

HW10: Assorted Topics

1 Speech and Machine Translation

1. Statistical approaches to both speech recognition and machine translation make use of the notion of *alignments*. Why are alignments important in these tasks? How do they differ between speech and MT?
2. **(Bonus for all)** The “recognition” (or “decoding”) step in speech recognition can be performed using HMM technology. This is not immediately true in machine translation. Why not?
3. Suppose we were to try to automatically induce alignments for the following data. What alignments do you imagine the system would settle on? Based on these, do you think there is enough information in this small amount of data to automatically induce alignments? What is left unknown? (Languages are labeled “A” and “B”.)

A1. bpm uiv miba B1. pqzc oi bijmzc	A2. bpm uiv ibm B2. pqzc oi bijmbi
A3. q ibm zqkm B3. jwsc oi owpiv ew bijmbi	A4. q lzqvs jmmz B4. jwsc oi jqzc ew vwuc

2 Machine Learning

Alice decides to build a naive Bayes classifier to distinguish between emails from Professor Bob and Professor Clarence. She has collected the following examples of emails from these two Professors. She simply uses words as features. Compute every parameter for the naive Bayes classifier using maximum likelihood and classify the final test examples.

Bob	all students did great on this assignment
Bob	students should come to my office
Bob	should you need help talk to the ta
Bob	the ta did great grading this assignment
Clarence	no one did this assignment on time
Clarence	all students should fail
Clarence	the assignment is graded by the ta

Test example 1: “you did great”

Test example 2: “no students should fail”

Did the classifier do what you think it should? If not, why not?

(Bonus for all) Recompute and reclassify the test example using Laplace smoothing rather than maximum likelihood. Did the classifier do what you think it should? If not, why not? Did the classifications change?