## Lecture 9: Hidden Markov Model

Instructor: Yen-Chi Chen

These notes are partially based on those of Mathias Drton.

### 9.1 Introduction

Hidden Markov Model (HMM) is a powerful tool to model a time series with multiple patterns. For a concrete example, consider the trace plot of an MCMC for a 2-Gaussian mixture model:


There seems to be two patterns in the plot-one centered at 0 and the other centered at 6 . The trajectory is oscillating between the two centers. The HMM is a powerful tool to model a data like this.

The HMM consists of two variables, the observed variable $Y_{t}$ and a hidden state variable $X_{t}$. The variables $\left\{Y_{1}, \cdots, Y_{n}\right\}$ are what we observed whereas the hidden variables $\left\{X_{1}, \cdots, X_{n}\right\}$ are unobserved states at each time point. In an HMM, the joint distribution of the observations $Y_{1}, \ldots, Y_{n}$ and the hidden states $X_{1}, \ldots, X_{n}$ factors according to the graph:


Namely, the observations $Y_{t}$ 's are dependent on each other because of the hidden variables. This implies that conditioned on the hidden states, observed variables $Y_{t}$ 's are independent. Moreover, the hidden states form a homogeneous Markov chain. Here we will assume that the number of states $S=\{1,2, \cdots, s\}$ is finite and we will also assume that the observed variables are categorical/discrete $Y_{t} \in M=\{1, \cdots, m\}$.

The HMM consists of 3 sets of parameters:

- Initial distribution (of hidden state): $\nu=(\nu(1), \cdots, \nu(s))$.
- Transition probability: $\mathbf{P}=\left\{p_{i j}\right\}, i, j=1, \cdots, s$.
- Emission probability: $\mathbf{E}=\{e(k \mid i)\}, i=1, \cdots, s$ and $k=1, \cdots, m$.

The first two parameters are the typical parameters of a Markov chain and the last one, the emission probability, describes how a hidden state is associated with the observed variables.

Let $\mathbf{y}_{-t}=\left(y_{1}, \cdots, y_{t-1}, y_{t+1}, \cdots, y_{n}\right)$ and $\mathbf{y}_{a: b}=\left(y_{a}, y_{a+1}, \cdots, y_{b}\right)$ and $\mathbf{y}=\left(y_{1}, \cdots, y_{n}\right)$. According to the graphical model, the variables of an HMM is characterizes by the parameters as

$$
\begin{aligned}
P\left(x_{1}, \cdots, x_{n}\right) & =\nu\left(x_{1}\right) \prod_{t=2}^{n} P\left(x_{t} \mid x_{t-1}\right)=\nu\left(x_{1}\right) \prod_{t=2}^{n} p_{x_{t-1}, x_{t}} \\
P\left(y_{t} \mid \mathbf{y}_{-t}, \mathbf{x}\right) & =P\left(y_{t} \mid x_{t}\right)=e\left(y_{t} \mid x_{t}\right)
\end{aligned}
$$

Beware, although the hidden variables $X_{1}, \cdots, X_{n}$ are Markov chain, the observed variables $Y_{1}, \cdots, Y_{n}$ may not be.

There are four common goals of an HMM:

- Likelihood Evaluation: we want to rapidly evaluate the likelihood value $L(\nu, \mathbf{P}, \mathbf{E} \mid \mathbf{y})=P(\mathbf{y})$.
- Parameter Estimation: we want to estimate the underlying parameters $\nu, \mathbf{P}, \mathbf{E}$.
- Hidden State Inference: given $\mathbf{y}$ and parameters $\nu, \mathbf{P}, \mathbf{E}$, we want to reconstruct the corresponding hidden states $\mathbf{x}$.
- Forecasting: we want to predict the future outcomes.


### 9.2 Likelihood Evaluation: Forward and Backward Algorithm

### 9.2.1 Backward algorithm

Assume that we observed $\mathbf{Y}=\mathbf{y}$. The likelihood function is the joint probability of $\mathbf{y}$, so

$$
\begin{aligned}
L(\nu, \mathbf{P}, \mathbf{E} \mid \mathbf{y})=P(\mathbf{y})= & \sum_{\mathbf{x}} P(\mathbf{y}, \mathbf{x}) \\
& =\sum_{\mathbf{x}} P(\mathbf{y} \mid \mathbf{x}) P(\mathbf{x}) \\
& =\sum_{\mathbf{x}} \prod_{t=1}^{n} P\left(y_{t} \mid \mathbf{y}_{1:(t-1)}, \mathbf{x}\right) \nu\left(x_{1}\right) \prod_{r=2}^{n} P\left(x_{r} \mid \mathbf{x}_{1:(r-1)}\right) \\
& =\sum_{\mathbf{x}} \nu\left(x_{1}\right) \prod_{t=1}^{n} e\left(y_{t} \mid x_{t}\right) \prod_{r=2}^{n} p\left(x_{r} \mid x_{r-1}\right) \\
& =\sum_{\mathbf{x}} \nu\left(x_{1}\right) e\left(y_{1} \mid x_{1}\right) \prod_{t=2}^{n} e\left(y_{t} \mid x_{t}\right) p\left(x_{t} \mid x_{t-1}\right) .
\end{aligned}
$$

Note that because each $x_{t}$ has $s$ possible states, computing the summation here requires an evaluation of $s^{n}$ elements, which is very large. Thus, we need to find a smart way to bypass this summation problem.

Here we introduce the backward algorithm, which will drastically reduce the computational cost. To illustrate the idea, we consider the case $n=3$. The likelihood function will be

$$
\begin{aligned}
L(\nu, \mathbf{P}, \mathbf{E} \mid \mathbf{y}) & =\sum_{x_{1}, x_{2}, x_{3}} \nu\left(x_{1}\right) e\left(y_{1} \mid x_{1}\right) e\left(y_{2} \mid x_{2}\right) p\left(x_{2} \mid x_{1}\right) e\left(y_{3} \mid x_{3}\right) p\left(x_{3} \mid x_{2}\right) \\
& =\underbrace{\sum_{x_{1}} \nu\left(x_{1}\right) e\left(y_{1} \mid x_{1}\right) \underbrace{\{\sum_{x_{2}} e\left(y_{2} \mid x_{2}\right) p\left(x_{2} \mid x_{1}\right) \underbrace{\left\{\sum_{x_{3}} e\left(y_{3} \mid x_{3}\right) p\left(x_{3} \mid x_{2}\right)\right\}}_{b_{2}\left(x_{2}\right)}\}}_{b_{1}\left(x_{1}\right)}}_{b_{0}} .
\end{aligned}
$$

What do we gain by doing this? Naively summing over all states requires an $O\left(n s^{n}\right)$ operations. But recursively evaluating $b_{t}(1), \cdots, b_{t}(s)$ is a lot cheaper using dynamic programming; the total cost is $O\left(n s^{2}\right)$. Now we take a close look at each $b_{t}\left(x_{t}\right)$.

$$
b_{2}\left(x_{2}\right)=\sum_{x_{3}} e\left(y_{3} \mid x_{3}\right) p\left(x_{3} \mid x_{2}\right)=\sum_{x_{3}} P\left(y_{3} \mid x_{3}\right) P\left(x_{3} \mid x_{2}\right)=P\left(y_{3} \mid x_{2}\right)
$$

And

$$
\begin{aligned}
b_{1}\left(x_{1}\right) & =\sum_{x_{2}} e\left(y_{2} \mid x_{2}\right) p\left(x_{2} \mid x_{1}\right) b_{2}\left(x_{2}\right) \\
& =\sum_{x_{2}} P\left(y_{2} \mid x_{2}\right) P\left(x_{2} \mid x_{1}\right) P\left(y_{3} \mid x_{2}\right) \\
& =\sum_{x_{2}} P\left(y_{3}, y_{2}, x_{2} \mid x_{1}\right) \\
& =P\left(y_{3}, y_{2} \mid x_{1}\right)
\end{aligned}
$$

You can generalize all the above derivation to more variable case. The general form of $b_{t}(i)$ is

$$
\begin{equation*}
b_{t}(i)=\sum_{j=1}^{s} p(j \mid i) e\left(y_{t+1} \mid j\right) b_{t+1}(j) \tag{9.1}
\end{equation*}
$$

and it represents the backward probability:

$$
\begin{equation*}
b_{t}(i)=P\left(\mathbf{y}_{(t+1): n} \mid X_{t}=i\right) \tag{9.2}
\end{equation*}
$$

and $b_{0}=P(\mathbf{y})=L(\nu, \mathbf{E}, \mathbf{P} \mid \mathbf{y})$ is the likelihood value.
The Backward algorithm works as follows:

1. Let $b_{n}(i)=1$ for all $i \in S$.
2. For $t=n-1, \cdots, 1$, compute

$$
b_{t}(i)=\sum_{j=1}^{s} p(j \mid i) e\left(y_{t+1} \mid j\right) b_{t+1}(j)
$$

for $i \in S$.
3. Finally, compute

$$
b_{0}=\sum_{j=1}^{s} \nu(j) e\left(y_{1} \mid j\right) b_{1}(j)=L(\nu, \mathbf{E}, \mathbf{P} \mid \mathbf{y})
$$

It is called the backward algorithm becomes we are taking a backward sweep from $t=n$ to $t=0$.
Here is derivation of the formula in equation (9.1) from equation (9.2). We assume that $b_{t}(i)=P\left(\mathbf{y}_{(t+1): n} \mid X_{t}=\right.$ $i)$. Then

$$
\begin{aligned}
b_{t}(i) & =P\left(\mathbf{y}_{(t+1): n} \mid X_{t}=i\right) \\
& =\sum_{j} P\left(y_{t+1}, X_{t+1}=j, \mathbf{y}_{(t+2): n} \mid X_{t}=i\right) \\
& =\sum_{j} P\left(y_{t+1} \mid X_{t+1}=j, \mathbf{y}_{(t+2): n}, X_{t}=i\right) P\left(\mathbf{y}_{(t+2): n} \mid X_{t+1}=j, X_{t}=i\right) P\left(X_{t+1}=j \mid X_{t}=i\right) \\
& =\text { using the graphical model } \\
& =\sum_{j} P\left(y_{t+1} \mid X_{t+1}=j\right) P\left(\mathbf{y}_{(t+2): n} \mid X_{t+1}=j\right) P\left(X_{t+1}=j \mid X_{t}=i\right) \\
& =\sum_{j} e\left(y_{t+1} \mid j\right) b_{t+1}(j) p(j \mid i) \\
& =(9.1)
\end{aligned}
$$

### 9.2.2 Forward algorithm

The forward algorithm relies on the idea of forward probability. To see how this idea works, we consider again our $n=3$ case but now we rewrite the probability as

$$
\begin{aligned}
L(\nu, \mathbf{P}, \mathbf{E} \mid \mathbf{y}) & =\sum_{x_{1}, x_{2}, x_{3}} \nu\left(x_{1}\right) e\left(y_{1} \mid x_{1}\right) e\left(y_{2} \mid x_{2}\right) p\left(x_{2} \mid x_{1}\right) e\left(y_{3} \mid x_{3}\right) p\left(x_{3} \mid x_{2}\right) \\
& =\sum_{x_{1}, x_{2}, x_{3}} e\left(y_{3} \mid x_{3}\right) p\left(x_{3} \mid x_{2}\right) e\left(y_{2} \mid x_{2}\right) p\left(x_{2} \mid x_{1}\right) \nu\left(x_{1}\right) e\left(y_{1} \mid x_{1}\right) \\
& =\sum_{x_{3}}^{e\left(y_{3} \mid x_{3}\right) \sum_{x_{2}} p\left(x_{3} \mid x_{2}\right) \underbrace{e\left(y_{2} \mid x_{2}\right) \sum_{x_{1}} p\left(x_{2} \mid x_{1}\right) \underbrace{\nu\left(x_{1}\right) e\left(y_{1} \mid x_{1}\right)}_{a_{2}\left(x_{2}\right)}}_{a_{3}\left(x_{3}\right)}}
\end{aligned}
$$

The quantity $a_{t}(i)$ is called the forward probability. Using a similar derivation as the backward probability, you can show that

$$
\begin{equation*}
a_{t}(i)=P\left(\mathbf{y}_{1: t}, x_{t}=i\right) \tag{9.3}
\end{equation*}
$$

The forward algorithm is a method based on the forward probability to evaluate the likelihood function. It works as follows:

1. Initialize $a_{1}(i)=\nu(i) e\left(y_{1} \mid i\right)$ for $i \in S$.

2 . For $t=1, \cdots, n-1$, compute

$$
\begin{equation*}
a_{t+1}(i)=e\left(y_{t+1} \mid i\right)\left\{\sum_{j=1}^{s} p(i \mid j) a_{t}(j)\right\} \tag{9.4}
\end{equation*}
$$

for $i \in S$.
3. Finally, we compute

$$
\sum_{j=1}^{s} a_{n}(j)=\sum_{j=1}^{s} P\left(\mathbf{y}, x_{n}=j\right)=P(\mathbf{y})=L(\nu, \mathbf{E}, \mathbf{P} \mid \mathbf{y})
$$

Here, we are sweeping from $t=1$ to $n$ so it is called the forward algorithm.

### 9.3 Parameter Estimation and the Baum-Welch algorithm

To infer $\theta=(\nu, \mathbf{P}, \mathbf{E})$, a simple approach is the MLE. However, there is no closed form solution to the MLE because we have hidden variables $\mathbf{X}$. Thus, we need to design an algorithm to find the MLE. As is mentioned in the previous lecture, we will use the EM algorithm in this case.

The EM algorithm for HMM is also called the Baum-Welch algorithm. The complete-data likelihood is

$$
P(\mathbf{x}, \mathbf{y} ; \theta)=\nu\left(x_{1}\right) \prod_{t=2}^{n} p\left(x_{t} \mid x_{t-1}\right) \prod_{r=1}^{n} e\left(y_{r} \mid x_{r}\right)
$$

and the corresponding log-likelihood is

$$
\log P(\mathbf{x}, \mathbf{y} ; \theta)=\log \nu\left(x_{1}\right)+\sum_{t=2}^{n} \log p\left(x_{t} \mid x_{t-1}\right)+\sum_{r=1}^{n} \log e\left(y_{r} \mid x_{r}\right)
$$

## E-step:

To derive the E-step, we first analyze the expected log-likelihood value given a parameter $\theta^{(k)}=\left(\nu^{(k)}, \mathbf{E}^{(k)}, \mathbf{P}^{(k)}\right)$ :

$$
\begin{aligned}
E\left\{\log P(\mathbf{x}, \mathbf{y} ; \theta) \mid \mathbf{y} ; \theta^{(k)}\right\} & =\sum_{\mathbf{x}}^{\sum_{F_{1}\left(\nu ; \theta^{(k)}\right)} \log P(\mathbf{x}, \mathbf{y} ; \theta) P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right)} \\
& =\underbrace{\sum_{\mathbf{x}} \log \nu\left(x_{1}\right) P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right)}_{F_{2}\left(\mathbf{P} ; \theta^{(k)}\right)}+\underbrace{\sum_{\mathbf{x}} \sum_{t=2}^{n} \log p\left(x_{t} \mid x_{t-1}\right) P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right)}_{F_{3}\left(\mathbf{E} ; \theta^{(k)}\right)} \\
& +\underbrace{}_{\sum_{\mathbf{x}}^{\sum_{t=1} \sum_{t}^{n} \log e\left(y_{t} \mid x_{t}\right) P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right)} .}
\end{aligned}
$$

The additive form makes the M-step a lot easier-we can separately maximize each term.

## M-step:

Updating $\nu$ :

$$
\begin{aligned}
F_{1}\left(\nu ; \theta^{(k)}\right) & =\sum_{\mathbf{x}} P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right) \log \nu\left(x_{1}\right) \\
& =\sum_{x_{1}} P\left(x_{1} \mid \mathbf{y} ; \theta^{(k)}\right) \log \nu\left(x_{1}\right) \\
& =\sum_{i} \gamma_{1}\left(i ; \theta^{(k)}\right) \log \nu(i),
\end{aligned}
$$

where

$$
\begin{equation*}
\gamma_{t}\left(i ; \theta^{(k)}\right)=P\left(x_{t}=i \mid \mathbf{y} ; \theta^{(k)}\right)=\frac{a_{t}\left(i ; \theta^{(k)}\right) b_{t}\left(i ; \theta^{(k)}\right)}{\sum_{j} a_{t}\left(j ; \theta^{(k)}\right) b_{t}\left(j ; \theta^{(k)}\right)} \tag{9.5}
\end{equation*}
$$

Note that $a_{t}(j ; \theta)$ and $b_{t}(j ; \theta)$ are the forward and backward probability assuming $\theta$ is the generating probability.

Using Lagrange multiplier, one can show that the maximizer is

$$
\nu^{(k+1)}(i)=\gamma_{1}\left(i ; \theta^{(k)}\right)=\text { imputed frequency of state } i \text { at time } 1
$$

## Updating $\mathbf{P}$ :

$$
\begin{aligned}
F_{2}\left(\mathbf{P} ; \theta^{(k)}\right) & =\sum_{\mathbf{x}} \sum_{t=2}^{n} \log p\left(x_{t} \mid x_{t-1}\right) P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right) \\
& =\sum_{t=2}^{n} \sum_{x_{t}, x_{t-1}} \log p\left(x_{t} \mid x_{t-1}\right) \sum_{\mathbf{x}_{1:(t-2)}, \mathbf{x}_{(t+1): n}} P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right) \\
& =\sum_{t=2}^{n} \sum_{i, j} P\left(X_{t}=j, X_{t-1}=i \mid \mathbf{y} ; \theta^{(k)}\right) \log p(j \mid i) \\
& =\sum_{i, j}\left\{\sum_{t=2}^{n} P\left(X_{t}=j, X_{t-1}=i \mid \mathbf{y} ; \theta^{(k)}\right)\right\} \log p(j \mid i) \\
& =\sum_{i, j} N_{t}\left(i, j ; \theta^{(k)}\right) \log p(j \mid i),
\end{aligned}
$$

where

$$
N\left(i, j ; \theta^{(k)}\right)=\sum_{t=2}^{n} P\left(X_{t}=j, X_{t-1}=i \mid \mathbf{y} ; \theta^{(k)}\right)
$$

The maximization problem is like the maximization of the Markov chain with the number of observed transitions from state $i$ to state $j$ is $N\left(i, j ; \theta^{(k)}\right)$. Thus, we know that the maximizer will be

$$
\begin{aligned}
p^{(k+1)}(j \mid i) & =\frac{N\left(i, j ; \theta^{(k)}\right)}{\sum_{j} N\left(i, j ; \theta^{(k)}\right)} \\
& =\frac{\text { expected number of transitions from } i \rightarrow j}{\text { expected number of transitions from } i}
\end{aligned}
$$

Note that there is an interesting relation between $N(i, j ; \theta)$ and $\gamma_{t}(i ; \theta)$ :

$$
\begin{aligned}
\sum_{j} N\left(i, j ; \theta^{(k)}\right) & =\sum_{j} \sum_{t=2}^{n} P\left(X_{t}=j, X_{t-1}=i \mid \mathbf{y} ; \theta^{(k)}\right) \\
& =\sum_{t=2}^{n} \sum_{j} P\left(X_{t}=j, X_{t-1}=i \mid \mathbf{y} ; \theta^{(k)}\right) \\
& =\sum_{t=2}^{n} P\left(X_{t-1}=i \mid \mathbf{y} ; \theta^{(k)}\right) \\
& =\sum_{t=2}^{n} \gamma_{t-1}\left(i ; \theta^{(k)}\right)
\end{aligned}
$$

This observation also gives a hint on how to efficiently compute $N\left(i, j ; \theta^{(k)}\right)=\sum_{t=2}^{n} P\left(X_{t}=j, X_{t-1}=\right.$ $\left.i \mid \mathbf{y} ; \theta^{(k)}\right)$ :

$$
\begin{aligned}
P\left(X_{t}=j, X_{t-1}=i \mid \mathbf{y} ; \theta^{(k)}\right) \propto & P\left(X_{t}=j, X_{t-1}=i, \mathbf{y} ; \theta^{(k)}\right) \\
= & P\left(\mathbf{y}_{(t+1): n} \mid X_{t}=j, X_{t-1}=i, \mathbf{y}_{1: t} ; \theta^{(k)}\right) P\left(\mathbf{y}_{1: t}, X_{t}=j, X_{t-1}=i ; \theta^{(k)}\right) \\
= & P\left(\mathbf{y}_{(t+1): n} \mid X_{t}=j ; \theta^{(k)}\right) P\left(y_{t} \mid \mathbf{y}_{1:(t-1)}, X_{t}=j, X_{t-1}=i ; \theta^{(k)}\right) \\
& \times P\left(X_{t}=j \mid X_{t-1}=i, \mathbf{y}_{1:(t-1)} ; \theta^{(k)}\right) P\left(X_{t-1}=i, \mathbf{y}_{1:(t-1)} ; \theta^{(k)}\right) \\
= & P\left(\mathbf{y}_{(t+1): n} \mid X_{t}=j ; \theta^{(k)}\right) e\left(y_{t} \mid j ; \theta^{(k)}\right) P\left(j \mid i ; \theta^{(k)}\right) P\left(\mathbf{y}_{1:(t-1)}, X_{t-1}=i ; \theta^{(k)}\right) \\
= & b_{t}\left(j ; \theta^{(k)}\right) e^{(k)}\left(y_{t} \mid j\right) p^{(k)}(j \mid i) a_{t-1}\left(i ; \theta^{(k)}\right) .
\end{aligned}
$$

Thus,

$$
P\left(X_{t}=j, X_{t-1}=i \mid \mathbf{y} ; \theta^{(k)}\right)=\frac{b_{t}\left(j ; \theta^{(k)}\right) e^{(k)}\left(y_{t} \mid j\right) p^{(k)}(j \mid i) a_{t-1}\left(i ; \theta^{(k)}\right)}{\sum_{i^{\prime}, j^{\prime}} b_{t}\left(j^{\prime} ; \theta^{(k)}\right) e^{(k)}\left(y_{t} \mid j^{\prime}\right) p^{(k)}\left(j^{\prime} \mid i^{\prime}\right) a_{t-1}\left(i^{\prime} ; \theta^{(k)}\right)} .
$$

We can compute this probability easily and then sum over $t$ to obtain $N\left(i, j ; \theta^{(k)}\right)$ and update $p^{(k+1)}(j \mid i)$.
Updating E:

$$
\begin{aligned}
F_{3}\left(\mathbf{E} ; \theta^{(k)}\right) & =\sum_{\mathbf{x}} \sum_{t=1}^{n} \log e\left(y_{t} \mid x_{t}\right) P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right) \\
& =\sum_{t=1}^{n} \sum_{x_{t}} \log e\left(y_{t} \mid x_{t}\right) \sum_{\mathbf{x}_{-t}} P\left(\mathbf{x} \mid \mathbf{y} ; \theta^{(k)}\right) \\
& =\sum_{t=1}^{n} \sum_{x_{t}} P\left(x_{t} \mid \mathbf{y} ; \theta^{(k)}\right) \log e\left(y_{t} \mid x_{t}\right) \\
& =\sum_{x_{t}} \sum_{t=1}^{n} \gamma_{t}\left(x_{t} ; \theta^{(k)}\right) \log e\left(y_{t} \mid x_{t}\right)
\end{aligned}
$$

Again, using Lagrange multiplier, the maximizer is

$$
e^{(k+1)}(\ell \mid i)=\frac{\sum_{t=1}^{n} \gamma_{t}\left(i ; \theta^{(k)}\right) I\left(y_{t}=\ell\right)}{\sum_{t=1}^{n} \gamma_{t}\left(i ; \theta^{(k)}\right)}
$$

With all the above updating equation, the Baum-Welch algorithm is

1. Start with $\theta^{(0)}=\left(\nu^{(0)}, \mathbf{P}^{(0)}, \mathbf{E}^{(0)}\right)$. Repeat the following steps until some convergence criterion is met:
(a) Compute the forward and backward probabilities $\left\{a_{t}\left(i \mid \theta^{(k)}\right)\right\}$ and $\left\{b_{t}\left(i \mid \theta^{(k)}\right)\right\}$.
(b) Update HMM parameters

$$
\begin{aligned}
\nu^{(k+1)}(i) & =\gamma_{1}\left(i ; \theta^{(k)}\right) \\
p^{(k+1)}(j \mid i) & =\frac{N\left(i, j ; \theta^{(k)}\right)}{\sum_{j} N\left(i, j ; \theta^{(k)}\right)} \\
e^{(k+1)}(\ell \mid i) & =\frac{\sum_{t=1}^{n} \gamma_{t}\left(i ; \theta^{(k)}\right) I\left(y_{t}=\ell\right)}{\sum_{t=1}^{n} \gamma_{t}\left(i ; \theta^{(k)}\right)} .
\end{aligned}
$$

### 9.4 Hidden State Reconstruction and Viterbi Algorithm

Note that here we assume that the parameters $\nu, \mathbf{P}, \mathbf{E}$ are given so we can easily evaluate the forward and the backward probabilities. We can estimate these parameters using the BW algorithm.

The forward and backward probabilities together imply a very interesting result:

$$
\begin{aligned}
P\left(X_{t}=i \mid \mathbf{y}\right) & =\frac{P\left(X_{t}=i, \mathbf{y}\right)}{P(\mathbf{y})} \\
& =\frac{P\left(\mathbf{y}_{1: t}, \mathbf{y}_{(t+1): n}, X_{t}=i\right)}{P(\mathbf{y})} \\
& =\frac{P\left(\mathbf{y}_{(t+1): n} \mid \mathbf{y}_{1: t}, X_{t}=i\right) P\left(\mathbf{y}_{1: t}, X_{t}=i\right)}{P(\mathbf{y})} \\
& =\text { using graphical model } \\
& =\frac{P\left(\mathbf{y}_{(t+1): n} \mid X_{t}=i\right) P\left(\mathbf{y}_{1: t}, X_{t}=i\right)}{P(\mathbf{y})} \\
& =\frac{a_{t}(i) b_{t}(i)}{\sum_{j=1}^{s} a_{t}(j) b_{t}(j)}
\end{aligned}
$$

Namely,

$$
P\left(X_{t}=i \mid \mathbf{y}\right) \propto a_{t}(i) b_{t}(i)
$$

so this gives us the marginal probability of a hidden state at time $t$. Thus, a method to reconstruct the hidden state is

$$
\widehat{x}_{t}=\operatorname{argmax}_{i} a_{t}(i) b_{t}(i)
$$

and by applying this to every $t$, we obtain

$$
\widehat{x}_{1}, \cdots, \widehat{x}_{n}
$$

However, this idea is based on the marginal probability, not the joint probability. So it may give us a path such that $P\left(\widehat{x}_{t+1} \mid \widehat{x}_{t}\right)=0$ !
A better way is to find the states according to the joint distribution function. Namely,

$$
\mathbf{x}^{*}=\operatorname{argmax}_{\mathbf{x}} P(\mathbf{x} \mid \mathbf{y})=\operatorname{argmax}_{\mathbf{x}} P(\mathbf{x}, \mathbf{y}) .
$$

In this case, $\mathbf{x}^{*}$ is the MAP estimator. Note that $\mathbf{x}^{*}$ involves $n$ elements so there are totally $s^{n}$ possible elements we need to search to obtain $\mathbf{x}^{*}$. Ideally, we want to avoid searching all possible configurations. Here we will use the idea that for a bivariate function $f(i, j)$,

$$
\max _{i, j} f(i, j)=\max _{j}\left\{\max _{i} f(i, j)\right\}
$$

The joint distribution function up to time point $t+1$ can be factored as

$$
\begin{aligned}
P\left(\mathbf{x}_{1:(t+1)}, \mathbf{y}_{1:(t+1)}\right) & =P\left(x_{t+1}, y_{t+1} \mid \mathbf{x}_{1: t}, \mathbf{y}_{1: t}\right) P\left(\mathbf{x}_{1: t}, \mathbf{y}_{1: t}\right) \\
& =P\left(y_{t+1} \mid x_{t+1}, \mathbf{x}_{1: t}, \mathbf{y}_{1: t}\right) P\left(x_{t+1} \mid \mathbf{x}_{1: t}, \mathbf{y}_{1: t}\right) P\left(\mathbf{x}_{1: t}, \mathbf{y}_{1: t}\right) \\
& =\text { using graphical model } \\
& =P\left(y_{t+1} \mid x_{t+1}\right) P\left(x_{t+1} \mid x_{t}\right) P\left(\mathbf{x}_{1: t}, \mathbf{y}_{1: t}\right) \\
& =e\left(y_{t+1} \mid x_{t+1}\right) P\left(x_{t+1} \mid x_{t}\right) P\left(\mathbf{x}_{1: t}, \mathbf{y}_{1: t}\right) .
\end{aligned}
$$

This implies a very interesting recursive relation. Define

$$
d_{t}(i)=\max _{\mathbf{x}_{1:(t-1)}} P\left(\mathbf{x}_{1:(t-1)}, x_{t}=i, \mathbf{y}_{1: t}\right)
$$

for $i \in S$. Then for $t=2, \cdots, n$ and $i=1, \cdots, s$,

$$
\begin{aligned}
d_{t}(i) & =\max _{\mathbf{x}_{1:(t-1)}} P\left(\mathbf{x}_{1:(t-1)}, x_{t}=i, \mathbf{y}_{1: t}\right) \\
& =\max _{\mathbf{x}_{1:(t-1)}} e\left(y_{t} \mid x_{t}=i\right) P\left(x_{t}=i \mid x_{t-1}\right) P\left(\mathbf{x}_{1:(t-1)}, \mathbf{y}_{1:(t-1)}\right) \\
& =e\left(y_{t} \mid x_{t}=i\right) \max _{\mathbf{x}_{1:(t-1)}} P\left(x_{t}=i \mid x_{t-1}\right) P\left(\mathbf{x}_{1:(t-1)}, \mathbf{y}_{1:(t-1)}\right) \\
& =e\left(y_{t} \mid x_{t}=i\right) \max _{x_{t-1}} P\left(x_{t}=i \mid x_{t-1}\right) \underbrace{\max _{1:(t-2)} P\left(\mathbf{x}_{1:(t-1)}, \mathbf{y}_{1:(t-1)}\right)}_{d_{t-1}\left(x_{t-1}\right)} \\
& =e\left(y_{t} \mid x_{t}=i\right)\left\{\max _{x_{t-1}} P\left(x_{t}=i \mid x_{t-1}\right) d_{t-1}\left(x_{t-1}\right)\right\} .
\end{aligned}
$$

Again, we obtain a recursive relation between each $d_{t}(i)$.
This suggests a recursive approach to find the joint maximizer, which is known as the Viterbi algorithm:

1. Compute $d_{1}(i)=\nu(i) e\left(y_{1} \mid i\right)$ for each $i \in S$.
2. For $t=2, \cdots, n$, compute

$$
\begin{aligned}
d_{t}(i) & =e\left(y_{t} \mid i\right) \cdot \max _{j} d_{t-1}(j) p(i \mid j) \quad \text { for all } i \in S \\
f_{t}(i) & =\operatorname{argmax}_{j} d_{t-1}(j) p(i \mid j) \quad \text { for all } i \in S
\end{aligned}
$$

3. Let $p^{*}=\max _{j} d_{n}(j)$ and $x_{n}^{*}=\operatorname{argmax}_{j} d_{n}(j)$.
4. (backtracking) For $t=n-1, \cdots, 1$ :

$$
x_{t}^{*}=f_{t+1}\left(x_{t+1}^{*}\right) .
$$

5. Return $\mathbf{x}^{*}=\left(x_{1}^{*}, \cdots, x_{n}^{*}\right)$.

The logic of Viterbi algorithm is that for a function $g(i, j)$

$$
\left(i^{*}, j^{*}\right)=\operatorname{argmax}_{i, j} g(i, j) \Rightarrow j^{*}=\operatorname{argmax}\left\{\max _{i} g(i, j)\right\}
$$

It is clear that $x_{n}^{*}=\operatorname{argmax}_{j} d_{n}(j)$ is indeed the maximizer at the last time point. To see that $x_{t}^{*}$ indeed maximizes the joint probability, we consider the case of $n-1$ (and you can generalize it to any $t$ using induction). Given $x_{n}^{*}$, the true maximizer of $x_{n-1}$ should be

$$
\begin{aligned}
x_{n-1}^{*} & =\operatorname{argmax}_{j} \max _{\mathbf{x}_{1:(n-2)}} P\left(\mathbf{x}_{1: n-2}, x_{n-1}=j, x_{n}=x_{n}^{*}, \mathbf{y}\right) \\
& =\operatorname{argmax}_{j} \max _{\mathbf{x}_{1:(n-2)}} e\left(y_{n} \mid x_{n}=x_{n}^{*}\right) P\left(x_{n}=x_{n}^{*} \mid x_{n-1}=j\right) P\left(\mathbf{x}_{1:(n-2)}, x_{n-1}=j, \mathbf{y}_{1:(n-1)}\right) \\
& =\operatorname{argmax}_{j} P\left(x_{n}=x_{n}^{*} \mid x_{n-1}=j\right) \max _{\mathbf{x}_{1:(n-2)}} P\left(\mathbf{x}_{1:(n-2)}, x_{n-1}=j, \mathbf{y}_{1:(n-1)}\right) \\
& =\operatorname{argmax}_{j} P\left(x_{n}=x_{n}^{*} \mid x_{n-1}=j\right) d_{n-1}(j) \\
& =f_{n}\left(x_{n}^{*}\right) .
\end{aligned}
$$

Thus, indeed the Viterbi algorithm finds the path that maximizes the joint probability.
Note that when $n$ is large, the rounding error of computing the probability could be a problem. To avoid this problem, one can move all calculation to the log scale and use

$$
\tilde{d}_{t}(i)=\max _{\mathbf{x}_{1:(t-1)}} \log P\left(\mathbf{x}_{1:(t-1)}, x_{t}=i, \mathbf{y}_{1: t}\right)
$$

We will obtain another recursive relation:

$$
\tilde{d}_{t}(i)=\log e\left(y_{t} \mid i\right)+\max _{j}\left\{\tilde{d}_{t-1}(j)+\log p(i \mid j)\right\}
$$

### 9.5 Forecasting

The forecasting problem is that case where we have observations up to time point $n$ and we would like to predict the possible outcomes at time point $n+h$ for $h=1,2, \cdots$. Specifically, we want to obtain a forecast probability of each possible outcome $y \in M$. Using the notion of probability, what we really want is $P\left(Y_{n+1}=y_{n+1} \mid \mathbf{y}\right)$ for each $y_{n+1} \in M$.

We start with the simple case $h=1$.

$$
\begin{aligned}
P\left(Y_{n+1}=y_{n+1} \mid \mathbf{y}\right) & =\sum_{i} P\left(Y_{n+1}=y_{n+1}, X_{n}=i \mid \mathbf{y}\right) \\
& =\sum_{i} P\left(Y_{n+1}=y_{n+1} \mid X_{n}=i\right) P\left(X_{n}=i \mid \mathbf{y}\right) \\
& =\sum_{i}\left\{\sum_{j} P\left(Y_{n+1}=y_{n+1}, X_{n+1}=j \mid X_{n}=i\right)\right\} P\left(X_{n}=i \mid \mathbf{y}\right) \\
& =\sum_{i}\left\{\sum_{j} P\left(Y_{n+1}=y_{n+1} \mid X_{n+1}=j\right) p(j \mid i)\right\} P\left(X_{n}=i \mid \mathbf{y}\right) .
\end{aligned}
$$

Recall that $P\left(X_{n}=i \mid \mathbf{y}\right)=\frac{b_{n}(i) a_{n}(i)}{\sum_{k} b_{n}(k) a_{n}(k)}=\frac{a_{n}(i)}{\sum_{k} a_{n}(k)}$ since $b_{n}(k)=1$ for all $k$. Using this fact, we can rewrite the above as

$$
\begin{aligned}
P\left(Y_{n+1}=y_{n+1} \mid \mathbf{y}\right) & =\sum_{i}\left\{\sum_{j} P\left(Y_{n+1}=y_{n+1} \mid X_{n+1}=j\right) P(j \mid i)\right\} P\left(X_{n}=i \mid \mathbf{y}\right) \\
& =\frac{\sum_{i}\left\{\sum_{j} e\left(y_{n+1} \mid j\right) p(j \mid i)\right\} a_{n}(i)}{\sum_{k} a_{n}(k)}
\end{aligned}
$$

With an estimate of $e$ and $P(j \mid i)$ and the corresponding forward probability, we can easily compute the forecast probability of each state $y_{n+1} \in M$.

For a general $h$ case (forecasting $h$-time point future), we have

$$
\begin{aligned}
P\left(Y_{n+h}=y_{n+h} \mid \mathbf{y}\right) & =\sum_{i} P\left(Y_{n+h}=y_{n+h}, X_{n}=i \mid \mathbf{y}\right) \\
& =\sum_{i}\left\{\sum_{j} P\left(Y_{n+h}=y_{n+h}, X_{n+h}=j \mid X_{n}=i, \mathbf{y}\right)\right\} P\left(X_{n}=i \mid \mathbf{y}\right) \\
& =\sum_{i}\left\{\sum_{j} P\left(Y_{n+h}=y_{n+h} \mid X_{n+h}=j\right) P\left(X_{n+h}=j \mid X_{n}=i\right)\right\} P\left(X_{n}=i \mid \mathbf{y}\right) \\
& =\sum_{i}\left\{\sum_{j} e\left(y_{n+h} \mid j\right) p^{(h)}(j \mid i)\right\} P\left(X_{n}=i \mid \mathbf{y}\right) \\
& =\frac{\sum_{i}\left\{\sum_{j} e\left(y_{n+h} \mid j\right) p^{(h)}(j \mid i)\right\} a_{n}(i)}{\sum_{k} a_{n}(k)},
\end{aligned}
$$

where $p^{(h)}(j \mid i)$ is the $h$-step transition probability of the Markov chain. Using this equation, we can make prediction about the likelihood of each outcome in the future.

### 9.6 Model Selection

In all of the above analysis, we assume that the total number of hidden states $s$ is known. However, in reality, we often do not know this number so we need to choose it using some smart trick. The study of selecting the right amount of hidden states is part of a general problem called model selection.

An intuitive way is to consider various $s$ and see which one gives the highest likelihood value. However, this idea will suffer the problem of overfitting-the more hidden states we fit to the data, the higher likelihood value we will obtain.

This overfitting problem is something we need to be very cautious about when analyzing the data. A common approach to model selection is to use a penalized criterion. For the HMM with $s$ hidden states, let $\ell_{n}^{*}(s)$ be the log-likelihood function corresponding to the MLE of this model. We then select $s$ by maximizing

$$
\ell_{n}^{*}(s)-\Psi(s)
$$

where $\Psi(s)$ is an increasing function with respect to $s$. We call $\Psi(s)$ the penalty or regularizer. In many cases, we frame the problem as the minimizer problem and choose $s$ by minimizing

$$
R(s)=-2 \ell_{n}^{*}(s)+2 \Psi(s)
$$

While there are many choices of $\Psi(s)$, two most popular ones are the Akaike information criterion (AIC) and Bayesian information criterion (BIC), which leads to

$$
R_{\mathrm{AIC}}(s)=-2 \ell_{n}^{*}(s)+2 p, \quad R_{\mathrm{BIC}}(s)=-2 \ell_{n}^{*}(s)+p \log n
$$

where $p=p(s)$ is the total number of parameters. In an HMM,

$$
p(s)=\underbrace{(s-1)}_{\sim \nu}+\underbrace{s(s-1)}_{\sim \mathbf{P}}+\underbrace{s(m-1)}_{\sim \mathbf{E}}=s^{2}-1+s(m-1) .
$$

AIC has a smaller penalty compared to BIC so it often selects a model with more variables. AIC aims at selecting the model that is the best for prediction under certain criterion whereas BIC aims at finding the true model according to Bayesian's point of view.

### 9.6.1 Intuition of AIC and BIC

Here we use an informal way of deriving AIC and BIC that gives us some intuition on how they are constructed ${ }^{1}$. Note that the notations in this subsection is independent of the notations in the previous sections. To illustrate the idea, we assume that we observed $X_{1}, \cdots, X_{n}$ that are IID from an unknown PDF $p_{0}(x)$. And we consider different models

$$
M_{j}=\left\{p\left(x ; \theta_{j}\right): \theta_{j} \in \Theta_{j}\right\}, \quad \operatorname{dim}\left(\Theta_{j}\right)=d_{j}
$$

for $j=1, \cdots, m$.
Intuition of AIC: Let $\widehat{p}_{j}(x)=p\left(x ; \widehat{\theta}_{j}\right)$ with $\widehat{\theta}_{j}$ be the MLE under parameter space $\Theta_{j}$. Recall that a common measure of closeness between a fitted model $\widehat{p}_{j} \in M_{j}$ and the true model is the KL divergence:

$$
K\left(p_{0}, \widehat{p}_{j}\right)=\int p_{0}(x) \log \left(\frac{p_{0}(x)}{\widehat{p}_{j}(x)}\right)=-\int p_{0}(x) \log \widehat{p}_{j}(x)+\int p_{0}(x) \log p_{0}(x) d x .
$$

Since the last term is independent of $j$, minimizing $K\left(p_{0}, \widehat{p}_{j}\right)$ across different $j$ is equivalent to maximizing

$$
K_{j}=\int p_{0}(x) \log \widehat{p}_{j}(x)=\int p_{0}(x) \log p\left(x ; \widehat{\theta}_{j}\right)
$$

AIC aims at choosing $j$ that minimizing the KL divergence so it is equivalent to choosing $j$ maximizing $K_{j}$. However, $K_{j}$ involves $p_{0}$ so we have to estimate it. A natural estimate of $K_{j}$ is to use the data, leading to

$$
\widehat{K}_{j}=\frac{1}{n} \sum_{i=1}^{n} \log p\left(X_{i} ; \widehat{\theta}_{j}\right)
$$

However, this method is a biased estimator since the data is used both in finding the MLE $\widehat{\theta}_{j}$ and evaluating $\widehat{K}_{j}$.

Akaike showed that the bias of $\widehat{K}_{j}$ is approximately $\frac{d_{j}}{n}$, which motivates us to use

$$
\widehat{K}_{j}-\frac{d_{j}}{n}
$$

as an estimate of $K_{j}$. Now, notice that the $\log$-likelihood function $\ell_{n}(\theta)=\sum_{i=1}^{n} \log p\left(X_{i} ; \theta\right)=n \widehat{K}_{j}$. So maximizing $\widehat{K}_{j}-\frac{d_{j}}{n}$ is equivalent to minimizing

$$
-2 n \widehat{K}_{j}+2 d_{j}=-2 \ell_{n}\left(\widehat{\theta}_{j}\right)+2 d_{j}
$$

which is the AIC criterion.
Intuition of BIC: The original idea of BIC is from the posterior probability-we want to choose the model with a highest posterior probability. But it is easier to understand it using the Bayesian evidence (very

[^0]similar to the Bayes factor). Let $\pi\left(M_{j}\right)$ be the posterior probability of $j$-th model. Recall that in terms of Bayesian evidence, we want to choose the model that maximizes
$$
p\left(X_{1}, \cdots, X_{n} \mid M_{j}\right)
$$

Assume that for model $j$, we have a prior distribution $\pi_{j}$ on the parameter $\theta_{j} \in \Theta_{j}$. Then the evidence can be written as

$$
\begin{aligned}
p\left(X_{1}, \cdots, X_{n} \mid M_{j}\right) & =\int p\left(X_{1}, \cdots, X_{n} \mid \theta_{j}\right) \pi_{j}\left(\theta_{j}\right) d \theta_{j} \\
& =\int L\left(\theta_{j} \mid X_{1}, \cdots, X_{n}\right) \pi_{j}\left(\theta_{j}\right) d \theta_{j} \\
& =\int e^{\sum_{i=1}^{n} \ell\left(\theta_{j} \mid X_{i}\right)} \pi_{j}\left(\theta_{j}\right) d \theta_{j}
\end{aligned}
$$

When the likelihood model is smooth, the expected log-likelihood function $\mathbb{E}\left(\ell\left(\theta_{j} \mid X_{1}\right)\right)$ is a quadratic function around its model. So when $n$ is large, Tayler expansion gives us

$$
\frac{1}{n} \sum_{i=1}^{n} \ell\left(\theta_{j} \mid X_{i}\right) \approx \frac{1}{n} \sum_{i=1}^{n} \ell\left(\widehat{\theta}_{j} \mid X_{i}\right)+\frac{1}{2}\left(\theta_{j}-\widehat{\theta}_{j}\right)^{T}\left(\frac{1}{n} \sum_{i=1}^{n} \ell^{\prime \prime}\left(\widehat{\theta}_{j} \mid X_{i}\right)\right)\left(\theta_{j}-\widehat{\theta}_{j}\right)
$$

Also,

$$
\sum_{i=1}^{n} \ell^{\prime \prime}\left(\widehat{\theta}_{j} \mid X_{i}\right) \approx n H\left(\widehat{\theta}_{j}\right)
$$

where $H(\theta)=\mathbb{E}\left(\ell^{\prime \prime}\left(\theta \mid X_{i}\right)\right)$. Thus, putting these into exponent, we obtain

$$
e^{\sum_{i=1}^{n} \ell\left(\theta_{j} \mid X_{i}\right)} \approx e^{\sum_{i=1}^{n} \ell\left(\widehat{\theta}_{j} \mid X_{i}\right)+\frac{1}{2}\left(\theta_{j}-\widehat{\theta}_{j}\right)^{T} n H\left(\widehat{\theta}_{j}\right)\left(\theta_{j}-\widehat{\theta}_{j}\right)} .
$$

Therefore, the logarithmic of the Bayesian evidence (taking log will not affect the maximizer) will be

$$
\begin{aligned}
\log p\left(X_{1}, \cdots, X_{n} \mid M_{j}\right) & \approx \log \int e^{\sum_{i=1}^{n} \ell\left(\widehat{\theta}_{j} \mid X_{i}\right)+\frac{1}{2}\left(\theta_{j}-\widehat{\theta}_{j}\right)^{T} n H\left(\widehat{\theta}_{j}\right)\left(\theta_{j}-\widehat{\theta}_{j}\right)} \pi_{j}\left(\theta_{j}\right) d \theta_{j} \\
& =\sum_{i=1}^{n} \ell\left(\widehat{\theta}_{j} \mid X_{i}\right)+\log \int e^{\frac{1}{2}\left(\theta_{j}-\widehat{\theta}_{j}\right)^{T} n H\left(\widehat{\theta}_{j}\right)\left(\theta_{j}-\widehat{\theta}_{j}\right)} \pi_{j}\left(\theta_{j}\right) d \theta_{j} \\
& \approx \sum_{i=1}^{n} \ell\left(\widehat{\theta}_{j} \mid X_{i}\right)+\log \left(\sqrt{(2 \pi / n)^{p} \operatorname{det}\left(H^{-1}\left(\widehat{\theta}_{j}\right)\right)}\right) \\
& \approx \sum_{i=1}^{n} \ell\left(\widehat{\theta}_{j} \mid X_{i}\right)-\frac{p}{2} \log n
\end{aligned}
$$

which is what the BIC criterion is about.


[^0]:    ${ }^{1}$ If you are interested in more details, I would recommend reading http://www.stat.cmu.edu/~larry/=stat705/Lecture16. pdf

