## CS440/ECE448 Lecture 20: Hidden Markov Models

Slides by Svetlana Lazebnik, 11/2016
Modified by Mark Hasegawa-Johnson, 3/2019
Hidden Markov Model $=\left[\begin{array}{lllll}\text { [ } & \text { i } & \mathrm{j} & \mathrm{k} & \cdots\end{array}\right]$
State Sequence $\mathrm{Q}=\left[\begin{array}{lllllllll}\mathrm{i} & \mathrm{i} & \mathrm{i} & \mathrm{j} & \mathrm{j} & \mathrm{k} & \mathrm{k} & \mathrm{k} & \ldots\end{array}\right]$
Observations $\mathrm{O}=\left[\begin{array}{lllllllll}\mathrm{o}_{1} & \mathrm{o}_{2} & \mathrm{o}_{3} & \mathrm{o}_{4} & \mathrm{o}_{5} & \mathrm{o}_{6} & o_{7} & \mathrm{o}_{8} & o_{9}\end{array}\right]$


## Probabilistic reasoning over time

- So far, we've mostly dealt with episodic environments
- Exceptions: games with multiple moves, planning
- In particular, the Bayesian networks we've seen so far describe static situations
- Each random variable gets a single fixed value in a single problem instance
- Now we consider the problem of describing probabilistic environments that evolve over time
- Examples: robot localization, human activity detection, tracking, speech recognition, machine translation,


## Hidden Markov Models

- At each time slice $t$, the state of the world is described by an unobservable variable $X_{t}$ and an observable evidence variable $\mathrm{E}_{t}$
- Transition model: distribution over the current state given the whole past history: $P\left(X_{t} \mid X_{0}, \ldots, X_{t-1}\right)=P\left(X_{t} \mid X_{0: t-1}\right)$
- Observation model: $P\left(E_{t} \mid X_{0: t}, E_{1: t-1}\right)$



## Hidden Markov Models

- Markov assumption (first order)
- The current state is conditionally independent of all the other states given the state in the previous time step
- What does $\mathrm{P}\left(\mathrm{X}_{t} \mid \mathrm{X}_{0: t-1}\right)$ simplify to?

$$
P\left(X_{t} \mid X_{0: t-1}\right)=P\left(X_{t} \mid X_{t-1}\right)
$$

- Markov assumption for observations
- The evidence at time $t$ depends only on the state at time $t$
- What does $\mathrm{P}\left(\mathrm{E}_{t} \mid \mathbf{X}_{0: t}, \mathrm{E}_{1: t-1}\right)$ simplify to?

$$
\mathrm{P}\left(\mathrm{E}_{t} \mid \mathbf{X}_{0: t}, \mathrm{E}_{1: t-1}\right)=\mathrm{P}\left(\mathrm{E}_{t} \mid \mathrm{X}_{t}\right)
$$



## Example



## Example

Transition model


## An alternative visualization



Transition
probabilities

|  | $R_{t}=T$ | $R_{t}=F$ |
| :---: | :---: | :---: |
| $R_{t-1}=T$ | 0.7 | 0.3 |
| $R_{t-1}=F$ | 0.3 | 0.7 |

Observation
(emission)
probabilities

|  | $U_{t}=T$ | $U_{t}=F$ |
| :---: | :---: | :---: |
| $R_{t}=T$ | 0.9 | 0.1 |
| $R_{t}=F$ | 0.2 | 0.8 |

## Another example

- States: $\mathrm{X}=$ \{home, office, cafe $\}$
- Observations: $\mathrm{E}=\{\mathrm{sms}$, facebook, email\}


Transition Probabilities

|  | home | office | cafe |
| ---: | :---: | :---: | :---: |
| home | 0.2 | 0.6 | 0.2 |
| office | 0.5 | 0.2 | 0.3 |
| cafe | 0.2 | 0.8 | 0.0 |


| Emission Probabilities |  |  |  |
| :---: | :---: | :---: | :---: |
|  sms facebook email <br> home 0.3 0.5 0.2 <br> office 0.1 0.1 0.8 <br> cafe 0.8 0.1 0.1 |  |  |  |

Slide credit: Andy White

## The Joint Distribution

- Transition model: $\mathrm{P}\left(\mathrm{X}_{t} \mid \mathrm{X}_{0: t-1}\right)=\mathrm{P}\left(\mathrm{X}_{t} \mid \mathrm{X}_{t-1}\right)$
- Observation model: $\mathrm{P}\left(\mathrm{E}_{t} \mid \mathrm{X}_{0: t}, \mathrm{E}_{1: t-1}\right)=\mathrm{P}\left(\mathrm{E}_{t} \mid \mathrm{X}_{t}\right)$
- How do we compute the full joint $P\left(X_{0: t}, E_{1: t}\right)$ ?

$$
P\left(\boldsymbol{X}_{0: t}, \boldsymbol{E}_{1: t}\right)=P\left(X_{0}\right) \prod_{i=1}^{t} P\left(X_{i} \mid X_{i-1}\right) P\left(E_{i} \mid X_{i}\right)
$$



## Review: Bayes net inference

- Inference:
- Trees: Sum-Product Algorithm (Textbook: "Variable Elimination" Algorithm)
- Other Nets: Junction Tree Algorithm (Textbook: "Join Tree" Algorithm)
- In General: NP-Complete, because clique size = graph size in general
- Parameter learning
- Fully observed: Count \# times each event occurs
- Partially observed: Expectation-Maximization algorithm
- Estimate Probability of each event at each time
- E[\# times event occurs] = sum_t(Probability event occurs at time t)


## Sum-Product Algorithm for HMMs

- An HMM is a tree!
- For example, suppose we want to find $\mathrm{P}(\mathrm{X} 3 \mid \mathrm{E} 1, \mathrm{E} 2, \mathrm{E} 3)$
- Product: P(X0,X1,E1)=P(X0)P(X1|X0)P(E1|X1)
- Sum: P(X1|E1)=P(X1,E1)/P(E1)
- Product: P(X1,X2,E2|E1)=P(X1|E1)P(X2|X1)P(E2|X2)
- Sum: P(X2|E1,E2)=P(X2,E2|E1)/P(E2|E1)
- ...



## HMM inference tasks

- Filtering: what is the distribution over the current state $X_{t}$ given all the evidence so far, $\mathbf{e}_{1: t}$ ?
- The forward algorithm = sum-product algorithm for Xt given e1:t



## HMM inference tasks

- Filtering: what is the distribution over the current state $X_{t}$ given all the evidence so far, $\mathbf{e}_{1: t}$ ?
- Smoothing: what is the distribution of some state $X_{k}$ given the entire observation sequence $\mathbf{e}_{1: \mathrm{t}}$ ?
- The forward-backward algorithm = sum-product algorithm for Xk given e1:t, when $1<k<t$
- Xk = query variable, unknown, need to consider all its possible values
- E1:t = evidence variables, known, only need to consider the given values



## HMM inference tasks

- Filtering: what is the distribution over the current state $X_{t}$ given all the evidence so far, $\mathbf{e}_{1: t}$ ?
- Smoothing: what is the distribution of some state $X_{k}$ given the entire observation sequence $\mathbf{e}_{1: \mathrm{t}}$ ?
- Evaluation: compute the probability of a given observation sequence $\mathbf{e}_{1: t}$



## HMM inference tasks

- Filtering: what is the distribution over the current state $X_{t}$ given all the evidence so far, $\mathbf{e}_{1: t}$
- Smoothing: what is the distribution of some state $X_{k}$ given the entire observation sequence $\mathbf{e}_{1: \mathrm{t}}$ ?
- Evaluation: compute the probability of a given observation sequence $\mathbf{e}_{1: t}$
- Decoding: what is the most likely state sequence $\mathbf{X}_{0: t}$ given the observation sequence $\mathbf{e}_{1: t}$ ?
- The Viterbi algorithm



## HMM Learning and Inference

- Inference tasks
- Filtering: what is the distribution over the current state $X_{t}$ given all the evidence so far, $\mathbf{e}_{1: \mathrm{t}}$
- Smoothing: what is the distribution of some state $X_{k}$ given the entire observation sequence $\mathbf{e}_{1: t}$ ?
- Evaluation: compute the probability of a given observation sequence $\mathbf{e}_{1: t}$
- Decoding: what is the most likely state sequence $\mathbf{X}_{0: t}$ given the observation sequence $\mathbf{e}_{1: t}$ ?
- Learning
- Given a training sample of sequences, learn the model parameters (transition and emission probabilities)
- EM algorithm


## Applications of HMMs

- Speech recognition HMMs:
- Observations are acoustic signals (continuous valued)
- States are specific positions in specific words (so, tens of thousands)

- Machine translation HMMs:
- Observations are words (tens of thousands)
- States are translation options

Google

Translate From: Latin $\quad \leftrightarrows \quad$ To: English •

- Robot tracking:
- Observations are range readings (continuous)
- States are positions on a map (continuous)



# Application of HMMs: Speech recognition 

- "Noisy channel" model of speech



## Speech feature extraction



## Speech feature extraction

Spectrogram


## Phonetic model

- Phones: speech sounds
- Phonemes: groups of speech sounds that have a unique meaning/function in a language (e.g., there are several different ways to pronounce " t ")


## Phonetic model

| IPA | ARPAbet |  | IPA | ARPAbet |
| :---: | :---: | :---: | :---: | :---: |
| Symbol | Symbol | Word | Transcription | Transcription |
| ［p］ | ［p］ | parsley | ［＇parsli］ | ［paarsliy］ |
| ［t］ | ［t］ | tarragon | ［＇tærəgan］ | ［t aeraxg aan］ |
| ［k］ | ［k］ | catnip | ［＇kætnip］ | ［k aetnix p］ |
| ［b］ | ［b］ | $\underline{\text { bay }}$ | ［ber］ | ［bey］ |
| ［d］ | ［d］ | dill | ［dil］ | ［dih I］ |
| ［g］ | ［g］ | garlic | ［＇garlik］ | ［gaarlix k］ |
| ［m］ | ［m］ | $\underline{\text { mint }}$ | ［mmt］ | ［ m ih nt ］ |
| ［ n ］ | ［n］ | nutmeg | ［＇nıtmeg］ | ［ n ahtmeh g |
| ［［］］ | ［ ng ］ | ginseng | ［＇dzmsin］ | ［jh ih nsix ng］ |
| ［f］ | ［f］ | fennel | ［＇fın！］］ | ［ feh nel ］ |
| ［v］ | ［v］ | clove | ［klouv］ | ［klow v］ |
| ［日］ | ［th］ | thistle | ［＇日is］］ | ［th ih sel］ |
| ［ $¢$ ］ | ［dh］ | heather | ［＇hとðə］ | ［ h eh dh axr］ |
| ［s］ | ［s］ | sage | ［serd3］ | ［s ey jh］ |
| ［z］ | ［z］ | hazelnut | ［＇herz｜nıt］ | ［heyzelnaht］ |
| ［［］ | ［sh］ | squash | ［skwaf］ | ［skwash］ |
| ［3］ | ［zh］ | ambrosia | ［æm＇brouzz］ | ［ae mbrow zh ax］ |
| ［t5］ | ［ch］ | chicory | ［＇tfikxi］ | ［ch ih k axriy ］ |
| ［d3］ | ［jh］ | sage | ［serd3］ | ［s ey jh］ |
| ［1］ | ［I］ | licorice | ［＇İkəif］ | ［lih k axr ix sh］ |
| ［w］ | ［w］ | kiwi | ［＇kiwi］ | ［kiy w iy］ |
| ［r］ | ［r］ | parsley | ［＇parsli］ | ［paarsliy］ |
| ［j］ | ［y］ | yew | ［yu］ | ［y uw］ |
| ［h］ | ［h］ | horseradish | ［＇horsrædis］ | ［h aorsraedih sh］ |
| ［？］ | ［q］ | uh－oh | ［？$\wedge$ ？ 0 ］ | ［ q ah q ow］ |
| ［r］ | ［dx］ | butter | ［＇bara］ | ［b ah dx axr ］ |
| ［ r$]$ | ［ nx ］ | wintergreen | ［wזัə ${ }^{\text {chin］}}$ | ［wihnxaxrgrin］ |
| ［1］ | ［el］ | thistle | ［＇Eisl］ | ［th ih sel ］ |



Figure 4．1 IPA and ARPAbet symbols for transcription of English consonants．

## HMM models for phones

HMM states in most speech recognition systems correspond to subsegments of triphones

- Triphone: the /b/ in "about" (ax-b+aw) sounds different from the /b/ in "Abdul" (ae-b+d). There are around 60 phones and as many as $60^{3}$ context-dependent triphones.
- Subsegments: /b/ has three subsegments: the closure, the silence, and the release. There are $3 \times 60^{3}$ subsegments of triphones.


Figure 7.11 An example of the context-dependent triphone $b(a x, a w)$ (the phone [b] preceded by a [ax] and followed by a [aw], as in the beginning of about, showing its left, middle, and right subphones.

## HMM models for words



Figure 7.5 Pronunciation networks for the words I, on, need, and the. All networks (especially the) are significantly simplified.

## Putting words together



- Given a sequence of acoustic features, how do we find the corresponding word sequence?


## The Viterbi Algorithm

$$
\begin{aligned}
& \max _{\boldsymbol{X}_{0: t}} P\left(\boldsymbol{X}_{0: t}, \boldsymbol{E}_{0: t}\right) \\
& =\max _{X_{t}} P\left(E_{t} \mid X_{t}\right) \max _{X_{t-1}} P\left(X_{t} \mid X_{t-1}\right) P\left(E_{t-1} \mid X_{t-1}\right) \max _{X_{t-2}} \ldots
\end{aligned}
$$

Complexity changes from $\mathrm{O}\left\{\mathrm{N}^{\wedge} \mathrm{T}\right\}$ to $\mathrm{O}\left\{\mathrm{TN}^{\wedge} 2\right\}$


## Decoding with the Viterbi algorithm



Figure 7.10 The entries in the individual state columns for the Viterbi algorithm. Each cell keeps the probability of the best path so far and a pointer to the previous cell along that path. Backtracing from the successful last word (the), we can reconstruct the word sequence I need the.

## For more information

- CS 447: Natural Language Processing
- ECE 417: Multimedia Signal Processing
- ECE 594: Mathematical Models of Language
- Linguistics 506: Computational Linguistics
- D. Jurafsky and J. Martin, "Speech and Language Processing," $2^{\text {nd }}$ ed., Prentice Hall, 2008

