10-701 Machine Learning, Spring 2011: Homework 4

Due: Tuesday March 1st at the begining of the class

Instructions There are two questions on this assignment. Please submit your writeup as two separate sets of pages according to questions, with your name and userid on each set.

1 EM and Bayes Nets [Yi Zhang, 40 points]

In the class we learned the famous expectation-maximization algorithm, which is widely used to estimate model parameters from partially observed data. In this question, we will use **[EM]** to denote the lecture slides on Feb 17th ("Graphical models 3: EM") from the course website. Since each page of [EM] contains two slides, we will use, e.g., the upper/lower slide on page 7 of [EM], as a pointer to the first/second slide on a specific page of lecture slides.

Given the set of observed variables $\mathbf{X}$ and the set of unobserved variables $\mathbf{Z}$, as shown by the lower slide on page 6 of [EM], the EM algorithm for estimating model parameters θ from training examples is the following procedure:

1. E-Step: for each example k, use $\mathbf{X}_k$ and the current θ to calculate $P(\mathbf{Z}_k|\mathbf{X}_k, \theta)$.
2. M-Step: re-estimate θ as $\theta \leftarrow \text{argmax}_{\theta'} E_{P(\mathbf{Z}|\mathbf{X},\theta)}[\log P(\mathbf{X}, \mathbf{Z}|\theta')]$ using the training examples.
3. Iterate until convergence.

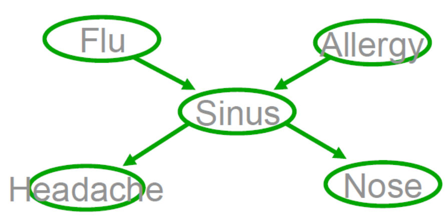

Figure 1: The Bayes net for the flu-allergy example

Now consider the familiar flu-allergy Bayes net as figure 1. Suppose F, A, H, N are observed (i.e., $\mathbf{X} = \{F, A, H, N\}$) and S is unobserved (i.e., $\mathbf{Z} = \{S\}$). Given a set of training examples $\{(f_k, a_k, h_k, n_k)\}_{k=1}^K$, the EM algorithm should proceed as follows:

1. E-Step: for each example k, use (f_k, a_k, h_k, n_k) and the current θ to calculate $P(S_k|f_k, a_k, h_k, n_k, \theta)$.
2. M-Step: re-estimate θ as $\theta \leftarrow \text{argmax}_{\theta'} \sum_{k=1}^K E_{P(S_k|f_k,a_k,h_k,n_k,\theta)}[\log P(f_k, a_k, h_k, n_k, S_k|\theta')]$.
3. Iterate until convergence.

1.1 A simplified EM algorithm

As we saw in the lower slide on page 8 of [EM], the EM algorithm can be simplified when the unobserved variable is of boolean values. In this case, the *simplified* EM algorithm is:

1. E-Step: for each example k, use $\mathbf{X}_k$ and the current θ to calculate the expected value $E(\mathbf{Z}_k|\mathbf{X}_k, \theta)$.
2. M-Step: re-estimate θ similarly to MLE on observed data, but replacing each count involving the unobserved variable by its expected count.
3. Iterate until convergence.

As a result, as shown by the upper slide on page 8 of [EM], the *simplified* EM algorithm for the flu-allergy Bayes net with a set of training examples $\{(f_k, a_k, h_k, n_k)\}_{k=1}^{K}$ is:

1. E-Step: for each example k, use (f_k, a_k, h_k, n_k) and the current θ to calculate the expected value $E(S_k) = E(S_k|f_k, a_k, h_k, n_k, \theta)$.
2. M-Step: re-estimate θ as MLE with expected counts, e.g., $\theta_{s|ij} \leftarrow \frac{\sum_{k=1}^{K} \delta(f_k=i, a_k=j) E(S_k)}{\sum_{k=1}^{K} \delta(f_k=i, a_k=j)}$. Recall that $\theta_{s|ij}$ stands for $P(S = 1|F = i, A = j)$.
3. Iterate until convergence.

Question [15 pts]: Given that S in the flu-allergy Bayes net is a boolean variable, prove that the simplified EM algorithm for estimating $\theta_{s|ij}$ is indeed equivalent to the standard EM algorithm. In other words, show that $\theta_{s|ij} \leftarrow \frac{\sum_{k=1}^{K} \delta(f_k=i, a_k=j) E(S_k)}{\sum_{k=1}^{K} \delta(f_k=i, a_k=j)}$ in the M-Step of the simplified EM indeed gives the corresponding parameter in $\theta \leftarrow \text{argmax}_{\theta'} \sum_{k=1}^{K} E_{P(S_k|f_k,a_k,h_k,n_k,\theta)}[\log P(f_k, a_k, h_k, n_k, S_k|\theta')]$ from the M-Step of the standard EM algorithm.

Note: you are required to prove this only for $\theta_{s|ij}$, and no need to worry about other parameters in θ.

Hint: the lower slide on page 4 of [EM] should be helpful. The slide shows how to derive the parameter $\theta_{s|ij}$ in $\theta \leftarrow \text{argmax}_{\theta'} \sum_{k=1}^{K} [\log P(f_k, a_k, h_k, n_k, s_k|\theta')]$ when all variables are observed.

1.2 Simulating the simplified EM algorithm

We are given the following $K = 8$ training examples as in Table 1, where only two examples contain unobserved values, namely, H_7 and N_8. We like to simulate a few steps of the simplified EM algorithm by hand. Note that this is not a programming question.

Table 1: Training examples for the flu-allergy Bayes net

k	F	A	S	H	N
$k = 1$	1	0	1	1	1
$k = 2$	0	1	1	1	0
$k = 3$	1	1	1	1	1
$k = 4$	0	0	0	0	0
$k = 5$	0	0	0	1	0
$k = 6$	0	0	0	0	1
$k = 7$	1	1	1	**?**	1
$k = 8$	1	1	1	1	**?**

Question A [7 pts]: given that all variables are boolean, how many parameters we need to estimate in the flu-allergy Bayes net? Also, list all the parameters we need to estimate, e.g., one parameter will be $\theta_{s|11}$, which stands for $P(S = 1|F = 1, A = 1)$.

Question B [6 pts]: Now we like to simulate the first E-step of the simplified EM algorithm. Before we start, we initialize all the parameters as 0.5, and then proceed to execute the E-step. What are the expectations we need to calculate in this E-step? Also, list the actual values for these expectations. (Note that only two examples contains unobserved variables).

Question C [6 pts]: Now we like to simulate the first M-step. List the estimated values of all model parameters we obtain in this M-step. (Note that we use the expected count only when the variable is unobserved in an example).

Question D [6 pts]: Last, let's simulate the second E-step. List the actual values for all the expectations we calculate in this E-step.

2 Midterm review questions [Tom Mitchell, 60 points]

This question contains some short questions adapted from previous midterm exams – a good way to review for our own on March 3.

1. Give a *one sentence* reason why [12 pts]:
   - we might prefer Decision Tree learning over Logistic Regression for a particular learning task.
   - we might prefer Logistic Regression over Naive Bayes for a particular learning task.
   - we choose parameters that minimize the sum of squared training errors in Linear Regression.
2. Suppose we train several classifiers to learn $f : X \rightarrow Y$, where X is the feature vector $X =< X_1, X_2, X_3 >$. Which of the following classifiers contains sufficient information to allow calculating $P(X_1, X_2, X_3, Y)$? If you answer yes, give a brief sketch of how. If you answer no, state what is missing. [12 pts]
   - Gaussian Naive Bayes
   - Logistic Regression
   - Linear Regression
3. True or False? If true, give a 1-2 sentence explanation. If false, a counterexample. *Your answer must fit into the space below the question* [12 pts].
   - As the number of data points grows to infinity, the MLE estimate of a parameter approaches the MAP estimate, for all possible priors.
   - The depth of a learned decision tree can be larger than the number of training examples used to create the tree.
   - There is *no* training data set for which a decision tree learner and logistic regression will output the same decision boundary.
4. In class we defined *conditional independence* by saying that random variable X is conditionally independent of Y given Z if and only if:
   $$P(X|Y,Z) = P(X|Z) \tag{1}$$
   Prove that if $P(XY|Z) = P(X|Z)P(Y|Z)$, then X is conditionally independent of Y given Z (*hint: this is a two-line proof*) [4 pts].
5. Consider the Bayes network below, defined over four Boolean variables [20 pts].

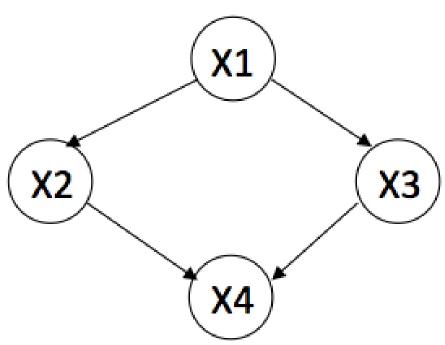

- How many parameters are needed to define $P(X1, X2, X3, X4)$ for this Bayes Net?
- Give the formula that calculates $P(X1 = 1, X2 = 0, X3 = 1, X4 = 0)$ using only the Bayes net parameters. Use notation like $P(X1 = 0|X2 = 1, X4 = 0)$ to refer to each Bayes net parameter you use in your formula.
- Give the formula that calculates $P(X1 = 1, X4 = 0)$ using only the Bayes net parameters.
- Give the formula that calculates $P(X2 = 1|X3 = 0)$ using only the Bayes net parameters.