10-701/15-781 Machine Learning, Fall 2003

Homework 7
Out: Nov 25, 2003 **Due: start of class Dec 4, 2003**

If you have questions, please contact Ning Hu <ninghu+781@cs.cmu.edu>.

Note: The goal of this assignment is to help you prepare for the final exam.

1. Gaussian Mixture Models and K-means

(a) You are given a Gaussian mixture model, and all its class probabilities and Gaussian mean locations are learned using EM, but the covariance matrices are forced to be the identity matrices for each class. Rather than using a mixture of Gaussians, you evaluate the probability of each of the K classes given the datapoint and take the the cluster with the highest probability to be the cluster that produced the point. Is this equivalent to doing using a K-means model?

No. GMM uses prior class probabilities $P(w_i)$ for evaluating the probability of each of the K classes given the datapoint, while K-means does not consider prior class probabilities (or it can be deemed as all the prior class probabilities are equal.)

(b) Suppose you've done K-means and your K is equal to your number of data points with each cluster defined by a single datapoint. Say that you classify test data points as part of the cluster that they would belong to according to your distance metric. What model is this equivalent to?

1-Nearest Neighbour

(c) Run K-means manually for the following dataset. Circles are data points and squares are the initial cluster centers. Draw the cluster centers and the decision boundaries that define each cluster. Use as many pictures as you need until convergence.

Note: Execute the algorithm such that if a mean has no points assigned to it, it stays where it is for that iteration.

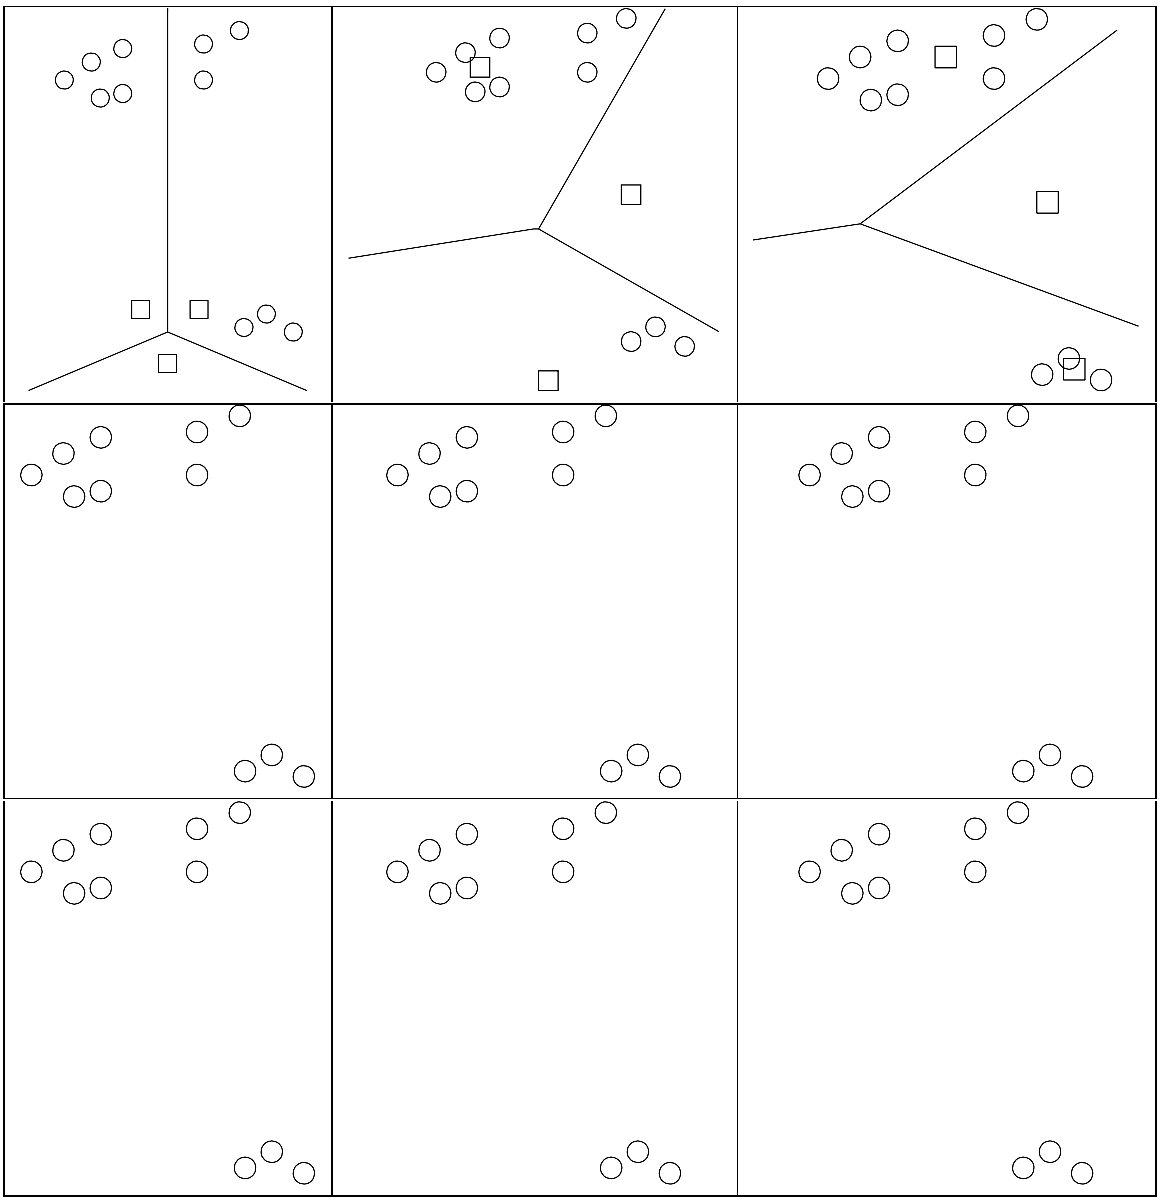

(d) Now run a Gaussian mixture model of three Gaussians on the same dataset. The initial cluster centers are the same as for the k-means problem, and those dashed-line ellipses represent the size and shape of the initial covariance matrices. Assume that the model puts no restrictions on the form of the covariance matrices and that EM updates both the means and covariance matrices. Draw (approximately) the cluster centers and the size/shape of the covariance matrices of the final converged GMM.

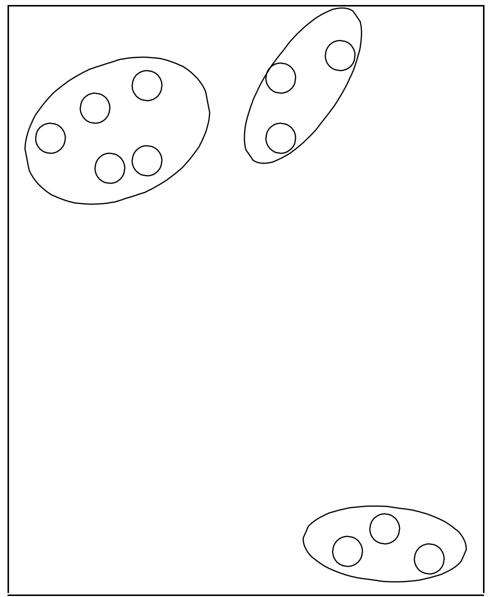

(e) Is the classification given by the mixture model the same as the classification given by k-means? Why or why not?

No. In the mixture model, soft associations (through the weights) are made with every data point by every gaussian, so it can't happen that the cluster center isn't associated with any data point. It was also ok to point out that the algorithms use different distance metrics, or that mixture models with full covariance matrices allow more flexibility in fitting a cluster. It wasn't enough to state the result of each algorithm on the example data (that doesn't say anything about why the result was like that).

2. Hidden Markov Models

(a) Imagine an HMM with two observation symbols: X and Y. Assuming the traditional notation of Hidden Markov Models, consider the following probability.

$$P(q_{t+2} = s_i \wedge q_{t+3} = s_j \wedge O_{t+3} = X \mid q_t = s_k \wedge q_{t+1} = s_m \wedge O_t = X)$$

Give an expression for this value in terms of appropriate elements of the transition and observation matrices, A and B.

$$a_{mi} a_{ij} b_j(X)$$

(b) Foxy lives a simple life. Some days she's Angry and some days she's Happy. But she hides her emotional state, and so all you can observe is whether she smiles, frowns or yells. We start on day 1 in the Happy state, and there is one transition per day.

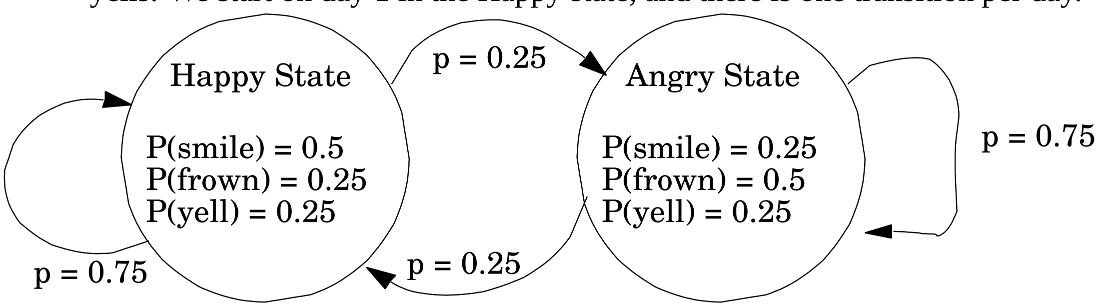

Definitions:
q_t = state on day t.
O_t = observation on day t.

i. What is $P(q_2 = Angry)$?

$$0.25$$

ii. What is $P(O_2 = smile)$?

$$\begin{aligned} P(O_2 = smile) &= P(O_2 = smile \wedge q_2 = Happy) + (O_2 = smile \wedge q_2 = Angry) \\ &= \frac{7}{16} \end{aligned}$$

iii. What is $P(q_2 = Angry | O_2 = smile)$?

By using Bayes rule and previous results,

$$P(q_2 = Angry | O_2 = smile) = \frac{1}{7} \simeq 0.143$$

iv. What is $P(O_{53} = yell)$?

0.25
As $P(Q_t = yell|q_t = Happy) = P(Q_t = yell|q_t = Angry) = 0.25$, the probability of observing $yell$ is always 0.25 no matter which state it is in at any time.

v. Write $\phi_t = P(q_t = Happy)$. Note that $\phi_1 = 1$. Then ϕ_{t+1} can be defined inductively from ϕ_t by an expression $\phi_{t+1} = X + Y\phi_t$. Give the numerical values of X and Y.

$$\phi_{t+1} = (1 - \phi) \times 0.25 + \phi \times 0.75 \quad = \quad \frac{1}{4} + \frac{1}{2}\phi_t$$

So $X = \frac{1}{4}$ and $Y = \frac{1}{2}$

vi. Assume that $O_1 = Frown$, $O_2 = Frown$, $O_3 = Frown$, $O_4 = Frown$, and $O_5 = Frown$. What is the most likely sequence of states?

HAAAA

3. Markov Decision Process

The path to enlightenment is difficult. In each of the states 1 (Normal), 3 (Hungry) and 5 (Starving), you have two possible actions: either Eat or Meditate. In the other states 2 (Over-Eating), 4 (Happy), 6 (Buddha), and 7(Dead) there is only one possible action. Notice that this MDP is almost completely deterministic, the only stochasticity appears when we are in state 5 (Starving) and choose to Eat. Then, with some probability 0.6, you are able to attain enlightenment and reach state 6 (Buddha), but with probability 0.4 your ascetic practice has gone too far and you die (state 7). Let $\gamma < 1$ be the discount factor, which indicates how much you evaluate delayed rewards with respect to immediate ones. The rewards are: $r(1) = 0$, $r(2) = r(3) = -1$, $r(4) = 10$, $r(5) = -10$, $r(6) = 100$, and $r(7) = -100$.

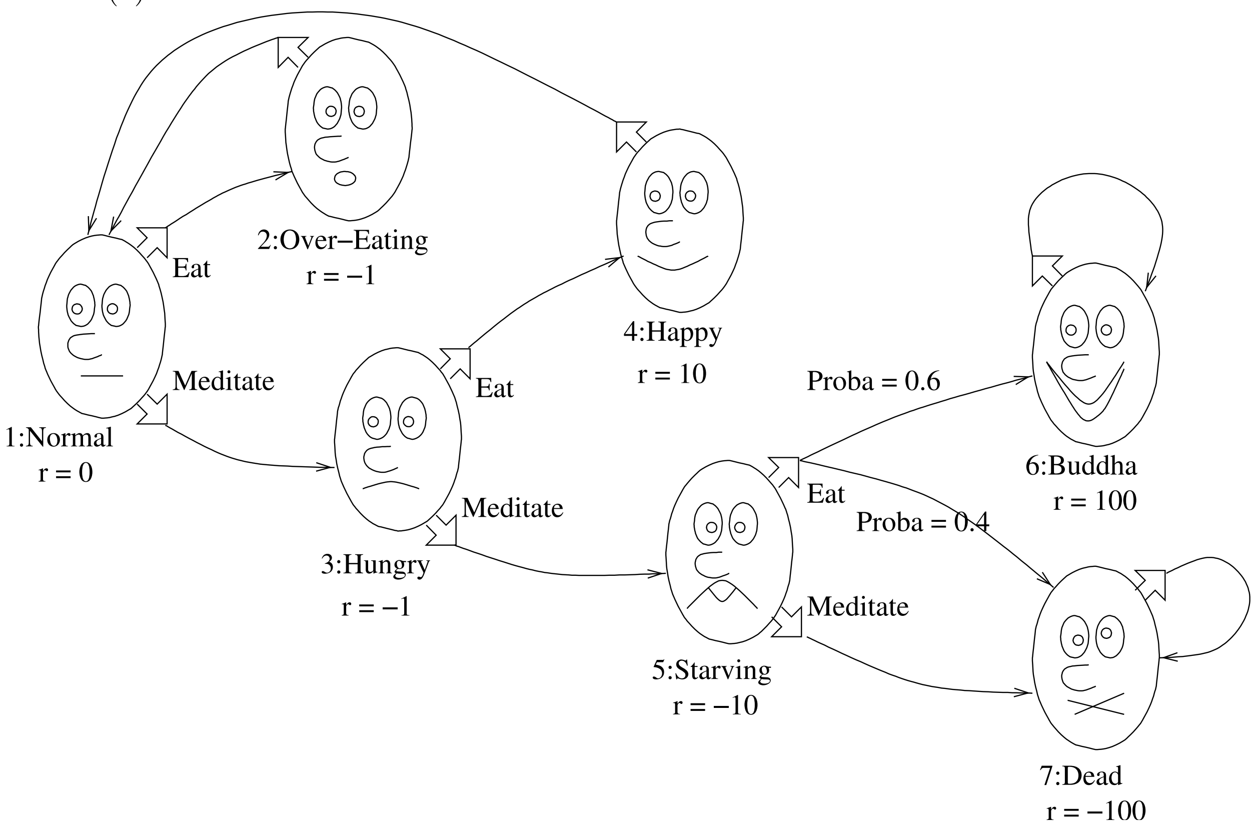

(a) We write: $J^*(i)$ = discounted sum of future rewards starting from state i and using the optimal policy π^*. Express $J^*(6)$ and $J^*(7)$ as a function of γ. (Do not use J^* values on the righthand side of your answer).

$J^*(6) = \frac{100}{1-\gamma}$

$J^*(7) = -\frac{100}{1-\gamma}$

(b) What is the optimal action in state 5?
$\pi^*(5) = Eat$

(c) Express $J^*(5)$ as a function of γ (Do not use J^* values on the righthand side of your answer).
$J^*(5) = -10 + \frac{20\gamma}{1-\gamma}$

(d) What is the optimal action in state 1?
$\pi^*(1) = Meditate$

(e) Write Bellman's equation for states 1, 3 and 4.
$J^*(1) = \gamma J^*(3)$

$J^*(3) = -1 + \gamma max[J^*(4), J^*(5)]$

$J^*(4) = 10 + \gamma J^*(1)$

(f) We consider an hedonistic person seeking nearly immediate rewards, for which $\gamma = 0.1$. What is the optimal action in state 3? And what are the values $J^*(1)$, $J^*(2)$, $J^*(3)$, $J^*(4)$, and $J^*(5)$?
$\pi^*(3) = Eat$

$J^*(1) = 0$

$J^*(2) = -1$

$J^*(3) = 0$

$J^*(4) = 10$

$J^*(5) = -\frac{70}{9} \simeq -7.78$

(g) Now we consider a self-disciplined person who is interested in delayed rewards, for which $\gamma = 0.9$. What is the optimal action in state 3?
$\pi^*(3) = Meditate$

(h) Suppose that $\gamma = 0.9999$, which method would you recommend to use for evaluating the J-values for a given policy: Matrix Inversion or Value Iteration?

Matrix Inversion