

CS 552/652: Hidden Markov Models for Speech Recognition
Spring, 2006
John-Paul Hosom ('hosom' at cslu.ogi.edu)

Schedule:

Mondays and Wednesdays from 12:00 pm to 1:30 pm in Room WCC 403

Grading:

| | |
|-------------------------------------|-----|
| Project 1: Dynamic Time Warping | 15% |
| Project 2: Viterbi Search | 20% |
| Project 3: Expectation-Maximization | 25% |
| Midterm (in-class) | 20% |
| Final (take-home) | 20% |

Web Access:

<http://www.cse.ogi.edu/class/cse552/>

Books:

Recommended Textbooks:

Fundamentals of Speech Recognition
Lawrence Rabiner and Biing-Hwang Juang
Prentice Hall, New Jersey, 1993.

Statistical Methods for Speech Recognition
Frederick Jelinek
The MIT Press, Cambridge, MA, 1999

Other Recommended Readings:

Spoken Language Processing: A Guide to Theory, Algorithm, and System Development
(Xuedong Huang, Alex Acero, Hsiao-Wuen Hon, Prentice Hall, 2001)

Large Vocabulary Continuous Speech Recognition
(Steve Young, 1996)

Survey of the State of the Art in Human Language Technology
(Cole et al., 1996) <http://cslu.cse.ogi.edu/HLTsurvey/>

Probability & Statistics for Engineering and the Sciences
(Jay L. Devore, Brooks/Cole Publishing, California, 1982)

Course Content

Lecture 1 April 3

- Course Overview
- Why Is Automatic Speech Recognition Difficult?
- Background: Speech Production, Representations of Speech, Models of Human Speech Recognition

Lecture 2 April 5

- General Issues in Developing ASR Systems
- Induction (review)
- DTW Motivation / Algorithm / Implementation
- DTW Examples

(Assign DTW Project)

Lecture 3 April 10

- Quick Review of Probability and Statistics
- What is a Markov Model?
- Log-Domain Mathematics

Lecture 4 April 12

- What is a Hidden Markov Model?
- HMM Topologies
- Vector Quantization

Lecture 5 April 17

- Gaussian Mixture Models
- Features: Spectrum, Cepstrum, LPC, Pre-Emphasis
- Features: Mel & Bark Scales, PLP, MFCC, EIH
- Noise Removal: CMS, RASTA

Lecture 6 April 19

- HMMs for Speech Recognition
- HMMs for Word Recognition
- Review: Framework of HMMs

Lecture 7 April 24

- Viterbi Search: Formula
- Viterbi Search: Algorithm

(DTW Project Due)

Lecture 8 April 26

- Duration modeling
- Initializing an HMM

(Assign Viterbi Project)

Lecture 9 May 1

- Review DTW Project
- Viterbi / K-Means Segmentation
- Multiple Training Files

Lecture 10 May 3

- Probability Evaluation: Alpha & Beta

- Forward Procedure
- Backward Procedure
- Forward and Backward Examples

Lecture 11 May 8

- In-Class Midterm

Lecture 12 May 10

- Gamma and Xi
- Forward-Backward Algorithm

Lecture 13 May 15

- Expectation-Maximization
 - Embedded Training
- (Viterbi Project Due)

Lecture 14 May 17

- Search: One-Pass Algorithm
 - Search Alternative: 2-Level Algorithm
 - Search Alternative: Level-Building Algorithm
- (Assign Forward-Backward Project)

Lecture 15 May 22

- Review Viterbi Project
- Null States
- Beam Search
- Grammar / Tree Search
- “On-Line” Processing
- Balancing Insertions and Deletions
- N-Best Outputs

Lecture 16 May 24

- Semi-Continuous HMMs
- State Tying/Clustering
- Cloning
- Pause Models

May 29 is Memorial Day; no class

Lecture 17 May 31

- N-Gram LMs
- LM Smoothing
- LM Discounting & Back-Off

Lecture 18 June 5

- Segment-Based Systems
- HMM/ANN Hybrids
- Experimental Systems: TRAPS, SPAM, etc

Lecture 19 June 7

- LM Look-Ahead
- ABNF Grammars
- Evaluation and State-of-the-Art
- Summary

(Forward-Backward Project Due, Hand Out Final Exam)