

Hidden Markov Models

Mark Hasegawa-Johnson
These slides are in the public domain

ECE 417: Multimedia Signal Processing

- 1 Hidden Markov Models
- 2 Recognition: the Forward Algorithm
- 3 Segmentation: the Backward Algorithm
- 4 Numerical Example
- 5 Summary
- 6 Written Example

Outline

- 1 Hidden Markov Models
- 2 Recognition: the Forward Algorithm
- 3 Segmentation: the Backward Algorithm
- 4 Numerical Example
- 5 Summary
- 6 Written Example


Notation: Inputs and Outputs

- Let's assume we have T consecutive observations,
 $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$.
- A “hidden Markov model” represents those probabilities by assuming some sort of “hidden” state sequence,
 $\mathbf{q} = [q_1, \dots, q_T]^T$, where q_t is the hidden (unknown) state variable at time t .

The idea is, can we model these probabilities well enough to solve problems like:


- 1 **Recognition:** What's $\Pr\{\mathbf{X}\}$ given the model?
- 2 **Segmentation:** What state is the model in at time t ?
- 3 **Training:** Can we learn a model to fit some data?

Notation: Inputs and Outputs



State Sequence $Q = [i \quad i \quad i \quad j \quad j \quad k \quad k \quad k \quad \dots]$

Observations $O = [o_1 \quad o_2 \quad o_3 \quad o_4 \quad o_5 \quad o_6 \quad o_7 \quad o_8 \quad o_9]$




HMM: Key Concepts

An HMM is a “generative model,” meaning that it models the joint probability $\Pr\{\mathbf{q}, \mathbf{X}\}$ using a model of the way in which those data might have been generated. An HMM pretends the following generative process:

- 1 Start in state $q_t = i$ with pmf $\pi_i = \Pr\{q_1 = i\}$.
- 2 Generate an observation, \mathbf{x} , with pdf $b_i(\mathbf{x}) = \Pr\{\mathbf{x}|q_t = i\}$.
- 3 Transition to a new state, $q_{t+1} = j$, according to pmf $a_{i,j} = \Pr\{q_{t+1} = j|q_t = i\}$.
- 4 Repeat.

HMM: Finite State Diagram



- ① Start in state $q_t = i$, for some $1 \leq i \leq N$.
- ② Generate an observation, \mathbf{x} , with pdf $b_i(\mathbf{x})$.
- ③ Transition to a new state, $q_{t+1} = j$, according to pmf a_{ij} .
- ④ Repeat steps #2 and #3, T times each.

Notation: Model Parameters

Solving an HMM is possible if you **carefully keep track of notation**. Here's standard notation for the parameters:

- $\pi_i = \Pr\{q_1 = i\}$ is called the **initial state probability**. Let N be the number of different states, so that $1 \leq i \leq N$.
- $a_{i,j} = \Pr\{q_t = j | q_{t-1} = i\}$ is called the **transition probability**, $1 \leq i, j \leq N$.
- $b_j(\mathbf{x}) = \Pr\{\mathbf{x}_t = \mathbf{x} | q_t = j\}$ is called the **observation probability**. It is usually estimated by a neural network, though Gaussians, GMMs, and even lookup tables are possible.
- Λ is the complete set of **model parameters**, including all the π_i 's and $a_{i,j}$'s, and the Gaussian, GMM, or neural net parameters necessary to compute $b_j(\mathbf{x})$.

The Three Problems for an HMM

1 Recognition:

- Suppose our ASR knows two different words. Each word is represented by a different set of model parameters: Λ_1 and Λ_2 .
- Suppose we have a test spectrogram, \mathbf{X} . We know that \mathbf{X} is one of those two words, but we don't know which.
- Which HMM was more likely to have produced \mathbf{X} ? In other words, is $\Pr\{\mathbf{X}|\Lambda_1\} > \Pr\{\mathbf{X}|\Lambda_2\}$?

2 Segmentation: What is $\Pr\{q_t = i|\mathbf{X}, \Lambda\}$?

3 Training: Given an initial HMM Λ , and an observation sequence \mathbf{X} , can we find Λ' such that $\Pr\{\mathbf{X}|\Lambda'\} > \Pr\{\mathbf{X}|\Lambda\}$?

Outline

- 1 Hidden Markov Models
- 2 Recognition: the Forward Algorithm**
- 3 Segmentation: the Backward Algorithm
- 4 Numerical Example
- 5 Summary
- 6 Written Example

The HMM Recognition Problem

- Given
 - $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$ and
 - $\Lambda = \{\pi_i, a_{i,j}, b_j(\mathbf{x}) \forall i, j\}$,
 what is $\Pr\{\mathbf{X}|\Lambda\}$?

- Let's solve a simpler problem first:

- Given
 - $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$ and
 - $\mathbf{q} = [q_1, \dots, q_T]^T$ and
 - $\Lambda = \{\pi_i, a_{i,j}, b_j(\mathbf{x}) \forall i, j\}$,
 what is $\Pr\{\mathbf{X}|\Lambda\}$?

Joint Probability of State Sequence and Observation Sequence

The joint probability of the state sequence and the observation sequence is calculated iteratively, from beginning to end:

- The probability that $q_1 = q_1$ is π_{q_1} .
- Given q_1 , the probability of \mathbf{x}_1 is $b_{q_1}(\mathbf{x}_1)$.
- Given q_1 , the probability of q_2 is $a_{q_1 q_2}$.
- ... and so on...

$$\Pr\{\mathbf{q}, \mathbf{X} | \Lambda\} = \pi_{q_1} b_{q_1}(\mathbf{x}_1) \prod_{t=2}^T a_{q_{t-1} q_t} b_{q_t}(\mathbf{x}_t)$$

Probability of the Observation Sequence

The probability of the observation sequence, alone, is somewhat harder, because we have to solve this sum:

$$\begin{aligned} \Pr\{\mathbf{X}|\Lambda\} &= \sum_{\mathbf{q}} \Pr\{\mathbf{q}, \mathbf{X}|\Lambda\} \\ &= \sum_{q_T=1}^N \cdots \sum_{q_1=1}^N \Pr\{\mathbf{q}, \mathbf{X}|\Lambda\} \end{aligned}$$

On the face of it, this calculation seems to have complexity $\mathcal{O}\{N^T\}$. So for a very small 100-frame utterance, with only 10 states, we have a complexity of $\mathcal{O}\{10^{100}\}$ = one google.

The Forward Algorithm

The solution is to use a kind of dynamic programming algorithm, called “the forward algorithm.” The forward probability is defined as follows:

$$\alpha_t(i) \equiv \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = i | \Lambda\}$$

Obviously, if we can find $\alpha_t(i)$ for all i and all t , we will have solved the recognition problem, because

$$\begin{aligned} \Pr\{\mathbf{X} | \Lambda\} &= \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_T | \Lambda\} \\ &= \sum_{i=1}^N \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_T, q_T = i | \Lambda\} \\ &= \sum_{i=1}^N \alpha_T(i) \end{aligned}$$

The Forward Algorithm

So, working with the definition $\alpha_t(i) \equiv \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = i | \Lambda\}$, let's see how we can actually calculate $\alpha_t(i)$.

1 Initialize:

$$\begin{aligned} \alpha_1(i) &= \Pr\{q_1 = i, \mathbf{x}_1 | \Lambda\} \\ &= \Pr\{q_1 = i | \Lambda\} \Pr\{\mathbf{x}_1 | q_1 = i, \Lambda\} \\ &= \pi_i b_i(\mathbf{x}_1) \end{aligned}$$

The Forward Algorithm

Definition: $\alpha_t(i) \equiv \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = i | \Lambda\}$.

① **Initialize:**

$$\alpha_1(i) = \pi_i b_i(\mathbf{x}_1), \quad 1 \leq i \leq N$$

② **Iterate:**

$$\alpha_t(j) = \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = j | \Lambda\}$$

$$= \sum_{i=1}^N \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_{t-1}, q_{t-1} = i\} \Pr\{q_t = j | q_{t-1} = i\} \Pr\{\mathbf{x}_t | q_t = j\}$$

$$= \sum_{i=1}^N \alpha_{t-1}(i) a_{i,j} b_j(\mathbf{x}_t)$$

The Forward Algorithm

So, working with the definition $\alpha_t(i) \equiv \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = i | \Lambda\}$, let's see how we can actually calculate $\alpha_t(i)$.

① **Initialize:**

$$\alpha_1(i) = \pi_i b_i(\mathbf{x}_1), \quad 1 \leq i \leq N$$


② **Iterate:**

$$\alpha_t(j) = \sum_{i=1}^N \alpha_{t-1}(i) a_{i,j} b_j(\mathbf{x}_t), \quad 1 \leq j \leq N, \quad 2 \leq t \leq T$$

③ **Terminate:**

$$\Pr\{\mathbf{X} | \Lambda\} = \sum_{i=1}^N \alpha_T(i)$$


Visualizing the Forward Algorithm using a Trellis



Using a trellis, the **initialize** step computes probabilities for the first column of the trellis:

$$\alpha_1(i) = \pi_i b_i(\mathbf{x}_1), \quad 1 \leq i \leq N$$


Visualizing the Forward Algorithm using a Trellis



The **iterate** step then computes the probabilities in the t^{th} column by adding up the probabilities in the $(t-1)^{\text{st}}$ column, each multiplied by the corresponding transition probability:

$$\alpha_t(j) = \sum_{i=1}^N \alpha_{t-1}(i) a_{i,j} b_j(\mathbf{x}_t), \quad 1 \leq j \leq N, \quad 2 \leq t \leq T$$

Visualizing the Forward Algorithm using a Trellis



The **terminate** step then computes the likelihood of the model by adding the probabilities in the last column:

$$\Pr\{\mathbf{X}|\Lambda\} = \sum_{i=1}^N \alpha_T(i)$$

The Forward Algorithm: Computational Complexity

Most of the computational complexity is in this step:

- **Iterate:**

$$\alpha_t(j) = \sum_{i=1}^N \alpha_{t-1}(i) a_{i,j} b_j(\mathbf{x}_t), \quad 1 \leq i, j \leq N, \quad 2 \leq t \leq T$$

Its complexity is:

- For each of $T - 1$ time steps, $2 \leq t \leq T, \dots$
- we need to calculate N different alpha-variables, $\alpha_t(j)$, for $1 \leq j \leq N, \dots$
- each of which requires a summation with N terms.

So the total complexity is $\mathcal{O}\{TN^2\}$. For example, with $N = 10$ and $T = 100$, the complexity is only $TN^2 = 10,000$ multiplies (much, much less than N^T !!)

Outline

- 1 Hidden Markov Models
- 2 Recognition: the Forward Algorithm
- 3 Segmentation: the Backward Algorithm**
- 4 Numerical Example
- 5 Summary
- 6 Written Example

The Segmentation Problem

There are different ways to define the segmentation problem. Let's define it this way:

- We want to find the most likely state, $q_t = i$, at time t, \dots
- given knowledge of the *entire* sequence $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]$, not just the current observation. So for example, we don't want to recognize state i at time t if the surrounding observations, \mathbf{x}_{t-1} and \mathbf{x}_{t+1} , make it obvious that this choice is impossible. Also, \dots
- given knowledge of the HMM that produced this sequence, Λ .

In other words, we want to find the **state posterior probability**, $\Pr\{q_t = i | \mathbf{X}, \Lambda\}$. Let's define some more notation for the state posterior probability, let's call it

$$\gamma_t(i) = \Pr\{q_t = i | \mathbf{X}, \Lambda\}$$

Use Bayes' Rule

Suppose we already knew the **joint probability**, $\Pr\{\mathbf{X}, q_t = i | \Lambda\}$.
Then we could find the state posterior using Bayes' rule:

$$\gamma_t(i) = \Pr\{q_t = i | \mathbf{X}, \Lambda\} = \frac{\Pr\{\mathbf{X}, q_t = i | \Lambda\}}{\sum_{j=1}^N \Pr\{\mathbf{X}, q_t = j | \Lambda\}}$$

Use the Forward Algorithm

Let's expand this:

$$\Pr\{\mathbf{X}, q_t = i | \Lambda\} = \Pr\{q_t = i, \mathbf{x}_1, \dots, \mathbf{x}_T | \Lambda\}$$

We already know about half of that:

$\alpha_t(i) = \Pr\{q_t = i, \mathbf{x}_1, \dots, \mathbf{x}_t | \Lambda\}$. We're only missing this part:

$$\Pr\{\mathbf{X}, q_t = i | \Lambda\} = \alpha_t(i) \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\}$$

Again, let's try the trick of "solve the problem by inventing new notation." Let's define

$$\beta_t(i) \equiv \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\}$$

The Backward Algorithm

Now let's use the definition $\beta_t(i) \equiv \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\}$, and see how we can compute that.

1 Initialize:

$$\beta_T(i) = 1, \quad 1 \leq i \leq N$$

This might not seem immediately obvious, but think about it. Given that there are no more \mathbf{x} vectors after time T , what is the probability that there are no more \mathbf{x} vectors after time T ? Well, 1, obviously.

The Backward Algorithm

Now let's use the definition $\beta_t(i) \equiv \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\}$, and see how we can compute that.

① **Initialize:**

$$\beta_T(i) = 1, \quad 1 \leq i \leq N$$

② **Iterate:**

$$\begin{aligned} \beta_t(i) &= \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\} \\ &= \sum_{j=1}^N \Pr\{q_{t+1} = j | q_t = i\} \Pr\{\mathbf{x}_{t+1} | q_{t+1} = j\} \Pr\{\mathbf{x}_{t+2}, \dots, \mathbf{x}_T | q_{t+1} = j\} \\ &= \sum_{j=1}^N a_{i,j} b_j(\mathbf{x}_{t+1}) \beta_{t+1}(j) \end{aligned}$$

The Backward Algorithm

Now let's use the definition $\beta_t(i) \equiv \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\}$, and see how we can compute that.

① **Initialize:**

$$\beta_T(i) = 1, \quad 1 \leq i \leq N$$

② **Iterate:**

$$\beta_t(i) = \sum_{j=1}^N a_{i,j} b_j(\mathbf{x}_{t+1}) \beta_{t+1}(j), \quad 1 \leq i \leq N, \quad 1 \leq t \leq T-1$$

③ **Terminate:**

$$\Pr\{\mathbf{X} | \Lambda\} = \sum_{i=1}^N \pi_i b_i(\mathbf{x}_1) \beta_1(i)$$

The Backward Algorithm: Computational Complexity

Most of the computational complexity is in this step:

- **Iterate:**

$$\beta_t(i) = \sum_{j=1}^N a_{i,j} b_j(\mathbf{x}_{t+1}) \beta_{t+1}(j), \quad 1 \leq i \leq N, \quad 2 \leq t \leq T$$

Its complexity is:

- For each of $T - 1$ time steps, $1 \leq t \leq T - 1, \dots$
- we need to calculate N different beta-variables, $\beta_t(i)$, for $1 \leq i \leq N, \dots$
- each of which requires a summation with N terms.

So the total complexity is $\mathcal{O}\{TN^2\}$.

Use Bayes' Rule

The segmentation probability is then

$$\begin{aligned}
 \gamma_t(i) &= \frac{\Pr\{\mathbf{X}, q_t = i | \Lambda\}}{\sum_{k=1}^N \Pr\{\mathbf{X}, q_t = k | \Lambda\}} \\
 &= \frac{\Pr\{\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = i | \Lambda\} \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\}}{\sum_{k=1}^N \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = k | \Lambda\} \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = k, \Lambda\}} \\
 &= \frac{\alpha_t(i) \beta_t(i)}{\sum_{k=1}^N \alpha_t(k) \beta_t(k)}
 \end{aligned}$$

What About State Sequences?

- Notice a problem: $\gamma_t(i)$ only tells us about one frame at a time! It doesn't tell us anything about the probability of a sequence of states, covering a sequence of frames!
- ... but we can extend the same reasoning to cover two or more consecutive frames. For example, let's define:

$$\xi_t(i, j) = \Pr\{q_t = i, q_{t+1} = j | \mathbf{X}, \Lambda\}$$

- We can solve for $\xi_t(i, j)$ using the same reasoning that we used for $\gamma_t(i)$! The result is

$$\xi_t(i, j) = \frac{\alpha_t(i) a_{ij} b_j(\mathbf{x}_{t+1}) \beta_{t=1}(j)}{\sum_{k=1}^N \alpha_t(k) a_{k,\ell} b_\ell(\mathbf{x}_{t+1}) \beta_{t=1}(\ell)}$$

Segmentation: The Backward Algorithm

In summary, we now have three new probabilities, all of which can be computed in $\mathcal{O}\{TN^2\}$ time:

1 The Backward Probability:

$$\beta_t(i) = \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\}$$

2 The State Posterior:

$$\gamma_t(i) = \Pr\{q_t = i | X, \Lambda\} = \frac{\alpha_t(i)\beta_t(i)}{\sum_{k=1}^N \alpha_t(k)\beta_t(k)}$$

3 The Segment Posterior:

$$\begin{aligned} \xi_t(i, j) &= \Pr\{q_t = i, q_{t+1} = j | X, \Lambda\} \\ &= \frac{\alpha_t(i)a_{i,j}b_j(\mathbf{x}_{t+1})\beta_{t+1}(j)}{\sum_{k=1}^N \sum_{\ell=1}^N \alpha_t(k)a_{k\ell}b_\ell(\mathbf{x}_{t+1})\beta_{t+1}(\ell)} \end{aligned}$$

Outline

- 1 Hidden Markov Models
- 2 Recognition: the Forward Algorithm
- 3 Segmentation: the Backward Algorithm
- 4 Numerical Example**
- 5 Summary
- 6 Written Example

Example: Gumball Machines



"Gumball machines in a Diner at Dallas, Texas, in 2008," Andreas Praefcke, public domain image.

Example: Gumball Machines

- **Observation Probabilities:** Suppose we have two gumball machines, $q = 1$ and $q = 2$. Machine #1 contains 60% Grapefruit gumballs, 40% Apple gumballs. Machine #2 contains 90% Apple, 10% Grapefruit.

$$b_1(x) = \begin{cases} 0.4 & x = A \\ 0.6 & x = G \end{cases}, \quad b_2(x) = \begin{cases} 0.9 & x = A \\ 0.1 & x = G \end{cases}$$

- **Initial State Probabilities:** My friend George flips a coin to decide which machine to use first.

$$\pi_i = 0.5, \quad i \in \{1, 2\}$$

- **Transition Probabilities:** After he's used a machine, George flips two coins, and he only changes machines if both coins come up heads.

$$a_{i,j} = \begin{cases} 0.75 & i = j \\ 0.25 & i \neq j \end{cases}$$

The Forward Algorithm: $t = 1$

Remember, the observation sequence is $X = (A, G, A)$.

$$\begin{aligned}\alpha_1(i) &= \pi_i b_1(i) \\ &= \begin{cases} (0.5)(0.4) = 0.2 & i = 1 \\ (0.5)(0.9) = 0.45 & i = 2 \end{cases}\end{aligned}$$

The Forward Algorithm: $t = 2$

Remember, the observation sequence is $X = (A, G, A)$.

$$\begin{aligned}
 \alpha_2(j) &= \sum_{i=1}^2 \alpha_1(i) a_{ij} b_j(x_2) \\
 &= \begin{cases} \alpha_1(1) a_{11} b_1(x_2) + \alpha_1(2) a_{21} b_1(x_2) & j = 1 \\ \alpha_1(1) a_{12} b_2(x_2) + \alpha_1(2) a_{22} b_2(x_2) & j = 2 \end{cases} \\
 &= \begin{cases} (0.2)(0.75)(0.6) + (0.45)(0.25)(0.6) = 0.04125 & j = 1 \\ (0.2)(0.25)(0.1) + (0.45)(0.75)(0.1) = 0.03875 & j = 2 \end{cases}
 \end{aligned}$$

The Backward Algorithm: $t = 3$

The backward algorithm always starts out with $\beta_T(i) = 1!$

$$\beta_3(i) = 1, \quad i \in \{1, 2\}$$

The Backward Algorithm: $t = 2$

Remember, the observation sequence is $X = (A, G, A)$.

$$\begin{aligned}
 \beta_2(i) &= \sum_{j=1}^2 a_{i,j} b_j(x_3) \beta_3(j) \\
 &= \begin{cases} a_{11} b_1(x_3) + a_{12} b_2(x_3) & i = 1 \\ a_{21} b_1(x_3) + a_{22} b_2(x_3) & i = 2 \end{cases} \\
 &= \begin{cases} (0.75)(0.4) + (0.25)(0.9) = 0.525 & j = 1 \\ (0.25)(0.4) + (0.75)(0.9) = 0.775 & j = 2 \end{cases}
 \end{aligned}$$

The Solution to the Puzzle

Given the observation sequence is $X = (A, G, A)$, the posterior state probability is

$$\begin{aligned} \gamma_2(i) &= \frac{\alpha_2(i)\beta_2(i)}{\sum_{k=1}^2 \alpha_2(k)\beta_2(k)} \\ &= \begin{cases} \frac{(0.04125)(0.525)}{(0.04125)(0.525)+(0.03875)(0.775)} = 0.42 & i = 1 \\ \frac{(0.03875)(0.775)}{(0.04125)(0.525)+(0.03875)(0.775)} = 0.58 & i = 2 \end{cases} \end{aligned}$$

So I should dismantle gumball machine #2, hoping to find my rare 1867 silver dollar. Good luck!

Outline

- ① Hidden Markov Models
- ② Recognition: the Forward Algorithm
- ③ Segmentation: the Backward Algorithm
- ④ Numerical Example
- ⑤ Summary**
- ⑥ Written Example

The Forward Algorithm

Definition: $\alpha_t(i) \equiv \Pr\{\mathbf{x}_1, \dots, \mathbf{x}_t, q_t = i | \Lambda\}$. Computation:

① **Initialize:**

$$\alpha_1(i) = \pi_i b_i(\mathbf{x}_1), \quad 1 \leq i \leq N$$

② **Iterate:**

$$\alpha_t(j) = \sum_{i=1}^N \alpha_{t-1}(i) a_{ij} b_j(\mathbf{x}_t), \quad 1 \leq j \leq N, \quad 2 \leq t \leq T$$

③ **Terminate:**

$$\Pr\{\mathbf{X} | \Lambda\} = \sum_{i=1}^N \alpha_T(i)$$

The Backward Algorithm

Definition: $\beta_t(i) \equiv \Pr\{\mathbf{x}_{t+1}, \dots, \mathbf{x}_T | q_t = i, \Lambda\}$. Computation:

① **Initialize:**

$$\beta_T(i) = 1, \quad 1 \leq i \leq N$$


② **Iterate:**

$$\beta_t(i) = \sum_{j=1}^N a_{i,j} b_j(\mathbf{x}_{t+1}) \beta_{t+1}(j), \quad 1 \leq i \leq N, \quad 1 \leq t \leq T-1$$

③ **Terminate:**

$$\Pr\{\mathbf{X} | \Lambda\} = \sum_{i=1}^N \pi_i b_i(\mathbf{x}_1) \beta_1(i)$$

Hidden Markov Model



- 1 Start in state $q_t = i$ with pmf π_i .
- 2 Generate an observation, \mathbf{x} , with pdf $b_i(\mathbf{x})$.
- 3 Transition to a new state, $q_{t+1} = j$, according to pmf a_{ij} .
- 4 Repeat.

Outline

- 1 Hidden Markov Models
- 2 Recognition: the Forward Algorithm
- 3 Segmentation: the Backward Algorithm
- 4 Numerical Example
- 5 Summary
- 6 Written Example**

Written Example

Joe's magic shop opens at random, and closes at random. To be more specific, if it's currently closed, the probability that it will open any time in the next hour is 10%; if it's currently open, the probability that it will close any time in the next hour is 10%. The shop is in a busy part of town; when the shop is open, the area gets even busier. If the shop is closed, the area is noisy with a probability of 40%. If it's open, the area is noisy with a probability of 70%.

At 1:00, you notice that the area is noisy, so you go to check; unfortunately, the shop is closed. At 2:00, the area is still noisy, but you decide that it's unlikely that the shop has opened in just one hour. At 3:00 the area is still noisy, and at 4:00, and at 5:00. How many hours in a row does the area need to be noisy before you decide that, with a probability of greater than 50%, the shop is open?