Psych156A/ Ling150
Spring 2016
Review Questions: Speech Segmentation

(1) Terms/concepts to know: speech segmentation, word segmentation, transitional probability, transitional probability minimum, isolated words, prosody, vowel harmony, proto-lexicon, psychological plausibility, precision, recall, F-score, CHILDES database, algebraic learning, Bayesian inference, oversegmentation, undersegmentation

(2) What was Saffran, Aslin, and Newport (1996)'s belief about the relation between transitional probability and word boundaries? That is, when did they believe transitional probability between syllables was likely to be higher and when did they believe it was likely to be lower?

(3) What is the difference between a transitional probability minimum and a low transitional probability? It is possible to have one without the other – briefly explain how.

(4) Why was it necessary for Saffran et al. (1996) to test children on real words vs. part-words from the artificial language? That is, why was the second experiment necessary to make their point about children's ability to track transitional probabilities for word segmentation?

(5) Are statistical cues (like transitional probability) alone enough for young infants to segment more realistic speech? (Hint: Are words in real languages all the same length? What happens when infants try to use just transitional probability to segment artificial languages that have words of different lengths?) What other cues seem to be helpful, especially when used in combination with statistical cues?

(6) How do English infants use stress to help them figure out where word boundaries might be? What about French infants? How would English and French infants end up coming up with slightly different strategies? (Hint: What do they base their bias on?) Unfortunately, this stress-based strategy doesn't work all the time – give an example in English where using stress the way infants do would fail to find the correct word boundaries.

(7) Which seems to come first – statistical learning strategies or stress-based strategies for word segmentation? What experimental evidence is there that demonstrates this?

(8) What does it mean for a learning algorithm to be psychologically plausible? Why would it be important for a computational model to use a psychologically plausible learning algorithm?

(9) Where did Gambell & Yang (2006) get the input for their model from? Is this a good source of realistic input that a child might hear? Why or why not?

(10) How did Gambell & Yang 2006 convert the transcripts of child-directed speech into spoken words that have stress patterns? Is the resulting pronounced speech likely to be the same as if someone were naturally speaking these words?

(11) Gambell & Yang found that tracking transitional probabilities failed to reliably segment child-directed speech data. What property of the data set caused the model to fail in Gambell & Yang's (2006) study? Why did this cause the transitional probability model to fail?

(12) Here are some words from the imaginary Guin language:

pengo   zu   pencrom   pentanor   az

(a) Where would an algebraic learner put word boundaries in the syllable sequence below?

Sequence: pen go az la to pen   crom   mer tem pen ta nor

Here are the same Guin words, along with their stress patterns:
(Note: pén = the "pen" syllable has stress)

péngo   zú   péncrom   péntanor   áz

(b) Where would an algebraic learner with knowledge of the Unique Stress Constraint put word boundaries in the syllable sequence below?

Sequence: pén go áz lá tó pén   crom   mér tem pén ta nor

(13) Does using statistical learning by itself always yield poor performance for word segmentation? Cite evidence to support your answer.

(14) Suppose a Bayesian learner is attempting to segment the utterance below and is trying to choose between three different segmentations:

Utterance: "Look at the pirate captain! Look at the flying boy!"

(i) Segmentation 1:
    "Lookat the pirate captain! Lookat the flying boy!"
    Words: lookat, the, pirate, captain, flying, boy
    # of words: 6
    length of words: between 1 and 2 syllables, average = 1.7
(ii) Segmentation 2:
    "Look atthe piratecaptain! Look atthe flyingboy!"
    Words: look, atthe, piratecaptain, flyingboy
    # of words: 4
    length of words: between 1 and 4 syllables, average = 2.5
(iii) Segmentation 3:
    "Lookat the pirate captain! Look atthe flying boy!"

Words: lookat, the, pirate, captain, look, atthe, flying, boy
# of words: 8
length of words: between 1 and 2 syllables, average = 1.25

(a) Assume that the Bayesian learner has two preferences: shorter words and fewer words. However, it values fewer words over shorter words. Given these preferences, would the Bayesian learner likely prefer segmentation 1 over segmentation 3, or instead prefer segmentation 3 over segmentation 1? Why? What about if it valued shorter words over fewer words?

(b) Suppose the Bayesian learner assigns the following probabilities to each of the segmentations:
     (i) p(segmentation 1) = 0.5
     (ii) p(segmentation 2) = 0.2
     (iii) p(segmentation 3) = 0.3
Which segmentation will the Bayesian learner choose? Why?

(15) What evidence is there that a Bayesian segmentation strategy is useful not just for English but for other languages as well?

(16) Why might we not want to evaluate a model's segmentation output against how an adult would segment the same data? (Hint: Think about how an adult's knowledge of language compares to a seven-month-old's.)

(17) What is an example of an oversegmentation that could still be useful to a learner of English? What about an example of an undersegmentation?

(18) If reasonable oversegmentations and undersegmentations are counted as correct, does the Bayesian segmentation strategy still succeed cross-linguistically? What about an algebraic learning strategy like the one Lignos (2012) investigated?