CS5560: Final Exam (EMDS batch...)

ROLL NO.________________

IMPORTANT INSTRUCTIONS:

1. Appropriately fill the blanks while **strictly** adhering to the instructions provided in the corresponding footnotes.
2. Unsimplified (correct) answers, (correct) answers in format other than that specified in the footnotes (wrong format), illegible writings, and those outside the blanks, will be ignored by the evaluator.
3. Also, please provide **concise** justification for your answer in the box corresponding to the blank. Note that if no box is provided for a blank, then you should not provide any justification.
4. In case a blank is not filled (or incomplete), the justification provided by you will be ignored. Illegible writings, small/tiny fonts (concise does NOT mean shrink font), and those outside the box will be ignored. The box size is an indicative of the amount of justification expected.
5. For blanks with boxes, you will be awarded full marks **if and only if** both your answer in the blank is correct and the justification in its box is precise.
6. It is strongly recommended that you first work out the solutions in rough sheets and then carefully plan and enter the blanks and the justifications. This way you can avoid over-writings etc.
7. Please do remember to write your Roll No. correctly on this sheet (above). Please return this question paper duly filled with your answers to the invigilator. Please do NOT submit your rough sheets.

FILL IN THE BLANKS

1. Of the following methods, the ones that view the parameters of a model as a random variable are B, C [1].

   (A) MLE

   (B) MAP

   (C) Bayesian Inference

   (D) Method of Moments

[1 Mark]

[1] Fill the blank with one or more of the alphabets "A", "B", "C", "D", separated by commas. If none of the choices satisfy the criteria, then write "none".

Technically correct ANS is MLE of covariance "does NOT exist".

2. Consider the training data $\left\{\begin{bmatrix}1\\5\end{bmatrix},\begin{bmatrix}2\\4\end{bmatrix},\begin{bmatrix}3\\3\end{bmatrix},\begin{bmatrix}4\\2\end{bmatrix},\begin{bmatrix}5\\1\end{bmatrix}\right\}$. With Gaussian model, $p_{\mu,\Sigma}(x,y)$, the MLE estimate of the mean is $\hat{\mu} = \begin{bmatrix}3\\3\end{bmatrix}$[2], and covariance is $\hat{\Sigma} = \boxed{\begin{bmatrix}2 & -2\\-2 & 2\end{bmatrix}}$[3]. The simplified expression for the mean of the posterior likelihood $\mathbb{E}_{p_{\hat{\mu},\hat{\Sigma}}}[Y/x] = \underline{6-x}$[4] So I am OK if you write so!

[0.5+1.5+3=5 Marks]

3. Consider the following Hierarchical[5] Bayesian model: $X_i \sim Ber(\theta_i)$, $i = 1,\ldots,n$, and $\theta_i \sim Beta(\alpha_1,\alpha_2)$, $i = 1,\ldots,n$. Assume, X_i is conditionally independent of X_j given $\theta \equiv (\theta_1,\ldots,\theta_n)$ for any $i \neq j$. Assume, X_i is conditionally independent of θ_j given θ_i for any $i \neq j$. Also, assume θ_i is conditionally independent of θ_j given $\alpha \equiv (\alpha_1,\alpha_2)$. Let $\mathcal{D}_i$ denote the set of m_i samples[6] from X_i. Let p_i denote the samples in $\mathcal{D}_i$ that are unity (rest are zero). Then the ML-II estimate of hyper-parameters α is given by

$$(\hat{\alpha}_1,\hat{\alpha}_2) = \text{argmax}_{\alpha_1>0,\alpha_2>0}\left\{\frac{\prod_{i=1}^{n} B(\alpha_1+p_i,\ \alpha_2+m_i-p_i)}{B^n(\alpha_1,\alpha_2)}\right\}^7.$$

The MAP estimate for θ_i (using the ML-II estimate $\hat{\alpha}_1,\hat{\alpha}_2$) is $\underline{\dfrac{p_i+\hat{\alpha}_1-1}{m_i+\hat{\alpha}_1+\hat{\alpha}_2-2}}$[8].

$p(\mathcal{D}) = \int p(\mathcal{D}/\theta)p(\theta)\,d\theta$ $= \prod_{i=1}^{n}\left[\int p(\mathcal{D}_i/\theta_i)\,p(\theta_i)\,d\theta_i\right]$ By assumptions	$p(\mathcal{D}_i/\theta_i) = \theta_i^{p_i}(1-\theta_i)^{m_i-p_i}$ $p(\theta_i) = \theta_i^{\alpha_1-1}(1-\theta_i)^{\alpha_2-1}/B(\alpha_1,\alpha_2)$

[4+2=6 Marks]

4. Consider the (multiple output) linear regression model: $p_W(y/x) \sim \mathcal{N}(W^\top x, I)$, where I denotes the identity matrix of appropriate size. Training data (input-output pairs) is

$$\left\{\left(\begin{bmatrix}1\\0\end{bmatrix},\begin{bmatrix}-1\\-1\end{bmatrix}\right),\left(\begin{bmatrix}1\\0\end{bmatrix},\begin{bmatrix}-1\\-2\end{bmatrix}\right),\left(\begin{bmatrix}1\\0\end{bmatrix},\begin{bmatrix}-2\\-1\end{bmatrix}\right),\left(\begin{bmatrix}0\\1\end{bmatrix},\begin{bmatrix}1\\1\end{bmatrix}\right),\left(\begin{bmatrix}0\\1\end{bmatrix},\begin{bmatrix}1\\2\end{bmatrix}\right),\left(\begin{bmatrix}0\\1\end{bmatrix},\begin{bmatrix}2\\1\end{bmatrix}\right)\right\}.$$

The MLE estimate of W from the above data is: $\underline{\begin{bmatrix}-1.33 & -1.33\\1.33 & 1.33\end{bmatrix}}$[9].

Can solve for each output individually. In lecture, formula was derived as: $\hat{w}_i = (XX^\top)^{-1}Xy_i$

[2] Fill the blank with a number/vector/matrix using decimal representation. If needed, you may round off number(s) to two decimal places.

[3] Fill the blank with a number/vector/matrix using decimal representation. If needed, you may round off number(s) to two decimal places.

[4] Your expression must NOT involve any parameters/variables other than x, and may involve numerical constants.

[5] Assume that both the parameters and the hyper-parameters are viewed as random variables in this Hierarchical Bayesian model.

[6] As usual in Bayesian modeling, assume that the samples are (mutually) conditionally independent given the corresponding parameters.

[7] Fill the blank with an expression that does NOT involve integrals. You may use mathematical symbols for known functions/numerical-constants defined in the lectures. Go home and think about a numerical procedure for solving this optimization problem!

[8] Fill the blank with the analytical solution of the MAP optimization problem. Your expression must be as simplified as possible.

[9] Fill the blank with a number/vector/matrix using decimal representation. If needed, you may round off number(s) to two decimal places.

[3 Marks]

5. Consider a classification problem with three classes: $1, 2, 3$. The class-conditionals for 1^{st}, 2^{nd}, 3^{rd} classes are respectively $\mathcal{N}(1, e^{-2}), \mathcal{N}(0, e^{2}), \mathcal{N}(-1, 1)$. The class priors are proportional to $e^{-2}, 1, e^{2}$ respectively. The resulting Bayes classifier will classify point $x = 0.5$ as class 3 [10].

$$\log p(1/x) \propto -\frac{(x-1)^2}{2e^{-2}} - 1 \overset{0.5}{<} 0$$
$$\log p(2/x) \propto -x^2/2e^{2} - 1 \overset{0.5}{<} 0$$
$$\log p(3/x) \propto -\frac{(x+1)^2}{2} + 2 \overset{0.5}{>} 0$$

[3 Marks]

6. Let X be a discrete random variable taking on c discrete values $v_1, \ldots, v_c$. Let $\mathcal{D}$ be a set of m samples from X. Let m_i be the number of samples in $\mathcal{D}$ that take the i^{th} discrete value ($\sum_{i=1}^{c} m_i = m$). It is proposed to use the Dirichlet-Multinoulli model with prior as the Dirichlet distribution with hyper-parameters as $\alpha_1, \ldots, \alpha_c$. Let $\alpha \equiv \sum_{i=1}^{m} \alpha_i$. Then the simplified expression for the posterior predictive likelihood $p(x = v_i, x' = v_j/\mathcal{D})$ is $(m_i+\alpha_i)(m_j+\alpha_j)/((m+\alpha+1)(m+\alpha))$ [11].

$$p(x=v_i, x'=v_j/D) = \int p(v_i/\theta)\, p(v_j/\theta)\, p(\theta/D)\, d\theta = \int \theta_j^{m_j+\alpha_j}\, \theta_i^{m_i+\alpha_i} \prod_{k \neq i,j} \theta_k^{m_k+\alpha_k-1}\, d\theta / B(m_1+\alpha_1, \ldots, m_c+\alpha_c)$$

[3 Marks]

7. Consider a model belonging to the exponential family defined by the likelihood: $p_w(x) = h(x)e^{w^\top \phi(x) - A(w)}$. The sufficient statistics, ψ, of the conjugate prior are defined by: $\psi(w) \equiv \begin{bmatrix} w \\ A(w) \end{bmatrix}$ [12].
Of the following, the statement(s) that is(are) true regarding the final simplified expression (denote by $\mathcal{E}$) for the posterior predictive likelihood, derived in the lecture, is(are): B, D [13].

(A) $\mathcal{E}$ involves A explicitly.

(B) $\mathcal{E}$ involves h explicitly.

(C) $\mathcal{E}$ involves ϕ explicitly.

(D) $\mathcal{E}$ involves the partition function of the posterior over the parameters explicitly.

[1+1=2 Marks]

8. Consider models $\mathcal{M}_1 \subsetneq \mathcal{M}_2 \subsetneq \ldots \subsetneq \mathcal{M}_n$. Let `Algo1` denote the model selection algorithm that selects the model with the highest likelihood of training data with the likelihood function corresponding to the MLE parameter. In case of a tie, among the equals the model with the highest index is selected by `Algo1`. Let `Algo2` denote the maximum marginal likelihood algorithm for model selection. Of the following, the statements that are necessarily true are: A [14]

[10] Fill the blank with one of the three numbers "1", "2", "3".

[11] Here, x, x' denote two samples "generated" from the Bayesian model. Your simplified expression must ONLY involve m_i, m_j, m, α and perhaps some numerical constants.

[12] Fill the blank with an appropriate mathematical expression involving some or all of h, w, ϕ, A.

[13] Fill the blank with one or more of the alphabets "A", "B", "C", "D", separated by commas. If none of the statements are true, then write "none".

[14] Fill the blank with one or more of the alphabets "A", "B", "C", "D", separated by commas. If none of the statements are true, then write "none".

(A) `Algo1` always selects $\mathcal{M}_n$.

(B) `Algo2` always selects $\mathcal{M}_1$.

(C) `Algo2` always selects $\mathcal{M}_n$.

(D) `Algo2` achieves better accuracy than `Algo1`.

[1 Mark]

9. Among the following statements, the ones that are necessarily true about the MCLE for logistic regression is(are): none [15].

(A) An optimal solution always exists for the maximum conditional likelihood optimization problem with ANY data (i.e., MCLE always exists).

(B) If an optimal solution exits for the maximum conditional likelihood optimization problem, then it is unique (i.e., MCLE is unique).

[1 Mark]

10. Consider the model defined by the likelihood function: $p_a(x) \propto x^{a-1}e^{-x}, x > 0$, parameterized by $a > 0$. The average value of sufficient statistics[16] with this model over the data $\{e^{-2}, e^{-1}, 1, e, e^2\}$ is 0 [17].

$\boxed{\phi(x) = \log(x)}$

[2 Marks]

11. Among the following, the statement(s) that is(are) necessarily true about the (vanilla) method of moments (MM) when employed for parameter estimation in exponential family models is(are): C, D [18].

(A) The moment matching equations are always feasible i.e., there will exist atleast one value of the canonical parameters where the moment matching equations are true.

(B) If the moment matching equations are feasible, then there will be a unique solution for them.

(C) MLE and MM always provide the same parameter estimates.

(D) MM cannot be directly applied for parameter estimation in discriminative models.

[1 Mark]

12. Among the following, the statement(s) that is(are) necessarily true about the MCLE for linear regression is(are): A [19].

[15] Fill the blank with one or more of the alphabets "A", "B" separated by commas. If none of the statements are true, then write "none".

[16] As always, we are interested only in the "minimal" sufficient statistics

[17] Fill the blank with a number/vector/matrix using decimal representation. If needed, you may round off number(s) to two decimal places.

[18] Fill the blank with one or more of the alphabets "A", "B", "C", "D", separated by commas. If none of the statements are true, then write "none".

[19] Fill the blank with one or more of the alphabets "A", "B" separated by commas. If none of the statements are true, then write "none".

(A) An optimal solution always exists for the maximum conditional likelihood optimization problem with ANY data (i.e., MCLE always exists).

(B) If an optimal solution exits for the maximum conditional likelihood optimization problem, then it is unique (i.e., MCLE is unique).

[1 Mark]

13. Let X be a multinoulli random variable taking on values $1, 2, \dots, 10$. A set of thousand samples from X are collected, however, the sample is shown only if it is of value 1 or 2. Else it is marked as "?". In summary, the training dataset ($\mathcal{D}$) is a set of 1000 samples each taking one of the three possible values "1", "2" and "?" (denoting that it is one of $3, \dots, 10$). Then, the statement "the Dirichlet distribution with (hyper) parameters $\alpha_1, \dots, \alpha_{10}$ is[20] a conjugate prior to the likelihood of $\mathcal{D}$ computed using the Multinoulli model (with 10 parameters)" is FALSE[21].

$p(D/\theta) = \theta_1^{m_1}\theta_2^{m_2}(\theta_3+\dots+\theta_{10})^{m-m_1-m_2}$; $p(\theta) = \prod_{i=1}^{10}\theta_i^{\alpha_i}$ | Forms do NOT match!

[2 Marks]

14. Consider the exponential family model defined by $p_v(y) = e^{\sum_{j=1}^n v_j\psi_j(y) - A(v)}$, where $\psi(y) \equiv \begin{bmatrix} \psi_1(y) \\ \vdots \\ \psi_n(y) \end{bmatrix}$ is the sufficient statistics. Now, let $p_w(y/x)$ be the likelihood for the corresponding GLM (discriminative model), where $v_i \equiv w_i^\top \phi(x)\ \forall\ i$ [22]. The simplified expression for the log-conditional-likelihood of the training data $\mathcal{D} = \{(x_1, y_1), \dots, (x_m, y_m)\}$ is $\sum_{i=1}^{m}\left[\sum_{j=1}^{n}(w_j^T\varphi(x_i))\psi_j(y_i) - A(w_1^T\varphi(x_i), \dots, w_c^T\varphi(x_i))\right]$ [23]. The simplified equality condition corresponding to "gradient wrt. w_j (in the earlier expression) is zero" is $\sum_{i=1}^{m} E[\psi_j(Y)/x_i]\,\varphi(x_i) = \sum_{i=1}^{m}\psi_j(y_i)\varphi(x_i)$ [24]

[2+2=4 Marks]

[20] Assume that $\alpha_i > 0\ \forall\ i = 1, \dots, 10$.

[21] Fill with either "True" or "False".

[22] In your textbook, this is refered to as the canonical link function. This is the same we employed in our lecture too.

[23] Fill the blank with an expression involving w_i, ϕ, A, ψ_j. Note that this expression is same as the objective in the MCLE optimization problem.

[24] Your expression must NOT involve A.