# Natural Language Processing 

Anoop Sarkar anoopsarkar.github.io/nlp-class

Simon Fraser University

October 18, 2018

# Natural Language Processing 

Anoop Sarkar<br>anoopsarkar.github.io/nlp-class

Simon Fraser University

Part 1: Introducing Hidden Markov Models

## Modelling pairs of sequences

| Input: sequ Input | ence of British | lerds | Outpu | sed | ence of Falkland | bels <br> Islands |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Output1 | N | N | V | P | N | N |
| Output2 | N | V | N | P | N | N |

N Noun, e.g. islands
$\checkmark$ Verb, e.g. leave, left
$P$ Preposition, e.g. on

## Modelling pairs of sequences

Input: sequence of words; Output: sequence of labels

| Input | British | left | waffles | on | Falkland | Islands |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| Output1 | N | N | V | P | N | N |
| Output2 | N | V | N | P | N | N |

- 3 states: $\mathcal{S}=\{N, V, P\}$
- Input sequence: $x_{1}, x_{2}, \ldots, x_{n}$
- Output sequence: $t_{1}, t_{2}, \ldots, t_{n}$ where $t_{i} \in \mathcal{S}$
- How many output sequences?
$\square$


## Modelling pairs of sequences

Input：sequence of characters；Output：sequence of labels Input 北京大学生比赛 7 chars Output1 BIBIIBI 7 labels Output2 BIIIBBI 7 labels<br>7 labels<br>B Begin word<br>I Inside word

BIBIIBI 北京—大学生—比赛（Beijing student competition）
BIIIBBI 北京大学—生—比赛（Peking University Health Competition）

## Hidden Markov Models

- Input: x
- Output space: $\mathcal{Y}(x)$
- Output: $y \in \mathcal{Y}(x)$
- We want to learn a function $f$ such that $f(x)=y$


## Hidden Markov Models

## Conditional model

- Construct function $f$ using a conditional probability:

$$
f(x)=\arg \max _{y \in \mathcal{Y}(x)} p(y \mid x)
$$

- We can construct this function $f$ using two principles:
- Discriminative learning: find the best output $y$ given input $x$
- Generative modelling: model the joint probability $p(x, y)$ to find $p(y \mid x)$


## Hidden Markov Models

## Generative Model

- Start from the joint probability $p(x, y)$ :

$$
p(x, y)=p(y) p(x \mid y)
$$

- Also:

$$
p(x, y)=p(x) p(y \mid x)
$$

Bayes Rule:

$$
p(y \mid x)=\frac{p(y) p(x \mid y)}{p(x)}
$$

## Hidden Markov Models

## Generative Model

- Bayes Rule:

$$
p(y \mid x)=\frac{p(y) p(x \mid y)}{p(x)}
$$

- where:

$$
p(x)=\sum_{y \in \mathcal{Y}(x)} p(x, y)=\sum_{y \in \mathcal{Y}(x)} p(y) p(x \mid y)
$$

- So using a generative model, we can find the best output $y$ using:

$$
p(y \mid x)=\frac{p(y) p(x \mid y)}{\sum_{y \in \mathcal{Y}(x)} p(y) p(x \mid y)}
$$

# Natural Language Processing 

Anoop Sarkar<br>anoopsarkar.github.io/nlp-class

Simon Fraser University

Part 2: Algorithms for Hidden Markov Models

## Hidden Markov Model

Model $\theta= \begin{cases}\pi_{i} & p(i): \text { starting at state } i \\ a_{i, j} & p(j \mid i): \text { transition to state } i \text { from state } j \\ b_{i}(o) & p(o \mid i): \text { output o at state } i\end{cases}$


## Hidden Markov Model Algorithms

- HMM as parser: compute the best sequence of states for a given observation sequence.
- HMM as language model: compute probability of given observation sequence.
- HMM as learner: given a corpus of observation sequences, learn its distribution, i.e. learn the parameters of the HMM from the corpus.
- Learning from a set of observations with the sequence of states provided (states are not hidden) [Supervised Learning]
- Learning from a set of observations without any state information. [Unsupervised Learning]


## HMM as Parser



$$
\pi=\begin{array}{|l|l|}
\hline A & 0.25 \\
\hline N & 0.75 \\
\hline
\end{array}
$$

$a=$| $a_{i, j}$ | $A$ | $N$ |
| :--- | :--- | :--- |
| $A$ | 0.0 | 1.0 |
| $N$ | 0.5 | 0.5 |


$b=$| $b_{i}(o)$ | clown | killer | problem | crazy |
| :--- | :--- | :--- | :--- | :--- |
| $A$ | 0 | 0 | 0 | 1 |
| $N$ | 0.4 | 0.3 | 0.3 | 0 |

The task: for a given observation sequence find the most likely state sequence. $a_{i, j}=p(j \mid i)$ and $b_{i}(o)=p(o \mid i)$

## HMM as Parser



- Find most likely sequence of states for killer clown
- Score every possible sequence of states: AA, AN, NN, NA
- $\mathrm{P}($ killer clown, AA$)=\pi_{A} \cdot b_{A}($ killer $) \cdot a_{A, A} \cdot b_{A}($ clown $)=0.0$
- $\mathrm{P}($ killer clown, AN$)=\pi_{A} \cdot b_{A}($ killer $) \cdot a_{A, N} \cdot b_{N}($ clown $)=0.0$
- $\mathrm{P}($ killer clown, NN$)=\pi_{N} \cdot b_{N}($ killer $) \cdot a_{N, N} \cdot b_{N}($ clown $)=$ $0.75 \cdot 0.3 \cdot 0.5 \cdot 0.4=0.045$
- $\mathrm{P}($ killer clown, NA $)=\pi_{N} \cdot b_{N}($ killer $) \cdot a_{N, A} \cdot b_{A}($ clown $)=0.0$
- Pick the state sequence with highest probability $(\mathrm{NN}=0.045)$.


## HMM as Parser

- As we have seen, for input of length 2, and a HMM with 2 states there are $2^{2}$ possible state sequences.
- In general, if we have $q$ states and input of length $T$ there are $q^{T}$ possible state sequences.
- Using our example HMM, for input killer crazy clown problem we will have $2^{4}$ possible state sequences to score.
- Our naive algorithm takes exponential time to find the best state sequence for a given input.
- The Viterbi algorithm uses dynamic programming to provide the best state sequence with a time complexity of $q^{2} \cdot T$


# Natural Language Processing 

Anoop Sarkar<br>anoopsarkar.github.io/nlp-class

Simon Fraser University

Part 3: Viterbi Algorithm for HMMs

## Viterbi Algorithm for HMMs

- For input of length $T: o_{1}, \ldots, o_{T}$, we want to find the sequence of states $s_{1}, \ldots, s_{T}$
- Each $s_{t}$ in this sequence is one of the states in the HMM.
- So the task is to find the most likely sequence of states:

$$
\arg \max _{s_{1}, \ldots, s_{T}} P\left(o_{1}, \ldots, o_{T}, s_{1}, \ldots, s_{T}\right)
$$

- The Viterbi algorithm solves this by creating a table $V[s, t]$ where $s$ is one of the states, and $t$ is an index between $1, \ldots, T$.


## Viterbi Algorithm for HMMs



- Consider the input killer crazy clown problem
- So the task is to find the most likely sequence of states:

$$
\arg \max _{s_{1}, s_{2}, s_{3}, s_{4}} P\left(\text { killer crazy clown problem, } s_{1}, s_{2}, s_{3}, s_{4}\right)
$$

- A sub-problem is to find the most likely sequence of states for killer crazy clown:

$$
\arg \max _{s_{1}, s_{2}, s_{3}} P\left(\text { killer crazy clown, } s_{1}, s_{2}, s_{3}\right)
$$

## Viterbi Algorithm for HMMs

- In our example there are two possible values for $s_{4}$ :
$\max _{s_{1}, \ldots, s_{4}} P\left(\right.$ killer crazy clown problem, $\left.s_{1}, s_{2}, s_{3}, s_{4}\right)=$

$$
\begin{aligned}
& \max \left\{\max _{s_{1}, s_{2}, s_{3}} P\left(\text { killer crazy clown problem, } s_{1}, s_{2}, s_{3}, N\right),\right. \\
& \left.\max _{s_{1}, s_{2}, s_{3}} P\left(\text { killer crazy clown problem, } s_{1}, s_{2}, s_{3}, A\right)\right\}
\end{aligned}
$$

- Similarly:

$$
\begin{aligned}
& \max _{s_{1}, \ldots, s_{3}} P\left(\text { killer crazy clown, } s_{1}, s_{2}, s_{3}\right)= \\
& \quad \max \left\{\max _{s_{1}, s_{2}} P\left(\text { killer crazy clown, } s_{1}, s_{2}, N\right),\right. \\
& \left.\max _{s_{1}, s_{2}} P\left(\text { killer crazy clown, } s_{1}, s_{2}, A\right)\right\}
\end{aligned}
$$

## Viterbi Algorithm for HMMs

- Putting them together:
$P\left(\right.$ killer crazy clown problem, $\left.s_{1}, s_{2}, s_{3}, N\right)=$ $\max \left\{P\left(\right.\right.$ killer crazy clown, $\left.s_{1}, s_{2}, N\right) \cdot a_{N, N} \cdot b_{N}($ problem $)$, $P\left(\right.$ killer crazy clown, $\left.s_{1}, s_{2}, A\right) \cdot a_{A, N} \cdot b_{N}($ problem $\left.)\right\}$
$P\left(\right.$ killer crazy clown problem, $\left.s_{1}, s_{2}, s_{3}, A\right)=$ $\max \left\{P\left(\right.\right.$ killer crazy clown, $\left.s_{1}, s_{2}, N\right) \cdot a_{N, A} \cdot b_{A}$ (problem), $P\left(\right.$ killer crazy clown, $\left.s_{1}, s_{2}, A\right) \cdot a_{A, A} \cdot b_{A}($ problem $\left.)\right\}$
- The best score is given by:
$\max _{s_{1}, \ldots, s_{4}} P\left(\right.$ killer crazy clown problem, $\left.s_{1}, s_{2}, s_{3}, s_{4}\right)=$ $\max _{N, A}\left\{\max _{s_{1}, s_{2}, s_{3}} P\left(\right.\right.$ killer crazy clown problem, $\left.s_{1}, s_{2}, s_{3}, N\right)$, $\max _{s_{1}, s_{2}, s_{3}} P\left(\right.$ killer crazy clown problem, $\left.\left.s_{1}, s_{2}, s_{3}, A\right)\right\}$


## Viterbi Algorithm for HMMs

- Provide an index for each input symbol:

1:killer 2:crazy 3:clown 4:problem
$V[N, 3]=\max _{s_{1}, s_{2}} P\left(\right.$ killer crazy clown, $\left.s_{1}, s_{2}, N\right)$
$V[N, 4]=\max _{s_{1}, s_{2}, s_{3}} P\left(\right.$ killer crazy clown problem, $\left.s_{1}, s_{2}, s_{3}, N\right)$

- Putting them together:

$$
\begin{aligned}
V[N, 4]= & \max \left\{V[N, 3] \cdot a_{N, N} \cdot b_{N}(\text { problem }),\right. \\
& \left.V[A, 3] \cdot a_{A, N} \cdot b_{N}(\text { problem })\right\} \\
V[A, 4]= & \max \left\{V[N, 3] \cdot a_{N, A} \cdot b_{A}(\text { problem }),\right. \\
& \left.V[A, 3] \cdot a_{A, A} \cdot b_{A}(\text { problem })\right\}
\end{aligned}
$$

- The best score for the input is given by: $\max \{V[N, 4], V[A, 4]\}$
- To extract the best sequence of states we backtrack (same trick as obtaining alignments from minimum edit distance)


## Viterbi Algorithm for HMMs

- For input of length $T: o_{1}, \ldots, o_{T}$, we want to find the sequence of states $s_{1}, \ldots, s_{T}$
- Each $s_{t}$ in this sequence is one of the states in the HMM.
- For each state $q$ we initialize our table: $V[q, 1]=\pi_{q} \cdot b_{q}\left(o_{1}\right)$
- Then compute for $t=1 \ldots T-1$ for each state $q$ :

$$
V[q, t+1]=\max _{q^{\prime}}\left\{V\left[q^{\prime}, t\right] \cdot a_{q^{\prime}, q} \cdot b_{q}\left(o_{t+1}\right)\right\}
$$

- After the loop terminates, the best score is $\max _{q} V[q, T]$


## Learning from Fully Observed Data

$$
\pi=\begin{array}{|l|l|}
\hline A & 0.25 \\
\hline N & 0.75 \\
\hline
\end{array} \quad a=\begin{array}{|l|l|l|}
\hline a_{i, j} & A & N \\
\hline A & 0.0 & 1.0 \\
\hline N & 0.5 & 0.5 \\
\hline
\end{array}
$$

$b=$| $b_{i}(o)$ | clown | killer | problem | crazy |
| :--- | :--- | :--- | :--- | :--- |
| $A$ | 0 | 0 | 0 | 1 |
| $N$ | 0.4 | 0.3 | 0.3 | 0 |

Viterbi algorithm:

| V | killer:1 | crazy:2 | clown:3 | problem:4 |
| :--- | :--- | :--- | :--- | :--- |
| A |  |  |  |  |
| N |  |  |  |  |

## Learning from Fully Observed Data

$$
\pi=\begin{array}{|l|l|}
\hline A & 0.25 \\
\hline N & 0.75 \\
\hline
\end{array}
$$

$$
a=\begin{array}{|l|l|l|}
\hline a_{i, j} & A & N \\
\hline A & 0.0 & 1.0 \\
\hline N & 0.5 & 0.5 \\
\hline
\end{array}
$$

$b=$| $b_{i}(o)$ | clown | killer | problem | crazy |
| :--- | :--- | :--- | :--- | :--- |
| $A$ | 0 | 0 | 0 | 1 |
| $N$ | 0.4 | 0.3 | 0.3 | 0 |

Viterbi algorithm:

| V | killer:1 | crazy:2 | clown:3 | problem:4 |
| :--- | :--- | :--- | :--- | :--- |
| A | 0 | 0.1125 | 0 | 0 |
| N | 0.225 | 0 | 0.045 | 0.00675 |

## Probability models of language

## Question

$$
\begin{aligned}
& \pi=\begin{array}{|c|c|}
\hline V & 0.25 \\
\hline N & 0.75 \\
\hline
\end{array} \\
& a=\begin{array}{|l|l|l|}
\hline a_{i, j} & V & N \\
\hline V & 0.5 & 0.5 \\
\hline N & 0.5 & 0.5 \\
\hline
\end{array}
\end{aligned}
$$

What is the best sequence of tags for each string below:

1. time
2. time flies
3. time flies can

# Natural Language Processing 

Anoop Sarkar<br>anoopsarkar.github.io/nlp-class

Simon Fraser University

Part 4: HMM as a Language Model

## Hidden Markov Model

Model $\theta= \begin{cases}\pi_{i} & p(i): \text { starting at state } i \\ a_{i, j} & p(j \mid i): \text { transition to state } i \text { from state } j \\ b_{i}(o) & p(o \mid i): \text { output o at state } i\end{cases}$


## Hidden Markov Model Algorithms

- HMM as parser: compute the best sequence of states for a given observation sequence.
- HMM as language model: compute probability of given observation sequence.
- HMM as learner: given a corpus of observation sequences, learn its distribution, i.e. learn the parameters of the HMM from the corpus.
- Learning from a set of observations with the sequence of states provided (states are not hidden) [Supervised Learning]
- Learning from a set of observations without any state information. [Unsupervised Learning]


## HMM as a Language Model



- Find $P($ killer clown $)=\sum_{y} P(y$, killer clown $)$
- $P($ killer clown $)=P(A A$, killer clown $)+P(A N$, killer clown $)+$ $P(N N$, killer clown $)+P(N A$, killer clown $)$


## HMM as a Language Model



- Consider the input killer crazy clown problem
- So the task is to find the sum over all sequences of states:

$$
\sum_{s_{1}, s_{2}, s_{3}, s_{4}} P\left(\text { killer crazy clown problem, } s_{1}, s_{2}, s_{3}, s_{4}\right)
$$

- A sub-problem is to find the most likely sequence of states for killer crazy clown:

$$
\sum_{s_{1}, s_{2}, s_{3}} P\left(\text { killer crazy clown, } s_{1}, s_{2}, s_{3}\right)
$$

## HMM as a Language Model

- In our example there are two possible values for $s_{4}$ :

$$
\begin{aligned}
& \sum_{s_{1}, \ldots, s_{4}} P\left(\text { killer crazy clown problem, } s_{1}, s_{2}, s_{3}, s_{4}\right)= \\
& \sum_{s_{1}, s_{2}, s_{3}} P\left(\text { killer crazy clown problem, } s_{1}, s_{2}, s_{3}, N\right)+
\end{aligned}
$$

$$
\sum_{s_{1}, s_{2}, s_{3}} P\left(\text { killer crazy clown problem, } s_{1}, s_{2}, s_{3}, A\right)
$$

- Very similar to the Viterbi algorithm. Sum instead of max, and that's the only difference!


## HMM as a Language Model

- Provide an index for each input symbol:

1:killer 2:crazy 3:clown 4:problem
$V[N, 3]=\sum_{s_{1}, s_{2}} P\left(\right.$ killer crazy clown, $\left.s_{1}, s_{2}, N\right)$
$V[N, 4]=\sum_{s_{1}, s_{2}, s_{3}} P\left(\right.$ killer crazy clown problem, $\left.s_{1}, s_{2}, s_{3}, N\right)$

- Putting them together:

$$
\begin{aligned}
V[N, 4]= & V[N, 3] \cdot a_{N, N} \cdot b_{N}(\text { problem })+ \\
& V[A, 3] \cdot a_{A, N} \cdot b_{N}(\text { problem }) \\
V[A, 4]= & V[N, 3] \cdot a_{N, A} \cdot b_{A}(\text { problem })+ \\
& V[A, 3] \cdot a_{A, A} \cdot b_{A}(\text { problem })
\end{aligned}
$$

- The best score for the input is given by: $V[N, 4]+V[A, 4]$


## HMM as a Language Model

- For input of length $T: o_{1}, \ldots, o_{T}$, we want to find

$$
P\left(o_{1}, \ldots, o_{T}\right)=\sum_{y_{1}, \ldots, y_{T}} P\left(y_{1}, \ldots, y_{T}, o_{1}, \ldots, o_{T}\right)
$$

- Each $y_{t}$ in this sequence is one of the states in the HMM.
- For each state $q$ we initialize our table: $V[q, 1]=\pi_{q} \cdot b_{q}\left(o_{1}\right)$
- Then compute recursively for $t=1 \ldots T-1$ for each state $q$ :

$$
V[q, t+1]=\sum_{q^{\prime}}\left\{V\left[q^{\prime}, t\right] \cdot a_{q^{\prime}, q} \cdot b_{q}\left(o_{t+1}\right)\right\}
$$

- After the loop terminates, the best score is $\sum_{q} V[q, T]$
- So: Viterbi with sum instead of max gives us an algorithm for HMM as a language model.
- This algorithm is sometimes called the forward algorithm.


# Natural Language Processing 

Anoop Sarkar<br>anoopsarkar.github.io/nlp-class

Simon Fraser University

Part 5: Supervised Learning for HMMs

## Hidden Markov Model Algorithms

- HMM as parser: compute the best sequence of states for a given observation sequence.
- HMM as language model: compute probability of given observation sequence.
- HMM as learner: given a corpus of observation sequences, learn its distribution, i.e. learn the parameters of the HMM from the corpus.
- Learning from a set of observations with the sequence of states provided (states are not hidden) [Supervised Learning]
- Learning from a set of observations without any state information. [Unsupervised Learning]


## Hidden Markov Model

Model $\theta= \begin{cases}\pi_{i} & \text { probability of starting at state } i \\ a_{i, j} & \text { probability of transition from state } i \text { to state } j \\ b_{i}(o) & \text { probability of output } o \text { at state } i\end{cases}$
Constraints : $\sum_{i} \pi_{i}=1, \sum_{j} a_{i, j}=1, \sum_{o} b_{i}(o)=1$


## HMM Learning from Labeled Data



- The task: to find the values for the parameters of the HMM:
- $\pi_{A}, \pi_{N}$
$-a_{A, A}, a_{A, N}, a_{N, N}, a_{N, A}$
- $b_{A}$ (killer), $b_{A}$ (crazy), $b_{A}$ (clown), $b_{A}$ (problem)
- $b_{N}($ killer $), b_{N}($ crazy $), b_{N}$ (clown), $b_{N}$ (problem)


## Learning from Fully Observed Data

## Labeled Data L

```
x1,y1: killer/N clown/N (x1 = killer,clown; y1 = N,N)
x2,y2: killer/N problem/N (x2 = killer,problem; y2 = N,N
x3,y3: crazy/A problem/N
x4,y4: crazy/A clown/N
x5,y5: problem/N crazy/A clown/N
x6,y6: clown/N crazy/A killer/N
```


## Learning from Fully Observed Data

- Let's say we have $m$ labeled examples:
$L=\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)$
- Each $\left(x_{\ell}, y_{\ell}\right)=\left\{o_{1}, \ldots, o_{T}, s_{1}, \ldots, s_{T}\right\}$
- For each $\left(x_{\ell}, y_{\ell}\right)$ we can compute the probability using the HMM:
- $\left(x_{1}=\right.$ killer, clown; $\left.y_{1}=N, N\right)$ :
$P\left(x_{1}, y_{1}\right)=\pi_{N} \cdot b_{N}($ killer $) \cdot a_{N, N} \cdot b_{N}($ clown $)$
- $\left(x_{2}=\right.$ killer, problem; $\left.y_{2}=N, N\right)$ :
$P\left(x_{2}, y_{2}\right)=\pi_{N} \cdot b_{N}($ killer $) \cdot a_{N, N} \cdot b_{N}($ problem $)$
- $\left(x_{3}=\right.$ crazy, problem; $\left.y_{3}=A, N\right)$ :
$P\left(x_{3}, y_{3}\right)=\pi_{A} \cdot b_{A}($ crazy $) \cdot a_{A, N} \cdot b_{N}($ problem $)$
- $\left(x_{4}=\right.$ crazy, clown; $\left.y_{4}=A, N\right)$ :
$P\left(x_{4}, y_{4}\right)=\pi_{A} \cdot b_{A}($ crazy $) \cdot a_{A, N} \cdot b_{N}($ clown $)$
- $\left(x_{5}=\right.$ problem, crazy, clown; $\left.y_{5}=N, A, N\right)$ :
$P\left(x_{5}, y_{5}\right)=\pi_{N} \cdot b_{N}($ problem $) \cdot a_{N, A} \cdot b_{A}($ crazy $) \cdot a_{A, N} \cdot b_{N}($ clown $)$
- $\left(x_{6}=\right.$ clown, crazy, killer; $\left.y_{6}=N, A, N\right)$ :
$P\left(x_{6}, y_{6}\right)=\pi_{N} \cdot b_{N}($ clown $) \cdot a_{N, A} \cdot b_{A}($ crazy $) \cdot a_{A, N} \cdot b_{N}($ killer $)$
- $\prod_{\ell} P\left(x_{\ell}, y_{\ell}\right)=\pi_{N}{ }^{4} \cdot \pi_{A}{ }^{2} \cdot a_{N, N}{ }^{2} \cdot a_{N, A}{ }^{2} \cdot a_{A, N}{ }^{4} \cdot a_{A, A}{ }^{0} \cdot b_{N}(\text { killer })^{3}$. $b_{N}(\text { clown })^{4} \cdot b_{N}(\text { problem })^{3} \cdot b_{A}(\text { crazy })^{4}$


## Learning from Fully Observed Data

- We can easily collect frequency of observing a word with a state (tag)
- $f(i, x, y)=$ number of times $i$ is the initial state in $(x, y)$
- $f(i, j, x, y)=$ number of times $j$ follows $i$ in $(x, y)$
- $f(i, o, x, y)=$ number of times $i$ is paired with observation o
- Then according to our HMM the probability of $x, y$ is:

$$
P(x, y)=\prod_{i} \pi_{i}^{f(i, x, y)} \cdot \prod_{i, j} a_{i, j}^{f(i, j, x, y)} \cdot \prod_{i, o} b_{i}(o)^{f(i, o, x, y)}
$$

## Learning from Fully Observed Data

- According to our HMM the probability of $x, y$ is:

$$
P(x, y)=\prod_{i} \pi_{i}^{f(i, x, y)} \cdot \prod_{i, j} a_{i, j}^{f(i, j, x, y)} \cdot \prod_{i, o} b_{i}(o)^{f(i, o, x, y)}
$$

- For the labeled data $L=\left(x_{1}, y_{1}\right), \ldots,\left(x_{\ell}, y_{\ell}\right), \ldots,\left(x_{m}, y_{m}\right)$

$$
\begin{aligned}
P(L) & =\prod_{\ell=1}^{m} P\left(x_{\ell}, y_{\ell}\right) \\
& =\prod_{\ell=1}^{m}\left(\prod_{i} \pi_{i}^{f\left(i, x_{\ell}, y_{\ell}\right)} \cdot \prod_{i, j} a_{i, j}^{f\left(i, j, x_{\ell}, y_{\ell}\right)} \cdot \prod_{i, o} b_{i}(o)^{f\left(i, o, x_{\ell}, y_{\ell}\right)}\right)
\end{aligned}
$$

## Learning from Fully Observed Data

- According to our HMM the probability of $x, y$ is:

$$
P(L)=\prod_{\ell=1}^{m}\left(\prod_{i} \pi_{i}^{f\left(i, x_{\ell}, y_{\ell}\right)} \cdot \prod_{i, j} a_{i, j}^{f\left(i, j, x_{\ell}, y_{\ell}\right)} \cdot \prod_{i, o} b_{i}(o)^{f\left(i, o, x_{\ell}, y_{\ell}\right)}\right)
$$

- The log probability of the labeled data $\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)$ according to HMM with parameters $\theta$ is:

$$
\begin{aligned}
L(\theta)= & \sum_{\ell=1}^{m} \log P\left(x_{\ell}, y_{\ell}\right) \\
= & \sum_{\ell=1}^{m} \sum_{i} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}+ \\
& \sum_{i, j} f\left(i, j, x_{\ell}, y_{\ell}\right) \log a_{i, j}+ \\
& \sum_{i, o} f\left(i, o, x_{\ell}, y_{\ell}\right) \log b_{i}(o)
\end{aligned}
$$

## Learning from Fully Observed Data

$$
\begin{aligned}
& L(\theta)=\sum_{\ell=1}^{m} \\
& \quad \sum_{i} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}+\sum_{i, j} f\left(i, j, x_{\ell}, y_{\ell}\right) \log a_{i, j}+\sum_{i, o} f\left(i, o, x_{\ell}, y_{\ell}\right) \log b_{i}(o)
\end{aligned}
$$

- $\theta=(\pi, a, b)$
- $L(\theta)$ is the log probability of the labeled data $\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)$
- We want to find a $\theta$ that will give us the maximum value of $L(\theta)$
- Find the $\theta$ such that $\frac{d L(\theta)}{d \theta}=0$


## Learning from Fully Observed Data

$$
\begin{aligned}
L(\theta) & =\sum_{\ell=1}^{m} \\
& \sum_{i} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}+\sum_{i, j} f\left(i, j, x_{\ell}, y_{\ell}\right) \log a_{i, j}+\sum_{i, o} f\left(i, o, x_{\ell}, y_{\ell}\right) \log b_{i}(o)
\end{aligned}
$$

- The values of $\pi_{i}, a_{i, j}, b_{i}(o)$ that maximize $L(\theta)$ are:

$$
\begin{aligned}
\pi_{i} & =\frac{\sum_{\ell} f\left(i, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{k} f\left(k, x_{\ell}, y_{\ell}\right)} \\
a_{i, j} & =\frac{\sum_{\ell} f\left(i, j, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{k} f\left(i, k, x_{\ell}, y_{\ell}\right)} \\
b_{i}(o) & =\frac{\sum_{\ell} f\left(i, o, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{o^{\prime} \in V} f\left(i, o^{\prime}, x_{\ell}, y_{\ell}\right)}
\end{aligned}
$$

## Learning from Fully Observed Data

## Labeled Data:

```
x1,y1: killer/N clown/N
x2,y2: killer/N problem/N
x3,y3: crazy/A problem/N
x4,y4: crazy/A clown/N
x5,y5: problem/N crazy/A clown/N
x6,y6: clown/N crazy/A killer/N
```


## Learning from Fully Observed Data

- The values of $\pi_{i}$ that maximize $L(\theta)$ are:

$$
\pi_{i}=\frac{\sum_{\ell} f\left(i, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{k} f\left(k, x_{\ell}, y_{\ell}\right)}
$$

- $\pi_{N}=\frac{2}{3}$ and $\pi_{A}=\frac{1}{3}$ because:

$$
\begin{aligned}
& \sum_{\ell} f\left(N, x_{\ell}, y_{\ell}\right)=4 \\
& \sum_{\ell} f\left(A, x_{\ell}, y_{\ell}\right)=2
\end{aligned}
$$

## Learning from Fully Observed Data

- The values of $a_{i, j}$ that maximize $L(\theta)$ are:

$$
a_{i, j}=\frac{\sum_{\ell} f\left(i, j, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{k} f\left(i, k, x_{\ell}, y_{\ell}\right)}
$$

$-a_{N, N}=\frac{1}{2} ; a_{N, A}=\frac{1}{2} ; a_{A, N}=1$ and $a_{A, A}=0$ because:

$$
\begin{array}{ll}
\sum_{\ell} f\left(N, N, x_{\ell}, y_{\ell}\right)=2 & \sum_{\ell} f\left(A, N, x_{\ell}, y_{\ell}\right)=4 \\
\sum_{\ell} f\left(N, A, x_{\ell}, y_{\ell}\right)=2 & \sum_{\ell} f\left(A, A, x_{\ell}, y_{\ell}\right)=0
\end{array}
$$

## Learning from Fully Observed Data

- The values of $b_{i}(o)$ that maximize $L(\theta)$ are:

$$
b_{i}(o)=\frac{\sum_{\ell} f\left(i, o, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{o^{\prime} \in V} f\left(i, o^{\prime}, x_{\ell}, y_{\ell}\right)}
$$

- $b_{N}($ killer $)=\frac{3}{10} ; b_{N}($ clown $)=\frac{4}{10} ; b_{N}($ problem $)=\frac{3}{10}$ and $b_{A}($ crazy $)=1$ because:

$$
\begin{aligned}
\sum_{\ell} f\left(N, \text { killer, } x_{\ell}, y_{\ell}\right) & =3 & \sum_{\ell} f\left(A, \text { killer, } x_{\ell}, y_{\ell}\right) & =0 \\
\sum_{\ell} f\left(N, \text { clown }, x_{\ell}, y_{\ell}\right) & =4 & \sum_{\ell} f\left(A, \text { clown }, x_{\ell}, y_{\ell}\right) & =0 \\
\sum_{\ell} f\left(N, \text { crazy, } x_{\ell}, y_{\ell}\right) & =0 & \sum_{\ell} f\left(A, \text { crazy }, x_{\ell}, y_{\ell}\right) & =4 \\
\sum_{\ell} f\left(N, \text { problem }, x_{\ell}, y_{\ell}\right) & =3 & \sum_{\ell} f\left(A, \text { problem, } x_{\ell}, y_{\ell}\right) & =0
\end{aligned}
$$

## Learning from Fully Observed Data

```
x1,y1: killer/N clown/N
x2,y2: killer/N problem/N
x3,y3: crazy/A problem/N
x4,y4: crazy/A clown/N
x5,y5: problem/N crazy/A clown/N
x6,y6: clown/N crazy/A killer/N
```

$$
\pi=\begin{array}{|l|l|}
\hline A & 0.25 \\
\hline N & 0.75 \\
\hline
\end{array}
$$

$a=$| $a_{i, j}$ | $A$ | $N$ |
| :--- | :--- | :--- |
| $A$ | 0.0 | 1.0 |
| $N$ | 0.5 | 0.5 |


$b=$| $b_{i}(o)$ | clown | killer | problem | crazy |
| :--- | :--- | :--- | :--- | :--- |
| $A$ | 0 | 0 | 0 | 1 |
| $N$ | 0.4 | 0.3 | 0.3 | 0 |

# Natural Language Processing 

Anoop Sarkar<br>anoopsarkar.github.io/nlp-class<br>Simon Fraser University

Part 6: Lagrange Multipliers

## Hidden Markov Model

Model $\theta= \begin{cases}\pi_{i} & p(i): \text { starting at state } i \\ a_{i, j} & p(j \mid i): \text { transition to state } i \text { from state } j \\ b_{i}(o) & p(o \mid i): \text { output o at state } i\end{cases}$


## Learning from Fully Observed Data

$$
\begin{aligned}
& L(\theta)=\sum_{\ell=1}^{m} \\
& \quad \sum_{i} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}+\sum_{i, j} f\left(i, j, x_{\ell}, y_{\ell}\right) \log a_{i, j}+\sum_{i, o} f\left(i, o, x_{\ell}, y_{\ell}\right) \log b_{i}(o)
\end{aligned}
$$

- $\theta=(\pi, a, b)$
- $L(\theta)$ is the log probability of the labeled data $\left(x_{1}, y_{1}\right), \ldots,\left(x_{m}, y_{m}\right)$
- We want to find a $\theta$ that will give us the maximum value of $L(\theta)$
- Find the $\theta$ such that $\frac{d L(\theta)}{d \theta}=0$


## Learning from Fully Observed Data

$$
\begin{aligned}
& L(\theta)=\sum_{\ell=1}^{m} \\
& \quad \sum_{i} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}+\sum_{i, j} f\left(i, j, x_{\ell}, y_{\ell}\right) \log a_{i, j}+\sum_{i, o} f\left(i, o, x_{\ell}, y_{\ell}\right) \log b_{i}(o)
\end{aligned}
$$

- Find the $\theta$ such that $\frac{d L(\theta)}{d \theta}=0$ and $\theta=(\pi, a, b)$
- Split up $L(\theta)$ into $L(\pi), L(a), L(b)$
- Let $\nabla L=\forall i, j, o: \frac{\partial L(\pi)}{\partial \pi_{i}}, \frac{\partial L(a)}{\partial a_{i, j}}, \frac{\partial L(b)}{\partial b_{i}(o)}$
- We must also obey constraints:

$$
\sum_{k} \pi_{k}=1, \sum_{k} a_{i, k}=1, \sum_{o} b_{i}(o)=1
$$

## Learning from Fully Observed Data

$$
L(\pi)=\sum_{\ell=1}^{m} \sum_{i} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}
$$

- Let us focus on $\nabla L(\pi)$ (the other two: $a$ and $b$ are similar)
- For the constraint $\sum_{k} \pi_{k}=1$ we introduce a new variable into our search for a maximum:

$$
L(\pi, \lambda)=L(\pi)+\lambda\left(1-\sum_{k} \pi_{k}\right)
$$

- $\lambda$ is called the Lagrange multiplier
- $\lambda$ penalizes any solution that does not obey the constraint
- The constraint ensures that $\pi$ is a probability distribution


## Learning from Fully Observed Data

$$
\frac{\partial L(\pi)}{\partial \pi_{i}}=\frac{\partial}{\partial \pi_{i}} \underbrace{\sum_{\ell=1}^{m} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}}_{\text {the only part with variable } \pi_{i}}+\underbrace{\sum_{\ell=1}^{m} \sum_{j: j \neq i} f\left(j, x_{\ell}, y_{\ell}\right) \log \pi_{j}}_{\text {no } \pi_{i} \text { so derivative is } 0}
$$

- We want a value of $\pi_{i}$ such that $\frac{\partial L(\pi, \lambda)}{\partial \pi_{i}}=0$

$$
\begin{aligned}
\frac{\partial}{\partial \pi_{i}} \sum_{\ell=1}^{m}\left(f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}+\lambda\left(1-\sum_{k} \pi_{k}\right)\right) & =0 \\
\frac{\partial}{\partial \pi_{i}} \sum_{\ell=1}^{m}(\underbrace{f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}}_{\frac{\partial}{\partial \pi_{i}}=\frac{f\left(i, \chi_{\ell}, y_{\ell}\right)}{\pi_{i}}}+\lambda-\underbrace{\lambda \pi_{i}}_{\frac{\partial}{\partial \pi_{i}}=\lambda}-\lambda \sum_{j: j \neq i} \pi_{j}) & =0
\end{aligned}
$$

## Learning from Fully Observed Data

$$
\frac{\partial L(\pi)}{\partial \pi_{i}}=\frac{\partial}{\partial \pi_{i}} \underbrace{\sum_{\ell=1}^{m} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}}_{\text {the only part with variable } \pi_{i}}+\underbrace{\sum_{\ell=1}^{m} \sum_{j: j \neq i} f\left(j, x_{\ell}, y_{\ell}\right) \log \pi_{j}}_{\text {no } \pi_{i} \text { so derivative is } 0}
$$

- We can obtain a value of $\pi_{i}$ wrt $\lambda$ :

$$
\begin{array}{r}
\frac{\partial L(\pi, \lambda)}{\partial \pi_{i}}=\underbrace{\sum_{\ell=1}^{m} \frac{f\left(i, x_{\ell}, y_{\ell}\right)}{\pi_{i}}-\lambda}_{\text {see previous slide }}=0 \\
\pi_{i}=\frac{\sum_{\ell=1}^{m} f\left(i, x_{\ell}, y_{\ell}\right)}{\lambda} \tag{1}
\end{array}
$$

- Combine $\pi_{i}$ s from Eqn (1) with constraint $\sum_{k} \pi_{k}=1$

$$
\lambda=\sum_{k} \sum_{\ell=1}^{m} f\left(k, x_{\ell}, y_{\ell}\right)
$$

## Learning from Fully Observed Data

$$
\frac{\partial L(\pi)}{\partial \pi_{i}}=\frac{\partial}{\partial \pi_{i}} \underbrace{\sum_{\ell=1}^{m} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}}_{\text {the only part with variable } \pi_{i}}+\underbrace{\sum_{\ell=1}^{m} \sum_{j: j \neq i} f\left(j, x_{\ell}, y_{\ell}\right) \log \pi_{j}}_{\text {no } \pi_{i} \text { so derivative is } 0}
$$

- The value of $\pi_{i}$ for which $\frac{\partial L(\pi, \lambda)}{\partial \pi_{i}}=0$ is Eqn (2) which can be combined with the value of $\lambda$ from Eqn (3).

$$
\begin{array}{r}
\pi_{i}=\frac{\sum_{\ell=1}^{m} f\left(i, x_{\ell}, y_{\ell}\right)}{\lambda} \\
\lambda=\sum_{k} \sum_{\ell=1}^{m} f\left(k, x_{\ell}, y_{\ell}\right)  \tag{3}\\
\pi_{i}=\frac{\sum_{\ell=1}^{m} f\left(i, x_{\ell}, y_{\ell}\right)}{\sum_{k} \sum_{\ell=1}^{m} f\left(k, x_{\ell}, y_{\ell}\right)}
\end{array}
$$

## Learning from Fully Observed Data

$$
\begin{aligned}
L(\theta) & =\sum_{\ell=1}^{m} \\
& \sum_{i} f\left(i, x_{\ell}, y_{\ell}\right) \log \pi_{i}+\sum_{i, j} f\left(i, j, x_{\ell}, y_{\ell}\right) \log a_{i, j}+\sum_{i, o} f\left(i, o, x_{\ell}, y_{\ell}\right) \log b_{i}(o)
\end{aligned}
$$

- The values of $\pi_{i}, a_{i, j}, b_{i}(o)$ that maximize $L(\theta)$ are:

$$
\begin{aligned}
\pi_{i} & =\frac{\sum_{\ell} f\left(i, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{k} f\left(k, x_{\ell}, y_{\ell}\right)} \\
a_{i, j} & =\frac{\sum_{\ell} f\left(i, j, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{k} f\left(i, k, x_{\ell}, y_{\ell}\right)} \\
b_{i}(o) & =\frac{\sum_{\ell} f\left(i, o, x_{\ell}, y_{\ell}\right)}{\sum_{\ell} \sum_{o^{\prime} \in V} f\left(i, o^{\prime}, x_{\ell}, y_{\ell}\right)}
\end{aligned}
$$

# Natural Language Processing 

Anoop Sarkar<br>anoopsarkar.github.io/nlp-class

Simon Fraser University

Part 7: Unsupervised Learning for HMMs

## Hidden Markov Model

Model $\theta= \begin{cases}\pi_{i} & p(i): \text { starting at state } i \\ a_{i, j} & p(j \mid i): \text { transition to state } i \text { from state } j \\ b_{i}(o) & p(o \mid i): \text { output o at state } i\end{cases}$


## Hidden Markov Model Algorithms

- HMM as parser: compute the best sequence of states for a given observation sequence.
- HMM as language model: compute probability of given observation sequence.
- HMM as learner: given a corpus of observation sequences, learn its distribution, i.e. learn the parameters of the HMM from the corpus.
- Learning from a set of observations with the sequence of states provided (states are not hidden) [Supervised Learning]
- Learning from a set of observations without any state information. [Unsupervised Learning]


## Learning from Unlabeled Data

## Unlabeled Data $U=x_{1}, \ldots, x_{m}$ :

x1: killer clown
x2: killer problem
x3: crazy problem
x4: crazy clown

- y1, y2, y3, y4 are unknown.
- But we can enumerate all possible values for y1, y2, y3, y4
- For example, for x1: killer clown

$$
\begin{array}{lll}
\mathrm{x} 1, \mathrm{y} 1,1: & \mathrm{killer} / \mathrm{A} \text { clown } / \mathrm{A} & p_{1}=\pi_{A} \cdot b_{A}(\text { killer }) \cdot a_{A, A} \cdot b_{A}(\text { clown }) \\
\mathrm{x} 1, \mathrm{y} 1,2: & \mathrm{killer} / \mathrm{A} \text { clown } / \mathrm{N} & p_{2}=\pi_{A} \cdot b_{A}(\text { killer }) \cdot a_{A, N} \cdot b_{N}(\text { clown }) \\
\mathrm{x} 1, \mathrm{y} 1,3: & \mathrm{killer} / \mathrm{N} \text { clown } / \mathrm{N} & p_{3}=\pi_{N} \cdot b_{N}(\text { killer }) \cdot a_{N, N} \cdot b_{N}(\text { clown }) \\
\mathrm{x} 1, \mathrm{y} 1,4: & \mathrm{killer} / \mathrm{N} \text { clown/A } & p_{4}=\pi_{N} \cdot b_{N}(\text { killer }) \cdot a_{N, A} \cdot b_{A}(\text { clown })
\end{array}
$$

## Learning from Unlabeled Data

- Assume some values for $\theta=\pi, a, b$
- We can compute $P\left(y \mid x_{\ell}, \theta\right)$ for any $y$ for a given $x_{\ell}$

$$
P\left(y \mid x_{\ell}, \theta\right)=\frac{P(x, y \mid \theta)}{\sum_{y^{\prime}} P\left(x, y^{\prime} \mid \theta\right)}
$$

- For example, we can compute $P(\mathrm{NN} \mid$ killer clown, $\theta)$ as follows:

$$
\frac{\pi_{N} \cdot b_{N}(\text { killer }) \cdot a_{N, N} \cdot b_{N}(\text { clown })}{\sum_{i, j} \pi_{i} \cdot b_{i}(\text { killer }) \cdot a_{i, j} \cdot b_{j}(\text { clown })}
$$

- $P\left(y \mid x_{\ell}, \theta\right)$ is called the posterior probability


## Learning from Unlabeled Data

- Compute the posterior for all possible outputs for each example in training:
- For x1: killer clown

$$
\begin{array}{lll}
\mathrm{x} 1, \mathrm{y} 1,1: & \text { killer/A clown/A } & P(\mathrm{AA} \mid \text { killer clown, } \theta) \\
\mathrm{x} 1, \mathrm{y} 1,2: & \text { killer/A clown/N } & P(\mathrm{AN} \mid \text { killer clown, } \theta) \\
\mathrm{x} 1, \mathrm{y} 1,3: & \text { killer/N clown/N } & P(\mathrm{NN} \mid \text { killer clown, } \theta) \\
\mathrm{x} 1, \mathrm{y} 1,4: & \text { killer/N clown/A } & P(\mathrm{NA} \mid \text { killer clown, } \theta)
\end{array}
$$

- For x2: killer problem
x2,y2,1: killer/A problem/A $\quad P(A A \mid k i l l e r ~ p r o b l e m, ~ \theta) ~$
x2,y2,2: killer/A problem/N $\quad P($ AN $\mid$ killer problem, $\theta)$
x2,y2,3: killer/N problem/N $\quad P(N N \mid k i l l e r ~ p r o b l e m, ~ \theta) ~$
$\mathrm{x} 2, \mathrm{y} 2,4$ : killer/N problem/A $\quad P(\mathrm{NA} \mid$ killer problem, $\theta)$
- Similarly for x3: crazy problem
- And x4: crazy clown


## Learning from Unlabeled Data

- For unlabeled data, the $\log$ probability of the data given $\theta$ is:

$$
\begin{aligned}
L(\theta) & =\sum_{\ell=1}^{m} \log \sum_{y} P\left(x_{\ell}, y \mid \theta\right) \\
& =\sum_{\ell=1}^{m} \log \sum_{y} P\left(y \mid x_{\ell}, \theta\right) \cdot P\left(x_{\ell} \mid \theta\right)
\end{aligned}
$$

- Unlike the fully observed case there is no simple solution to finding $\theta$ to maximize $L(\theta)$
- We instead initialize $\theta$ to some values, and then iteratively find better values of $\theta: \theta^{0}, \theta^{1}, \ldots$ using the following formula:

$$
\begin{aligned}
\theta^{t} & =\arg \max _{\theta} Q\left(\theta, \theta^{t-1}\right) \\
& =\sum_{\ell=1}^{m} \sum_{y} P\left(y \mid x_{\ell}, \theta^{t-1}\right) \cdot \log P\left(x_{\ell}, y \mid \theta\right)
\end{aligned}
$$

## Learning from Unlabeled Data

$$
\begin{aligned}
\theta^{t}= & \arg \max _{\theta} Q\left(\theta, \theta^{t-1}\right) \\
Q\left(\theta, \theta^{t-1}\right)= & \sum_{\ell=1}^{m} \sum_{y} P\left(y \mid x_{\ell}, \theta^{t-1}\right) \cdot \log P\left(x_{\ell}, y \mid \theta\right) \\
= & \sum_{\ell=1}^{m} \sum_{y} P\left(y \mid x_{\ell}, \theta^{t-1}\right) \\
& \left(\sum_{i} f\left(i, x_{\ell}, y\right) \cdot \log \pi_{i}\right. \\
& +\sum_{i, j} f\left(i, j, x_{\ell}, y\right) \cdot \log a_{i, j} \\
& \left.+\sum_{i, o} f\left(i, o, x_{\ell}, y\right) \cdot \log b_{i}(o)\right)
\end{aligned}
$$

## Learning from Unlabeled Data

$$
\begin{aligned}
g\left(i, x_{\ell}\right)= & \sum_{y} P\left(y \mid x_{\ell}, \theta^{t-1}\right) \cdot f\left(i, x_{\ell}, y\right) \\
g\left(i, j, x_{\ell}\right)= & \sum_{y} P\left(y \mid x_{\ell}, \theta^{t-1}\right) \cdot f\left(i, j, x_{\ell}, y\right) \\
g\left(i, o, x_{\ell}\right)= & \sum_{y} P\left(y \mid x_{\ell}, \theta^{t-1}\right) \cdot f\left(i, o, x_{\ell}, y\right) \\
\theta^{t}= & \arg \max _{\pi, a, b} \sum_{\ell=1}^{m} \sum_{i} g\left(i, x_{\ell}\right) \cdot \log \pi_{i} \\
& +\sum_{i, j} g\left(i, j, x_{\ell}\right) \cdot \log a_{i, j} \\
& +\sum_{i, o} g\left(i, o, x_{\ell}\right) \cdot \log b_{j}(o)
\end{aligned}
$$

## Learning from Unlabeled Data

$$
\begin{aligned}
& Q\left(\theta, \theta^{t-1}\right)=\sum_{\ell=1}^{m} \\
& \quad \sum_{i} g\left(i, x_{\ell}\right) \log \pi_{i}+\sum_{i, j} g\left(i, j, x_{\ell}\right) \log a_{i, j}+\sum_{i, o} g\left(i, o, x_{\ell}\right) \log b_{i}(o)
\end{aligned}
$$

- The values of $\pi_{i}, a_{i, j}, b_{i}(o)$ that maximize $L(\theta)$ are:

$$
\begin{aligned}
\pi_{i} & =\frac{\sum_{\ell} g\left(i, x_{\ell}\right)}{\sum_{\ell} \sum_{k} g\left(k, x_{\ell}\right)} \\
a_{i, j} & =\frac{\sum_{\ell} g\left(i, j, x_{\ell}\right)}{\sum_{\ell} \sum_{k} g\left(i, k, x_{\ell}\right)} \\
b_{i}(o) & =\frac{\sum_{\ell} g\left(i, o, x_{\ell}\right)}{\sum_{\ell} \sum_{o^{\prime} \in V} g\left(i, o^{\prime}, x_{\ell}\right)}
\end{aligned}
$$

## EM Algorithm for Learning HMMs

- Initialize $\theta^{0}$ at random. Let $t=0$.
- The EM Algorithm:
- E-step: compute expected values of $y, P(y \mid x, \theta)$ and calculate $g(i, x), g(i, j, x), g(i, o, x)$
- M-step: compute $\theta^{t}=\arg \max _{\theta} Q\left(\theta, \theta^{t-1}\right)$
- Stop if $L\left(\theta^{t}\right)$ did not change much since last iteration. Else continue.
- The above algorithm is guaranteed to improve likelihood of the unlabeled data.
- In other words, $L\left(\theta^{t}\right) \geq L\left(\theta^{t-1}\right)$
- But it all depends on $\theta^{0}$ !


## Acknowledgements

Many slides borrowed or inspired from lecture notes by Michael Collins, Chris Dyer, Kevin Knight, Philipp Koehn, Adam Lopez, Graham Neubig and Luke Zettlemoyer from their NLP course materials.

All mistakes are my own.
A big thank you to all the students who read through these notes and helped me improve them.

