associated with six brightly colored shapes
(sadly greyscale in the paper)

Figure from paper

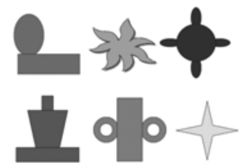

What the shapes are probably more like

Smith & Yu (2008): Experiment

Training: 30 slides with 2 objects named with two words (total time: 4 min)

Example training slides

manu
colat

bosa
manu

Smith & Yu (2008): Experiment

Testing: 12 trials with one word repeated 4 times and 2 objects (correct one and distracter) present

Which one does the infant think is *manu*? That should be the one the infant prefers to look at.

manu
manu
manu
manu

Smith & Yu (2008): Experiment

Results: Infants preferentially look at target over distracter, and 14-month-olds looked longer than 12-month-olds. This means they were able to tabulate distributional information across situations.

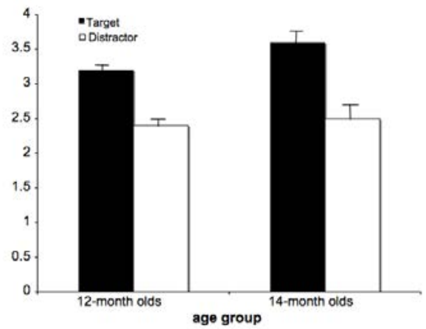

Implication: 12 and 14-month-old infants can do cross-situational learning

Xu & Tenenbaum 2007

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word ("blick", "fep", or "dax") during training. ("This is a blick/fep/dax") There were three classes of stimuli: vegetables, vehicles, and animals.

The vegetable class had these levels:

subordinate: green pepper
basic: pepper
superordinate: vegetable

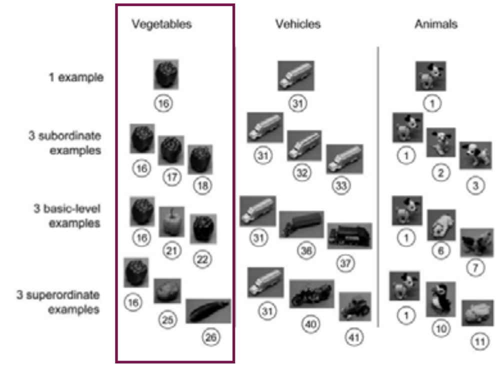

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word ("blick", "fep", or "dax") during training. ("This is a blick/fep/dax") There were three classes of stimuli: vegetables, vehicles, and animals.

The vehicle class had these levels:

subordinate: yellow truck
basic: truck
superordinate: vehicle

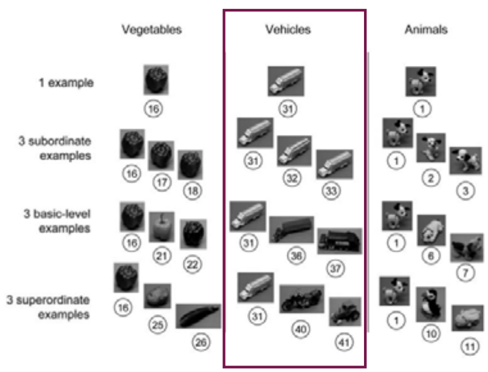

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word ("blick", "fep", or "dax") during training. ("This is a blick/fep/dax") There were three classes of stimuli: vegetables, vehicles, and animals.

The animal class had these levels:

subordinate: terrier
basic: dog
superordinate: animal

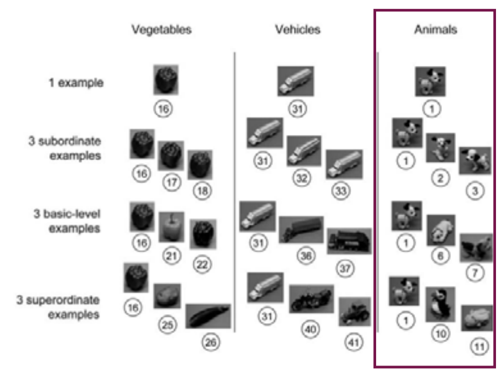

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word ("blick", "fep", or "dax") during training. ("This is a blick/fep/dax") There were three classes of stimuli: vegetables, vehicles, and animals.

There were four conditions:

The 1-example condition presented the same object & label three times.

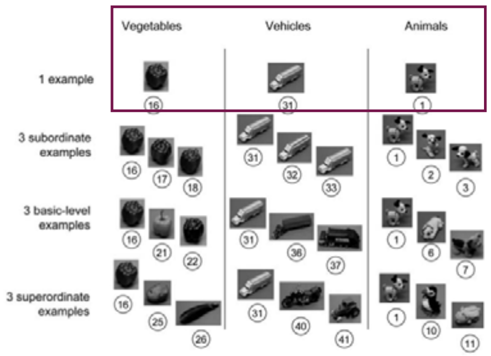

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word ("blick", "fep", or "dax") during training. ("This is a blick/fep/dax") There were three classes of stimuli: vegetables, vehicles, and animals.

There were four conditions:

The 1-example condition presented the same object & label three times.

"This is a fep."

"This is a fep."

"This is a fep."

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word ("blick", "fep", or "dax") during training. ("This is a blick/fep/dax") There were three classes of stimuli: vegetables, vehicles, and animals.

There were four conditions:

The 3-subordinate example condition presented a subordinate object & label three times.

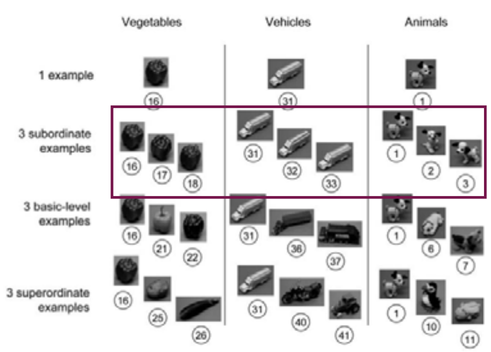

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word ("blick", "fep", or "dax") during training. ("This is a blick/fep/dax") There were three classes of stimuli: vegetables, vehicles, and animals.

There were four conditions:

The 3-basic-level example condition presented a basic-level object & label three times.

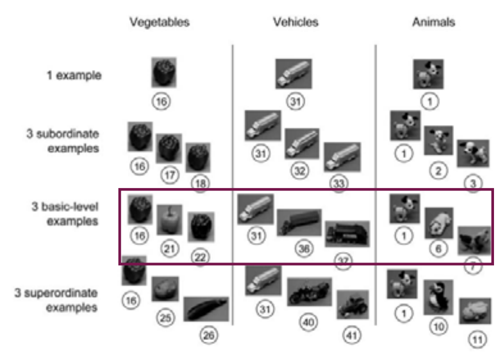

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word (“blick”, “fep”, or “dax”) during training. (“This is a blick/fep/dax”) There were three classes of stimuli: vegetables, vehicles, and animals.

There were four conditions:

The 3-basic-level example condition presented a basic-level object & label three times.

“This is a fep.”

“This is a fep.”

“This is a fep.”

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word (“blick”, “fep”, or “dax”) during training. (“This is a blick/fep/dax”) There were three classes of stimuli: vegetables, vehicles, and animals.

There were four conditions:

The 3-superordinate example condition presented a superordinate object & label three times.

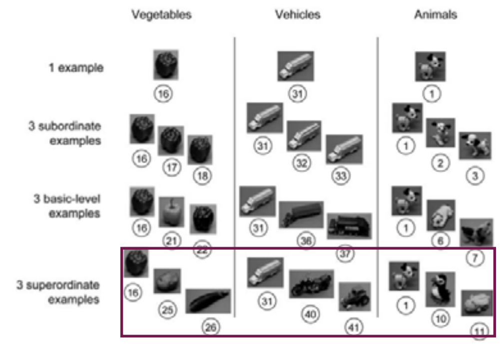

Testing children

Subjects: 3- and 4-year-old children

Task, part 1: Children were presented with three examples of a novel word (“blick”, “fep”, or “dax”) during training. (“This is a blick/fep/dax”) There were three classes of stimuli: vegetables, vehicles, and animals.

There were four conditions:

The 3-superordinate example condition presented a superordinate object & label three times.

“This is a fep.”

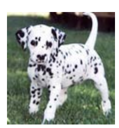

“This is a fep.”

“This is a fep.”

Testing children

Subjects: 3- and 4-year-old children

Task, part 2: generalization (asked to help Mr.Frog identify only things that are “blicks”/ “feps”/ “daxes” from a set of new objects)

There were three kinds of matches available:

Subordinate matches (which were the least general, given the examples the children were trained on)

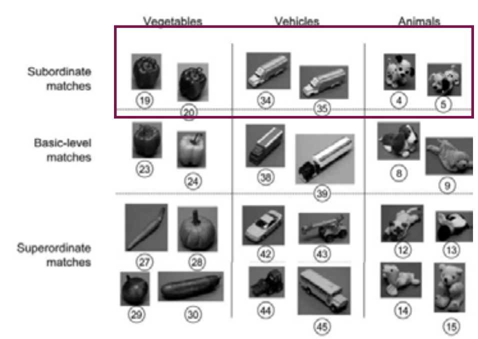

Testing children

Subjects: 3- and 4-year-old children

Task, part 2: generalization (asked to help Mr.Frog identify only things that are "blicks"/ "feps"/ "daxes" from a set of new objects)

There were three kinds of matches available:

Basic-level matches (which were more general, given the examples the children were trained on)

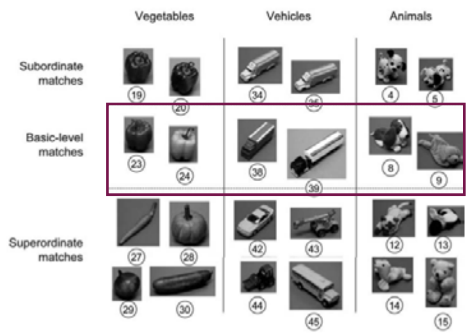

Testing children

Subjects: 3- and 4-year-old children

Task, part 2: generalization (asked to help Mr.Frog identify only things that are "blicks"/ "feps"/ "daxes" from a set of new objects)

There were three kinds of matches available:

Superordinate-level matches (which were the most general, given the examples the children were trained on)

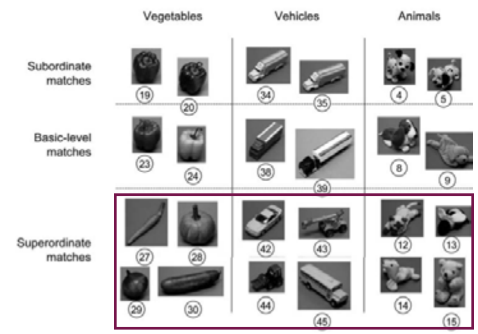

Children's generalizations

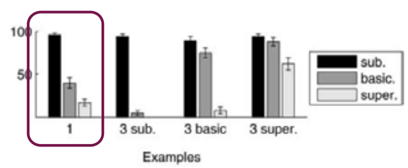

When children heard a single example three times, they readily generalized to the subordinate class, but were less likely to generalize to the basic-level, and even less likely to generalize to the superordinate level. This shows that young children are fairly conservative in their generalization behavior.

"This is a fep."

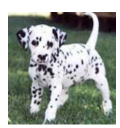

"This is a fep."

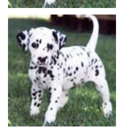

"This is a fep."

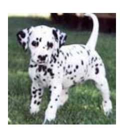

Children's generalizations

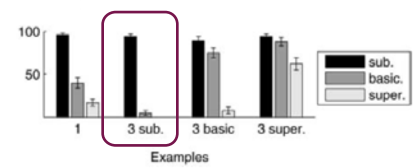

When children had only subordinate examples as input, they readily generalized to the subordinate class, but almost never generalized beyond that. They were sensitive to the suspicious coincidence, and chose the least-general hypothesis compatible with the data.

"This is a fep."

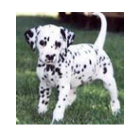

"This is a fep."

"This is a fep."

Children's generalizations

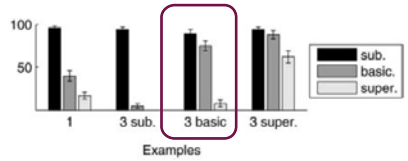

When children had basic-level examples as input, they readily generalized to the subordinate class and the basic-level class, but almost never generalized beyond that. They were again sensitive to the suspicious coincidence, and chose the least-general hypothesis compatible with the data.

"This is a fep."

"This is a fep."

"This is a fep."

Children's generalizations

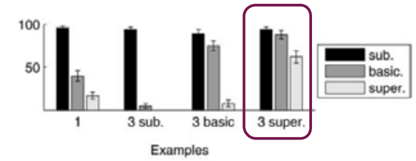

When children had superordinate-level examples as input, they readily generalized to the subordinate class and the basic-level class, and often generalized to the superordinate class. They were again sensitive to the suspicious coincidence, though they were still a little uncertain how far to extend the generalization.

"This is a fep."

"This is a fep."

"This is a fep."