CS5560: Exam3

9am-12pm, 23-11-2018 (Fri).

Note: Please write **relevant**, **precise**, and **concise** answers. Please write **legibly**. Do not forget to write your ROLL NO. on the answer sheet.

1. It is proposed to evaluate the popularity of **CSE**@IITH faculty from recordings of "like" and "dislike" (binary votes) from students in each of the course-offerings they make. However, for reasons of privacy and related policies, only the difference between the number of "like" votes and the number of "dislike" votes in a random subset of course-offerings for each faculty is made public. Ideally, the faculty with higher expected difference (of "like" and "dislike" votes), must be declared to be more popular.

   More formally, let V_i denote the rv that is equal to the number of "like" votes minus the number of "dislike" votes in a random course-offering by the i^{th} faculty. Ideally, faculty i must be declared more popular than faculty j, iff $\mathbb{E}[V_i] > \mathbb{E}[V_j]$. However, these expectations are unknown and the only published data for the i^{th} faculty is $\mathcal{D}_i = \{v_{i1}, \ldots, v_{im_i}\}$, where v_{ik} denotes the number of "like" votes minus the number of "dislike" votes faculty i actually received in course-offering k. Needless to say, to rank faculty as per popularity, one needs to estimate $\mathbb{E}[V_i]\ \forall\ i$ from the data $\mathcal{D}_i\ \forall\ i$.

   Consider the following Hierarchical Bayesian model for solving the above estimation problem: $V_i \sim \mathcal{N}(\mu_i, \sigma^2),\ \mu_i \sim \mathcal{N}(\mu_o, \tau_o^2)\ \forall\ i$. The parameters are μ_i (assume σ^2 is known). And, the hyperparameters are μ_o, τ_o^2. Derive a highly simplified formula for the estimate of hyperparameters using the maximum marginal likelihood (i.e., ML-II estimation, or Empirical Bayes estimation) algorithm.

   [**7.5**Marks]

   Derive a highly simplified formula for the estimate of μ_i using the MAP algorithm.

   [**5**Marks]

   Note that the expressions for these estimates must be written in terms[1] of the published (training) data and σ^2.

[1] Note that if this was not a Hierarchical Bayesian model and : $V_i \sim \mathcal{N}(\mu_i, \sigma^2) \forall\ i$ are the only modeling assumptions, then the answer is trivial $\mathbb{E}[V_i] \equiv \mu_i \approx \frac{\sum_{k=1}^{m_i} v_{ik}}{m_i}$. After the exam, think about situations where the proposed Hierarchical model could be advantageous over this baseline, and post your thoughts on G-Classroom.

2. In lectures it was described how a HMM can be used to solve the problem of converting an audio recording of a legitimate speech utterance of some sentence in a particular language like, say, English, into the corresponding text in that language. We also discussed how such a HMM can be trained in a completely unsupervised fashion using a set of speech utterances only (not tagged with the text written in them).

   Now consider the generalization where the utterance/text can be in any of the three languages English, Hindi, and Telugu[2]. The problem is to now convert a speech utterance into text without knowing the language in which it is was spoken. Extending the discussion in the lecture, it seems a "Mixture of HMMs" is the natural choice for the model.

   More specifically, let $X = (X_1, \ldots, X_T)$ be the rv denoting the sequence of feature vectors corresponding to a (random) speech utterance. Let $Y = (Y_1, \ldots, Y_T)$ be the rv denoting the sequence of phonemes in a (random) sentence. Let L be the rv denoting the language identity. Though it is clear that some unknown joint distribution, $p^*(x, y, l)$, governs the mechanism of speech generation, only samples from X are given for the purpose of training. We are interested in getting a good estimate of $p(l/x)$ and $p(y/x)$ from these samples.

   Consider the following modeling assumptions (in short "Mixture of HMMs"):

   (a) the joint over X, Y, L is modelled using a mixture model with (X, Y) as the observed variables and L as the hidden variable. The number of components is hence three.

   (b) each conditional joint $X, Y/L$ is modelled using a standard HMM, with X as the observed and Y as the hidden variable. Assume that the state-space for each of these three HMMs is the same, and the common number of states is c. Assume that the emission distributions are all modeled using an exponential family model defined by the same sufficient statistics, $\phi : \mathcal{X} \mapsto \mathbb{R}$.

   Draw a DAG over the rvs $X_1, \ldots, X_T, Y_1, \ldots, Y_T, L$ that is a perfect[3] IMAP for these modeling assumptions. Clearly write down the independence assumptions that you want to make (if any) across the samples[4].

   [**5**Marks]

   Clearly list down all the parameters in this "mixture of HMMs" model and present rough details of an EM algorithm for estimating these parameters using MLE with the (unsupervised) data.

   [**10**Marks]

[2] not uncommon in Telangana & AP

[3] $\mathcal{G}$ is a perfect IMAP to distribution p iff $I_g(\mathcal{G}) = I(p)$.

[4] If convenient, then you may use a "plate model" to depict independence conditions across samples

Now, instead of MLE, let us assume appropriate priors over the parameters are defined for inference using the Bayesian paradigm. In this case, draw a DAG over all the rvs involved that is a perfect IMAP for these Bayesian modeling assumptions. Also, in this case, list the (conditional) independence across the samples (if any).

[**5**Marks]

3. Consider the graph $\mathcal{G}$: 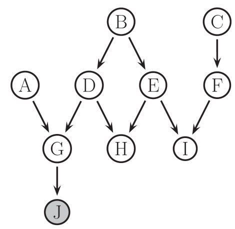

. Assuming the topological ordering is alphabetical (ascending), write down any three[5] (conditional) independence conditions in each of $I_o(\mathcal{G}), I_l(\mathcal{G}), I_g(\mathcal{G})$ i.e., "Ordered Markov property, Local Markov Property, Global Markov Property" (totally six conditions). Choose your six conditions such that one does not imply the other. The conditions in $I_g(\mathcal{G})$ that you write must involve A and use J as the observed (given) variable.

[**5**Marks]

4. Consider jointly Gaussian distributed rvs (X, Y, Z) with covariance matrix $\begin{bmatrix} 3 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 3 \end{bmatrix}$. Let this joint distribution be $p(x, y, z)$. Draw an undirected graph, H, over these rvs, such that $\mathcal{H}$ is a perfect[6] IMAP for p. Also, list all the (conditional) independence conditions that hold in p.

[**7.5**Marks]

5. Consider the graph $\mathcal{H}$: 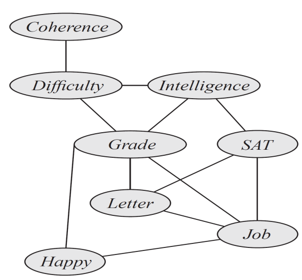

. Specify a valid Markov Network structure for $\mathcal{H}$ that induces the least number of factors.

[**5**Marks]

[5]In case in any category, if there are less than three independence conditions that are true, then mention so and list down all that are true in that category.

[6]$\mathcal{H}$ is a perfect IMAP to distribution p iff $I_g(\mathcal{H}) = I(p)$.