
CSE 150. Homework 6

Out: *Tue Aug 26*

Due: *Tue Sept 2*

Supplementary reading:

- Russell & Norvig, Chapter 15.
- L. R. Rabiner (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* 77(2):257–286.

6.1 Viterbi algorithm

In this problem, you will decode an English phrase from a long sequence of non-text observations. To do so, you will implement the same algorithm used in modern engines for automatic speech recognition. In a speech recognizer, these observations would be derived from real-valued measurements of acoustic waveforms. Here, for simplicity, the observations only take on binary values, but the high-level concepts are the same.

Consider a discrete HMM with $n = 26$ hidden states $S_t \in \{1, 2, \dots, 26\}$ and binary observations $O_t \in \{0, 1\}$. Download the ASCII data files from the course web site for this assignment. These files contain parameter values for the initial state distribution $\pi_i = P(S_1 = i)$, the transition matrix $a_{ij} = P(S_{t+1} = j | S_t = i)$, and the emission matrix $b_{ik} = P(O_t = k | S_t = i)$, as well as a long bit sequence of $T = 48000$ observations.

Use the Viterbi algorithm to compute the most probable sequence of hidden states conditioned on this particular sequence of observations. As always, you may program in the language of your choice. Turn in the following:


- a hard-copy print-out of your source code**
- a plot of the most likely sequence of hidden states versus time.**

To check your answer: suppose that the hidden states $\{1, 2, \dots, 26\}$ represent the letters $\{a, b, \dots, z\}$ of the English alphabet. If you have implemented the Viterbi algorithm correctly, the most probable sequence of hidden states (ignoring repeated elements) will reveal a challenge many people conquered this summer.

6.2 Conditional independence

Consider the hidden Markov model (HMM) shown below, with hidden states S_t and observations O_t for times $t \in \{1, 2, \dots, T\}$. State whether the following statements of conditional independence are true or false.

- _____ $P(S_t|S_{t-1}) = P(S_t|S_{t-1}, O_t)$
- _____ $P(S_t|S_{t-1}) = P(S_t|S_{t-1}, S_{t+1})$
- _____ $P(S_t|S_{t-1}) = P(S_t|S_{t-1}, O_{t-1})$
- _____ $P(S_t|O_{t-1}) = P(S_t|O_1, O_2, \dots, O_{t-1})$
- _____ $P(O_t|S_{t-1}) = P(O_t|S_{t-1}, O_{t-1})$
- _____ $P(O_t|O_{t-1}) = P(O_t|O_1, O_2, \dots, O_{t-1})$
- _____ $P(O_1, O_2, \dots, O_T) = \prod_{t=1}^T P(O_t|O_1, \dots, O_{t-1})$
- _____ $P(S_2, S_3, \dots, S_T|S_1) = \prod_{t=2}^T P(S_t|S_{t-1})$
- _____ $P(O_1, O_2, \dots, O_T|S_1, S_2, \dots, S_T) = \prod_{t=1}^T P(O_t|S_t)$
- _____ $P(S_1, S_2, \dots, S_T, O_1, O_2, \dots, O_T) = \prod_{t=1}^T P(S_t, O_t)$



6.3 Inference in HMMs

Consider a discrete HMM with hidden states S_t , observations O_t , transition matrix $a_{ij} = P(S_{t+1} = j | S_t = i)$ and emission matrix $b_{ik} = P(O_t = k | S_t = i)$. In class, we defined the forward-backward probabilities:

$$\begin{aligned}\alpha_{it} &= P(o_1, o_2, \dots, o_t, S_t = i), \\ \beta_{it} &= P(o_{t+1}, o_{t+2}, \dots, o_T | S_t = i),\end{aligned}$$

for a particular observation sequence $\{o_1, o_2, \dots, o_T\}$ of length T . In terms of these probabilities, which you may assume to be given, as well as the transition and emission matrices of the HMM, show how to (efficiently) compute the following posterior probabilities:

- (a) $P(S_{t+1} = j | S_t = i, o_1, o_2, \dots, o_T)$
- (b) $P(S_t = i | S_{t+1} = j, o_1, o_2, \dots, o_T)$
- (c) $P(S_{t-1} = i, S_t = j, S_{t+1} = k | o_1, o_2, \dots, o_T)$

In all these problems, you may assume that $t > 1$ and $t < T$; in particular, you are *not* asked to consider the boundary cases.

6.4 Policy improvement

Consider the Markov decision process (MDP) with two states $s \in \{0, 1\}$, two actions $a \in \{0, 1\}$, discount factor $\gamma = \frac{1}{2}$, and rewards and transition matrices as shown below:

s	$R(s)$
0	5
1	2

s	s'	$P(s' s, a=0)$
0	0	$\frac{1}{2}$
0	1	$\frac{1}{2}$
1	0	$\frac{1}{2}$
1	1	$\frac{1}{2}$

s	s'	$P(s' s, a=1)$
0	0	$\frac{3}{4}$
0	1	$\frac{1}{4}$
1	0	$\frac{1}{4}$
1	1	$\frac{3}{4}$

- (a) Consider the policy π that chooses the action $a = 0$ in each state. For this policy, solve the linear system of Bellman equations (by hand) to compute the state-value function $V^\pi(s)$ for $s \in \{0, 1\}$. Your answers should complete the following table. **Show your work for full credit.**

s	$\pi(s)$	$V^\pi(s)$
0	0	
1	0	

- (b) Compute the greedy policy $\pi'(s)$ with respect to the state-value function $V^\pi(s)$ from part (a). Your answers should complete the following table. **Show your work for full credit.**

s	$\pi(s)$	$\pi'(s)$
0	0	
1	0	