

Psych 156A/ Ling 150:
Acquisition of Language II

Lecture 6
Words in Fluent Speech I

Announcements

HW1 due today by the end of class

HW2 now available (not due till after midterm)

Review questions on word segmentation now available

Midterm review: in class on 5/3/12

Midterm: during class on 5/8/12


Computational Problem

Divide spoken speech into individual words

tuðəkæ̀səlbiꞑándðəgáblɪnsíri

Computational Problem

Divide spoken speech into individual words

tuðəkæ̀səlbiꞑándðəgáblɪnsíri 

↓

tu ðə kæ̀səl biꞑánd ðə gáblɪn síri
to the castle beyond the goblin city

Word Segmentation

“One task faced by all language learners is the segmentation of fluent speech into words. This process is particularly difficult because word boundaries in fluent speech are marked inconsistently by discrete acoustic events such as pauses...it is not clear what information is used by infants to discover word boundaries...there is no invariant cue to word boundaries present in all languages.”

- Saffran, Aslin, & Newport (1996)

Pauses between words don't really happen

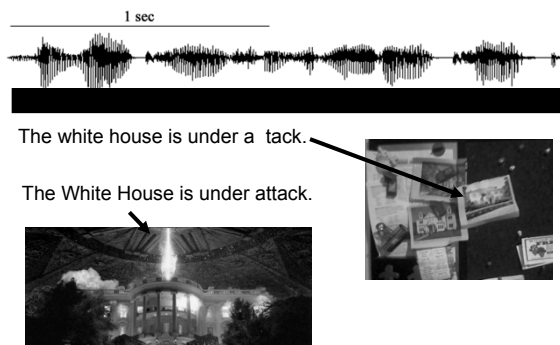
Word boundaries are not necessarily evident in the acoustic waveform



Segmentation mistakes from children

- Two dults
- [Two adults]
- I don't want to go to your ami!
- [I don't want to go to Miami]
- I am being have!
- [I am behaving!] (in response to "Behave!")
- Oh say can you see by the donzerly light?
- [Oh say can you see by the dawn's early light?]

Top-down influence



Top-down influence



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

The sky is falling!

OR

This guy is falling!

- **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?



Statistical Information Available

Maybe infants are sensitive to the statistical patterns contained in sequences of sounds.

"Over a corpus of speech 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

Maybe infants are sensitive to the statistical patterns contained in sequences of sounds.

"Over a corpus of speech 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 + st/e* is a common sound sequence

to the castle beyond the goblin city

Statistical Information Available

Maybe infants are sensitive to the statistical patterns contained in sequences of sounds.

“Over a corpus of speech 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)

Transitional Probability = Conditional Probability

$$\text{TrProb}(AB) = \text{Prob}(B | A)$$

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

$$\text{TrProb}(\text{"gob" "lin"}) = \text{Prob}(\text{"lin" | "gob"})$$

Read as “the probability of ‘lin’, given that ‘gob’ has just been encountered”

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)

Transitional Probability = Conditional Probability

$$\text{TrProb}(\text{"gob" "lin"}) = \text{Prob}(\text{"lin" | "gob"})$$

Example of how to calculate TrProb:

gob...

...ble, ...bler, ...bledygoon, ...let, ...lin, ...stopper
(6 options for what could follow “gob”)

$$\text{TrProb}(\text{"gob" "lin"}) = \text{Prob}(\text{"lin" | "gob"}) = 1/6$$

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)

Idea: $\text{Prob}(\text{"stle" | "ca"}) = \text{high}$

Why? “ca” is usually followed by “stle”

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)

Idea: Prob("be" | "stle") = lower
Why? "stle" is not usually followed by "be"

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)

Prob("yond" | "be") = higher
Why? "be" is commonly followed by "yond", among other options

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)


Prob("be" | "stle") < Prob("stle" | "ca")
Prob("be" | "stle") < Prob("yond" | "be")

to the castle **beyond** the goblin city

TrProb learner posits word boundary here,
at the minimum of the transitional probabilities

Important: doesn't matter what the probability actually is, so long as it's a minimum when compared to the probabilities surrounding it

Transitional Probability Example




$$0.1 < 0.5 \quad 0.1 < 0.3$$

0.1 = Transitional probability minimum,
compared with surrounding transitional
probabilities (0.5, 0.3)

Word boundary is here

Another Transitional Probability Example



0.7 = Transitional probability minimum, compared with surrounding transitional probabilities (0.8, 0.9)

Word boundary is here

8-month-old statistical learning

Saffran, Aslin, & Newport 1996

Familiarization-Preference Procedure (Jusczyk & Aslin 1995)

Habituation:

Infants exposed to auditory material that serves as potential learning experience

Test stimuli (tested immediately after familiarization):

(familiar) Items contained within auditory material

(novel) Items not contained within auditory material, but which are nonetheless highly similar to that material

8-month-old statistical learning

Saffran, Aslin, & Newport 1996

Familiarization-Preference Procedure (Jusczyk & Aslin 1995)

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), 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

Condition A:

tupiro, golabu, bidaku, padoti

Condition B:

dapiku, tilado, burobi, pagotu

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 intonation was monotone. There were no acoustic indicators of word boundaries.

Sample monotone speech: 

http://whyfiles.org/058language/images/baby_stream.aiff

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

$\text{TrProb}(\text{"tu" "pi"}) = 1.0$

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

$\text{TrProb}(\text{"tu" "pi"}) = 1.0 = \text{TrProb}(\text{"go" "la"}), \text{TrProb}(\text{"pa" "do"})$

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

$\text{TrProb}(\text{"ro" "go"}) < 1.0$ (0.3333...)

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

$\text{TrProb}(\text{"ro" "go"}), \text{TrProb}(\text{"ro" "pa"}) = 0.3333... <$
 $1.0 = \text{TrPrb}(\text{"pi" "ro"}), \text{TrProb}(\text{"go" "la"}), \text{TrProb}(\text{"pa" "do"})$

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

word boundary

word boundary

Testing Infant Sensitivity

Saffran, Aslin, & Newport 1996

Expt 1, test trial:

Each infant presented with repetitions of 1 of 4 words

2 were "real" words

(ex: *tupiro, golabu*)

2 were "fake" words whose syllables were jumbled up

(ex: *ropitu, bulago*)

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

Expt 1, test trial:

Each infant presented with repetitions of 1 of 4 words

2 were "real" words

(ex: *tupiro, golabu*)

2 were "fake" words whose syllables were jumbled up

(ex: *ropitu, bulago*)

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

Expt 1, results:

Infants listened longer to novel items (non-words)
(7.97 seconds for real words, 8.85 seconds for non-words)

Implication: Infants noticed the difference between real words and non-words from the artificial language after only 2 minutes of listening time!

Testing Infant Sensitivity

Saffran, Aslin, & Newport 1996

Expt 1, results:

Infants listened longer to novel items (non-words)
(7.97 seconds for real words, 8.85 seconds for non-words)

Implication: Infants noticed the difference between real words and non-words from the artificial language after only 2 minutes of listening time!

But why?

Could be that they just noticed a familiar sequence of sounds ("tupiro" familiar while "ropitu" never appeared), and didn't notice the differences in transitional probabilities.

Testing Infant Sensitivity

Saffran, Aslin, & Newport 1996

Expt 2, 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

Expt 2, 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

Expt 2, 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

Expt 2, 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.

Recap: Saffran, Aslin, & Newport (1996)

Experimental evidence suggests that 8-month-old infants can track statistical information such as the transitional probability between syllables. This can help them solve the task of word segmentation.

Evidence comes from testing children in an artificial language paradigm, with very short exposure time.



Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

- Infants use the prosody (rhythm) of an utterance to help them identify likely boundaries for words (sequences that cross utterance or clause boundaries are less likely to be words).
[Gout et al. 2004; Hirsh-Pasek et al. 1987; Jusczyk et al. 1992; Gerken et al. 1994; Nazzi et al. 2000; Seidl 2007]

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

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

- Infants use the prosody (rhythm) of an utterance to help them identify likely boundaries for words (sequences that cross utterance or clause boundaries are less likely to be words).
[Gout et al. 2004; Hirsh-Pasek et al. 1987; Jusczyk et al. 1992; Gerken et al. 1994; Nazzi et al. 2000; Seidl 2007]

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

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

- Infants use the prosody (rhythm) of an utterance to help them identify likely boundaries for words (sequences that cross utterance or clause boundaries are less likely to be words).
[Gout et al. 2004; Hirsh-Pasek et al. 1987; Jusczyk et al. 1992; Gerken et al. 1994; Nazzi et al. 2000; Seidl 2007]

"I went to the castle beyond the goblin city, which was very hard to get to, I saw the goblin king."
{pause} Not crossing a clause or utterance boundary - more likely to be a word

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

- Infants use the prosody (rhythm) of an utterance to help them identify likely boundaries for words (sequences that cross utterance or clause boundaries are less likely to be words).
[Gout et al. 2004; Hirsh-Pasek et al. 1987; Jusczyk et al. 1992; Gerken et al. 1994; Nazzi et al. 2000; Seidl 2007]

"I went to the castle beyond the goblin city, which was very hard to get to, I saw the goblin king."
{pause} Crossing a clause boundary - less likely to be a word

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

- Infants use the prosody (rhythm) of an utterance to help them identify likely boundaries for words (sequences that cross utterance or clause boundaries are less likely to be words).
[Gout et al. 2004; Hirsh-Pasek et al. 1987; Jusczyk et al. 1992; Gerken et al. 1994; Nazzi et al. 2000; Seidl 2007]

"I went to the castle beyond the goblin city, which was very hard to get to, I saw the goblin king."
{pause} Crossing an utterance boundary - less likely to be a word

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

- Infants distinguish between stressed and unstressed syllables, and they learn language-specific biases. English infants prefer words to begin with stress (Jusczyk et al. 1993, Jusczyk et al. 1999) while French infants prefer words to end with stress (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."
{pause}

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

- Infants distinguish between stressed and unstressed syllables, and they learn language-specific biases. English infants prefer words to begin with stress (Jusczyk et al. 1993, Jusczyk et al. 1999) while French infants prefer words to end with stress (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."
{pause} Pretty good strategy for English...

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

- Infants distinguish between stressed and unstressed syllables, and they learn language-specific biases. English infants prefer words to begin with stress (Jusczyk et al. 1993, Jusczyk et al. 1999) while French infants prefer words to end with stress (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."
{pause} ...though it's not perfect

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

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.

Other useful strategies

In addition to statistical information, infants appear to also use other cues to help them identify words in fluent speech.

Some evidence that this is the right sequence of events:

Thiessen & Saffran (2003) found that 6-month-olds prefer to segment using statistical cues (like transitional probability), but 9-month-olds prefer to use lexical stress cues. This suggests that infants first rely on statistical cues, and use the proto-lexicon derived from these statistical cues to infer the appropriate lexical stress bias.

Recap: Other useful strategies

Besides statistical cues to word segmentation, infants are apparently sensitive to prosodic cues such as clause and utterance boundaries, and also lexical stress patterns.

It seems that some of the lexical stress cues infants use are language-specific, so these cues are probably not used initially. Instead, these cues may be derived from the proto-lexicons infants have after using statistical cues.



Questions?



You should be able to do up through question 3 on HW2 and up through question 6 on the word segmentation review questions.