EE E6820: Speech \& Audio Processing \& Recognition

## Lecture 10: <br> ASR: Sequence Recognition

(1) Signal template matching
(2) Statistical sequence recognition
(3) Acoustic modeling
(4) The Hidden Markov Model (HMM)


## Signal template matching

- Framewise comparison of unknown word and stored templates:


Test


- distance metric?
- comparison between templates?
- constraints?


## Dynamic Time Warp (DTW)

- Find lowest-cost constrained path:
- matrix d(i,j) of distances between input frame $f_{i}$ and reference frame $r_{j}$
- allowable predecessors \& transition costs $\mathrm{T}_{\mathrm{xy}}$

Lowest cost to (i,j)


- Best path via traceback from final state
- have to store predecessors for (almost) every (i,j)


## DTW-based recognition

- Reference templates for each possible word
- Isolated word:
- mark endpoints of input word
- calculate scores through each template (+prune)
- choose best
- Continuous speech
- one matrix of template slices; special-case constraints at word ends



## DTW-based recognition (2)

+ Successfully handles timing variation
+ Able to recognize speech at reasonable cost
- Distance metric?
- pseudo-Euclidean space?
- Warp penalties?
- How to choose templates?
- several templates per word?
- choose 'most representative'?
- align and average?
$\rightarrow$ need a rigorous foundation...


## Outline

(1) Signal template matching
(2) Statistical sequence recognition

- state-based modeling
(3) Acoustic modeling
(4) The Hidden Markov Model (HMM)



## (2) Statistical sequence recognition

- DTW limited because it's hard to optimize
- interpretation of distance, transition costs?
- Need a theoretical foundation: Probability
- Formulate as MAP choice among models:

$$
M^{*}=\underset{M_{j}}{\operatorname{argmax}} p\left(M_{j} \mid X, \Theta\right)
$$

- $X=$ observed features
- $M_{j}=$ word-sequence models
- $\Theta=$ all current parameters


## Statistical formulation (2)

- Can rearrange via Bayes' rule (\& drop $p(X)$ ):

$$
\begin{aligned}
M^{*} & =\underset{M_{j}}{\operatorname{argmax}} p\left(M_{j} \mid X, \Theta\right) \\
& =\underset{M_{j}}{\operatorname{argmax}} p\left(X \mid M_{j}, \Theta_{A}\right) p\left(M_{j} \mid \Theta_{L}\right)
\end{aligned}
$$

- $p\left(X \mid M_{j}\right)=$ likelihood of obs'v'ns under model
- $p\left(M_{j}\right)=$ prior probability of model
- $\Theta_{A}=$ acoustics-related model parameters
- $\Theta_{L}=$ language-related model parameters
- Questions:
- what form of model to use for $p\left(X \mid M_{j}, \Theta_{A}\right)$ ?
- how to find $\Theta_{A}$ (training)?
- how to solve for $M_{j}$ (decoding)?


## State-based modeling

- Assume discrete-state model for the speech:
- observations are divided up into time frames
- model $\rightarrow$ states $\rightarrow$ observations:

Model $M_{j}$


- Probability of observations given model is:

$$
p\left(X \mid M_{j}\right)=\sum_{\text {all } Q_{k}} p\left(X_{1}^{N} \mid Q_{k}, M_{j}\right) \cdot p\left(Q_{k} \mid M_{j}\right)
$$

- sum over all possible state sequences $Q_{k}$
- How do observations depend on states? How do state sequences depend on model?


## The speech recognition chain

- After classification, still have problem of classifying the sequences of frames:

- Questions
- what to use for the acoustic classifier?
- how to represent 'model’ sequences?
- how to score matches?


## Outline

(1) Signal template matching
(2) Statistical sequence recognition
(3) Acoustic modeling

- defining targets
- neural networks \& Gaussian models
(4) The Hidden Markov Model (HMM)



## Acoustic Modeling

- Goal: Convert features into probabilities of particular labels:
i.e find $p\left(q_{n}^{i} \mid X_{n}\right)$ over some state set $\left\{q^{i}\right\}$
- conventional statistical classification problem
- Classifier construction is data-driven
- assume we can get examples of known good Xs for each of the $q^{i}$ s
- calculate model parameters by standard training scheme
- Various classifiers can be used
- GMMs model distribution under each state
- Neural Nets directly estimate posteriors
- Different classifiers have different properties
- features, labels limit ultimate performance


## Defining classifier targets

- Choice of $\left\{q^{i}\right\}$ can make a big difference
- must support recognition task
- must be a practical classification task
- Hand-labeling is one source...
- 'experts' mark spectrogram boundaries
- ...Forced alignment is another
- 'best guess' with existing classifiers, given words
- Result is targets for each training frame:



## Forced alignment

- Best labeling given existing classifier constrained by known word sequence



## Gaussian Mixture Models vs. Neural Nets

- GMMs fit distribution of features under states:
- separate 'likelihood' model for each state $q^{i}$
$p\left(\mathbf{x} \mid q^{k}\right)=\frac{1}{(\sqrt{2 \pi})^{d}\left|\Sigma_{k}\right|^{1 / 2}} \cdot \exp \left[-\frac{1}{2}\left(\mathbf{x}-\boldsymbol{\mu}_{k}\right)^{T} \Sigma_{k}^{-1}\left(\mathbf{x}-\mu_{k}\right)\right]$
- match any distribution given enough data
- Neural nets estimate posteriors directly

$$
p\left(q^{k} \mid \mathbf{x}\right)=F\left[\sum_{j} w_{j k} \cdot F\left[\sum_{j} w_{i j} x_{i}\right]\right]
$$

- parameters set to discriminate classes
- Posteriors \& likelihoods related by Bayes' rule:

$$
p\left(q^{k} \mid \mathbf{x}\right)=\frac{p\left(\mathbf{x} \mid q^{k}\right) \cdot \operatorname{Pr}\left(q^{k}\right)}{\sum_{j} p\left(\mathbf{x} \mid q^{j}\right) \cdot \operatorname{Pr}\left(q^{j}\right)}
$$

## Outline

(1) Signal template matching
(2) Statistical sequence recognition
(3) Acoustic classification
(4) The Hidden Markov Model (HMM)

- generative Markov models
- hidden Markov models
- model fit likelihood
- HMM examples



## 3 Markov models

- A (first order) Markov model is a finite-state system whose behavior depends only on the current state
- E.g. generative Markov model:


SAAAAAAAABBBBBBBBBCCCCBBBBBBCE

## Hidden Markov models

- Markov models where state sequence $Q=\left\{q_{n}\right\}$ is not directly observable (= 'hidden')
- But, observations $X$ do depend on $Q$ :
- $x_{n}$ is rv that depends on current state: $p(x \mid q)$

State sequence


AAAAAAAABBBBBBBBBBBCCCCBBBBBBBC



- can still tell something about state seq...


## (Generative) Markov models (2)

- HMM is specified by:
- transition probabilities $p\left(q_{n}^{j} \mid q_{n-1}^{i}\right) \equiv a_{i j}$
- (initial state probabilities $\left.p\left(q_{1}^{i}\right) \equiv \pi_{i}\right)$
- emission distributions $p\left(x \mid q^{i}\right) \equiv b_{i}(x)$
- states $q^{i}$
$\odot$ ( ${ }^{\text {® (a) (t) } \odot ~}$
- transition
probabilities $a_{i j}$


|  | k | a | t | $\cdot$ |
| :---: | :---: | :---: | :---: | :---: |
| $\cdot$ | 1.0 | 0.0 | 0.0 | 0.0 |
| k | 0.9 | 0.1 | 0.0 | 0.0 |
| a | 0.0 | 0.9 | 0.1 | 0.0 |
| t | 0.0 | 0.0 | 0.9 | 0.1 |

- emission
distributions $b_{i}(x)$



## Markov models for speech

- $\quad$ Speech models $M_{j}$
- typ. left-to-right HMMs (sequence constraint)
- observation \& evolution are conditionally independent of rest given (hidden) state $q_{n}$


- self-loops for time dilation



## Markov models for sequence recognition

- Independence of observations:
- observation $x_{n}$ depends only current state $q_{n}$

$$
\begin{aligned}
p(X \mid Q) & =p\left(x_{1}, x_{2}, \ldots x_{N} \mid q_{1}, q_{2}, \ldots q_{N}\right) \\
& =p\left(x_{1} \mid q_{1}\right) \cdot p\left(x_{2} \mid q_{2}\right) \cdot \ldots p\left(x_{N} \mid q_{N}\right) \\
& =\prod_{n=1}^{N} p\left(x_{n} \mid q_{n}\right)=\prod_{n=1}^{N} b_{q_{n}}\left(x_{n}\right)
\end{aligned}
$$

- Markov transitions:
- transition to next state $q_{i+1}$ depends only on $q_{i}$

$$
\begin{aligned}
& p(Q \mid M)=p\left(q_{1}, q_{2}, \ldots q_{N} \mid M\right) \\
& =p\left(q_{N} \mid q_{1} \ldots q_{N-1}\right) p\left(q_{N-1} \mid q_{1} \ldots q_{N-2}\right) \ldots p\left(q_{2} \mid q_{1}\right) p\left(q_{1}\right) \\
& =p\left(q_{N} \mid q_{N-1}\right) p\left(q_{N-1} \mid q_{N-2}\right) \ldots p\left(q_{2} \mid q_{1}\right) p\left(q_{1}\right) \\
& =p\left(q_{1}\right) \prod_{n=2}^{N} p\left(q_{n} \mid q_{n-1}\right)=\pi_{q_{1}} \prod_{n=2}^{N} a_{q_{n-1} q_{n}}
\end{aligned}
$$

## Model fit calculation

- From 'state-based modeling':

$$
p\left(X \mid M_{j}\right)=\sum_{\text {all } Q_{k}} p\left(X_{1}^{N} \mid Q_{k}, M_{j}\right) \cdot p\left(Q_{k} \mid M_{j}\right)
$$

- For HMMs:

$$
\begin{aligned}
p(X \mid Q) & =\prod_{n=1}^{N} b_{q_{n}}\left(x_{n}\right) \\
p(Q \mid M) & =\pi_{q_{1}} \cdot \prod_{n=2}^{N} a_{q_{n-1} q_{n}}
\end{aligned}
$$

- Hence, solve for $M^{*}$ :
- calculate $p\left(X \mid M_{j}\right)$ for each available model, scale by prior $p\left(M_{j}\right) \rightarrow p\left(M_{j} \mid X\right)$
- Sum over all $Q_{k}$ ???


## Summing over all paths



All possible 3-emission paths $Q_{k}$ from S to E


## The 'forward recursion'

- Dynamic-programming-like technique to calculate sum over all $Q_{k}$
- Define $\alpha_{n}(i)$ as the probability of getting to state $q^{i}$ at time step $n$ (by any path):

$$
\alpha_{n}(i)=p\left(x_{1}, x_{2}, \ldots x_{n}, q_{n}=q^{i}\right) \equiv p\left(X_{1}^{n}, q_{n}^{i}\right)
$$

- Then $\alpha_{n+1}(j)$ can be calculated recursively:



## Forward recursion (2)

- Initialize $\alpha_{1}(i)=\pi_{i} \cdot b_{i}\left(x_{1}\right)$
- Then total probability $p\left(X_{1}^{N} \mid M\right)=\sum_{i=1}^{S} \alpha_{N}(i)$
$\rightarrow$ Practical way to solve for $p\left(X \mid M_{j}\right)$ and hence perform recognition



## Optimal path

- May be interested in actual $q_{n}$ assignments
- which state was 'active' at each time frame
- e.g. phone labelling (for training?)
- Total probability is over all paths...
- ... but can also solve for single best path
= "Viterbi" state sequence
- Probability along best path to state $q_{n+1}^{j}$ :

$$
\alpha_{n+1}^{*}(j)=\left[\max _{i}\left\{\alpha_{n}^{*}(i) a_{i j}\right\}\right] \cdot b_{j}\left(x_{n+1}\right)
$$

- backtrack from final state to get best path
- final probability is product only (no sum)
$\rightarrow$ log-domain calculation just summation
- Total probability often dominated by best path:

$$
p\left(X, Q^{*} \mid M\right) \approx p(X \mid M)
$$

## Interpreting the Viterbi path

- Viterbi path assigns each $x_{n}$ to a state $q^{i}$
- performing classification based on $b_{i}(x)$
- ... at the same time as applying transition constraints $a_{i j}$

Inferred classification


Viterbi labels: ааааааавввввввввввссссвввввввс

- Can be used for segmentation
- train an HMM with 'garbage' and 'target' states
- decode on new data to find 'targets', boundaries
- Can use for (heuristic) training
- e.g. train classifiers based on labels...


## Recognition with HMMs

- Isolated word
- choose best $p(M \mid X) \propto p(X \mid M) p(M)$

- Continuous speech
- Viterbi decoding of one large HMM gives words



## HMM example: Different state sequences

Model $M_{1}$


Model $M_{2}$


Emission distributions


## Model inference:

 Emission probabilities

Model $M_{2}$


## Model inference:

 Transition probabilities

state alignment
log obs.l'hood $\log p\left(X \mid Q^{*}, M\right)=-26.0$


Model $M^{\prime}{ }_{1}$| $\log p(X$ | $M)=-32.2$ |
| ---: | :--- |
| $\log p\left(X, Q^{*}\right.$ | $M)=-33.6$ | log trans.prob $\log p\left(Q^{*} \mid M\right)=-7.6$



Model $\left.M^{\prime}{ }_{2}^{\log p(X} \begin{array}{r}M) \\ \log p\left(X, Q^{*}\right.\end{array} \right\rvert\, \begin{aligned} & M)=-34.5\end{aligned}$ log trans.prob

\[
\log p\left(Q^{*} \mid M\right)=-8.9

\]| 0 |
| :--- |
| -1 |
| -2 |
| -2 |

## Validity of HMM assumptions

- Key assumption is conditional independence: Given $q^{i}$, future evolution \& obs. distribution are independent of previous events
- duration behavior: self-loops imply exponential distribution


- independence of successive $x_{n} \mathrm{~s}$


$$
p(X)=\prod_{p\left(x_{n} \mid q^{i}\right) ?}
$$

## Recap: Recognizer Structure



- Know how to execute each state
- .. training HMMs?
- .. language/word models


## Summary

- Speech is modeled as a sequence of features
- need temporal aspect to recognition
- best time-alignment of templates = DTW
- Hidden Markov models are rigorous solution
- self-loops allow temporal dilation
- exact, efficient likelihood calculations

