IEOR 165 – Midterm
March 14, 2017

Name:	
Overall:	/50

Instructions:

1. Show your intermediate steps.
2. You are allowed a single 8.5x11 inch note sheet.
3. Calculators are allowed.

1	/10
2	/10
3	/10
4	/10
5	/10

1. Suppose we have measured stress and deflection: $s_i = \{10, 11, 12, 13, 14\}$ and $d_i = \{7, 8, 3, 10, 15\}$. Estimate the parameters of a linear model $d = m \cdot s + b$. (10 points)

   Using linear regression:
   $\overline{s} = 12$
   $\overline{d} = 8.6$
   $\overline{sd} = 106.8$
   $\overline{s^2} = 146$
   $\hat{m} = (\overline{sd} - \overline{s} \cdot \overline{d})/(\overline{s^2} - \overline{s}^2) = (106.8 - 12 \cdot 8.6)/(146 - 12^2) = 1.8$
   $\hat{b} = \overline{d} - \hat{m} \cdot \overline{s} = 8.6 - 1.8 \cdot 12 = -13$
   $d = 1.8 \cdot s - 13$

2. Suppose we make iid measurements from a random variable X_i with exponential distribution with rate λ. Recall that the mean of an exponential distribution with rate λ is $1/\lambda$. If our measurements are $X_i = \{3, 1, 4, 9, 2\}$, then estimate λ. (10 points)

   Using method of moments:
   $\mathbb{E}(X) = 1/\lambda \Rightarrow \lambda = 1/\mathbb{E}(X)$
   $\hat{\lambda} = n / \sum_{i=1}^{n} X_i = 5/(3+1+4+9+2) = 0.2632$

3. Suppose the relationship between a response y and predictor x is $k^y = \exp(bx)$. If we have iid data $x_i = \{1, 3, 4, 2, 4\}$ and $y_i = \{3, 1, 9, 10, 11\}$, then estimate k, b. (10 points)

   Taking the logarithm: $y \log k = bx \Rightarrow y = (b/\log k)x$
   Using the method of moments, $\mathbb{E}(y) = (b/\log k)\mathbb{E}(x)$
   $\overline{y} = (3 + 1 + 9 + 10 + 11)/5 = 6.8000$
   $\overline{x} = (1 + 3 + 4 + 2 + 4)/5 = 2.8000$
   So our estimate of $b/\log k$ is $\overline{y}/\overline{x} = 6.8000/2.8000 = 2.43$

4. Suppose $X_i \sim \mathcal{N}(\mu, 1)$ for $i = 1, \ldots, n$ are iid random variables.

(a) Compute the MLE for μ. (5 points)

$\hat{\mu} = \arg\max \prod_{i=1}^n \frac{1}{\sqrt{2\pi}} \exp(-(X_i - \mu)^2/2)$
Taking the negative logarithm and discarding constants:
$\hat{\mu} = \arg\min \sum_{i=1}^n (X_i - \mu)^2$
Setting the derivative equal to zero gives:
$2\sum_{i=1}^n (X_i - \mu) = 0 \Rightarrow \hat{\mu} = \frac{1}{n}\sum_{i=1}^n X_i = \overline{X}$

(b) Suppose the prior for μ is a uniform distribution with support $[a, b]$. Compute the MAP for μ. (5 points)

The prior is $g(\mu) = \mathbf{1}(a \leq \mu \leq b)$. So the MAP is:
$\hat{\mu} = \arg\max \mathbf{1}(a \leq \mu \leq b) \cdot \prod_{i=1}^n \frac{1}{\sqrt{2\pi}} \exp(-(X_i - \mu)^2/2)$
Taking the negative logarithm and discarding constants:
$\hat{\mu} = \arg\min \sum_{i=1}^n (X_i - \mu)^2 + \begin{cases} +\infty, & \mu < a \text{ or } \mu > 0 \\ 0, & a \leq \mu \leq b \end{cases}$
This is equivalent to:
$\hat{\mu} = \arg\min\{\sum_{i=1}^n (X_i - \mu)^2 \mid a \leq \mu \leq b\}$
Using the solution from part (a), we know the solution to the above problem is:
$\hat{\mu} = \begin{cases} a, & \text{if } \overline{X} < a \\ \overline{X}, & \text{if } a \leq \overline{X} \leq b \\ b, & \text{if } \overline{X} > b \end{cases}.$

5. Suppose $X \sim \mathcal{N}(\mu, \sigma_x^2)$ and $Y \sim \mathcal{N}(\mu, \sigma_y^2)$ are independent. Compute an unbiased estimate of the form $Z = aX + bY$ that minimizes the expected squared loss. (10 points)
   Hint: Recall that $Z \sim \mathcal{N}(a\mu + b\mu, a^2\sigma_x^2 + b^2\sigma_y^2)$.

   An unbiased estimate has $a\mu + b\mu = \mu$ or $a + b = 1$ or $a = 1 - b$. Using the bias-variance decomposition, we have: $\mathbb{E}((\mu - Z)^2) = (a\mu + b\mu - \mu)^2 + a^2\sigma_x^2 + b^2\sigma_y^2 = (1-b)^2\sigma_x^2 + b^2\sigma_y^2$
   Taking the derivative with respect to b and setting to zero gives:
   $-2(1-b)\sigma_x^2 + 2b\sigma_y^2 = 0 \Rightarrow b = \sigma_x^2/(\sigma_x^2 + \sigma_y^2)$
   So $a = \sigma_y^2/(\sigma_x^2 + \sigma_y^2)$