

Math 189r Fall 2016 Midterm

Name: _____ Time start/end: _____

Harvey Mudd College's Honor Code is in effect for all students taking this exam. If you feel unclear about any of the following instructions, please ask for clarification.

- This exam should be completed in a contiguous **three hour** period. If you so desire, you may insert a **twenty minute break** at the **one an a half hour mark** to get food, etc., but you are still restricted to the resources available to you during the exam (no internet). You may not write on your exam during the break if you choose to take one.
- No notes, books, computers or calculators will be allowed during the exam or break period except for one sheet of notes (8.5×11 ", front & back) that you have prepared yourself. When you are finished, staple the note sheet you used to the back of the exam. Please turn in your exam paper into the box outside Professor Gu's office (SHAN3420) no later than Monday, October, 6:30pm.
- Points may be deducted for answers that are not explained clearly.
- Please pace yourself as this exam has four questions, some with multiple parts.

Problem 1	/ 17 points
Problem 2	/ 24 points
Problem 3	/ 34 points
Problem 4	/ 25 points
Total	/ 100 points

Potentially Useful Equations

$$\mathbb{P}_{emp}(x) = \frac{1}{n} \sum_i \delta(x - x_i)$$

$$\int \delta(x - t) f(x) dx = f(t)$$

$$\|\mathbf{x}\|_1 = \sum_i |\mathbf{x}_i|$$

$$\|\mathbf{x}\|_2^2 = \mathbf{x}^\top \mathbf{x}$$

$$\text{cov}[\mathbf{x}] = \mathbb{E} [(\mathbf{x} - \mathbb{E}[\mathbf{x}])(\mathbf{x} - \mathbb{E}[\mathbf{x}])^\top]$$

$$\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{D/2} |\boldsymbol{\Sigma}|^{1/2}} \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right]$$

$$\mathbb{P}(\mathbf{x}_1|\mathbf{x}_2) = \mathcal{N}(x_1|\boldsymbol{\mu}_{1|2}, \boldsymbol{\Sigma}_{1|2})$$

$$\boldsymbol{\mu}_{1|2} = \boldsymbol{\mu}_1 + \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{22}^{-1} (\mathbf{x}_2 - \boldsymbol{\mu}_2)$$

$$\boldsymbol{\Sigma}_{1|2} = \boldsymbol{\Sigma}_{11} - \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{22}^{-1} \boldsymbol{\Sigma}_{21}$$

$$\hat{\boldsymbol{\mu}} = \frac{1}{N} \sum_i \mathbf{x}_i$$

$$\hat{\boldsymbol{\Sigma}} = \sum_i (\mathbf{x}_i - \hat{\boldsymbol{\mu}})(\mathbf{x}_i - \hat{\boldsymbol{\mu}})^\top$$

$$\text{Cat}(\mathbf{x}|\boldsymbol{\theta}) = \prod_{j=1}^K \theta_j^{\mathbb{I}(x_j=1)}$$

$$\mathbb{P}(\mathbf{x}|\boldsymbol{\theta}) = \frac{1}{Z(\boldsymbol{\theta})} h(\mathbf{x}) \exp(\boldsymbol{\theta}^\top \boldsymbol{\phi}(\mathbf{x}))$$

$$\mathbb{KL}(p||q) = \int p(x) \log \frac{p(x)}{q(x)} dx$$

$$\nabla_{\mathbf{x}} (\mathbf{x}^\top A \mathbf{x} + \mathbf{b}^\top \mathbf{x}) = (A + A^\top) \mathbf{x} + \mathbf{b}$$

$$\nabla_{\mathbf{x}}^2 (\mathbf{x}^\top A \mathbf{x} + \mathbf{b}^\top \mathbf{x}) = (A + A^\top)$$

1. (17 Points) Consider the Ridge Regression optimization problem

$$\text{minimize : } f(\mathbf{x}) = \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \|\mathbf{\Gamma x}\|_2^2.$$

(a) Solve for the optimal \mathbf{x}^* in closed form.

(b) Consider a dataset $\mathcal{D} = \{(0, 1), (1, 2), (2, 1), (3, 2)\}$. Construct a matrix A and \mathbf{b} from this dataset \mathcal{D} and compute the Ridge estimate with $\mathbf{\Gamma} = \mathbf{I}$. Note you may leave your solution in the form of the inverse of a matrix times a vector instead of computing the inverse by hand. All other operations must be simplified.

2. (24 Points - 3 each) Indicate whether each of the following statements is true (T) or not necessarily true (X). Be careful with your answers as no partial credit will be awarded on this problem.

- (a) The following is true: $\text{cov}[\mathbf{A}\mathbf{x} + \mathbf{b}] = \mathbf{A}^\top \text{cov}[\mathbf{x}]\mathbf{A}$.
- (b) Consider minimizing $-\ell(\mathbf{w}, \mathcal{D}_{train}) + \lambda\|\mathbf{w}\|_2^2$ where $\ell(\mathbf{w}, \mathcal{D}) = \sum_i y_i \log \sigma(\mathbf{w}^\top \mathbf{x}_i) + (1 - y_i) \log(1 - \sigma(\mathbf{w}^\top \mathbf{x}_i))$ is the average log-likelihood on a dataset \mathcal{D} for the ℓ_2 -regularized logistic regression model. If the training data is linearly separable, might some weights w_j become infinite if $\lambda = 0$? (Murphy 8.6.c)
- (c) Adding a zero-mean Gaussian Prior on each of the weights in a Logistic Regression model will encourage sparser weights than a zero-mean Laplace prior in a Maximum-a-Posteriori estimation of the optimal parameters.
- (d) Minimizing $\|X\mathbf{w} - \mathbf{y}\|_2^2$ and maximizing the likelihood $\prod_i \mathcal{N}(\mathbf{y}_i | \mathbf{w}^\top \mathbf{x}_i, \sigma^2)$ all result in the Normal Equations, the optimal solution $X^\top X \mathbf{w}^* = X^\top \mathbf{y}$, in the context of linear regression.
- (e) A covariance matrix Σ could potentially have a negative eigenvalue.
- (f) The level sets of a multivariate Gaussian density are always ellipses of the form $\{\mathbf{x} : \mathbf{x}^\top \mathbf{A} \mathbf{x} = k\}$.
- (g) The Multivariate Normal Distribution is not in the exponential family.
- (h) The goal of the Support Vector Machine is to maximize the margin, defined as the distance of the closest examples from the decision boundary.

3. (34 Points) Consider a Poisson distributed $X \sim \text{Poi}(\lambda)$ defined over $X \in \{0, 1, 2, \dots\}$ with probability mass function

$$\text{Poi}(x|\lambda) = e^{-\lambda} \frac{\lambda^x}{x!}.$$

Note that $\mathbb{E}[X] = \lambda$.

- (a) Show that the Poisson distribution is in the Exponential Family.

- (b) Consider creating a Generalized Linear Model from the Poisson distribution to model some count data. This will work since we showed that the distribution is an exponential family distribution. What is the distribution of the predicted value \hat{y} given a datapoint \mathbf{x}_i and weights \mathbf{w} ?

- (c) Derive an expression for the log-likelihood of a dataset \mathcal{D} , $\log \mathbb{P}(\mathcal{D}|\mathbf{w})$. Assume the data is identically and independently distributed.
- (d) Suppose we place an isotropic ($\Sigma = \sigma^2 \mathbf{I}$) Gaussian prior on the weights \mathbf{w} . Derive an expression which, when maximized, would maximize $\log \mathbb{P}(\mathbf{w}|\mathcal{D})$ (basically ignore constants). *Hint:* think about problem 1.a from homework 1. You may use results from that part of the assignment. You may introduce an auxiliary variable λ which somehow relates to σ^2 (and you don't need to define that relationship exactly).
- (e) Suppose we want to compute a maximum-a-posteriori estimate of \mathbf{w} given the prior from the previous problem. State the optimization problem we are trying to solve. Is this solvable in closed form? If so, solve it. Otherwise, compute the gradient of the objective function with respect to \mathbf{w} .

4. (25 Points) Suppose we want to model where we have some input data \mathbf{X} , with each datapoint corresponding to some observed function output \mathbf{f} . We also have points \mathbf{X}_* with which we want to predict what the function output \mathbf{f}_* will be. Our modelling assumption is this:

(1) The function outputs are jointly normal with mean $\boldsymbol{\mu}$ and covariance \mathbf{K} such that

$$\begin{pmatrix} \mathbf{f} \\ \mathbf{f}_* \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \boldsymbol{\mu} \\ \boldsymbol{\mu}_* \end{pmatrix}, \begin{pmatrix} \mathbf{K} & \mathbf{K}_* \\ \mathbf{K}_*^\top & \mathbf{K}_{**} \end{pmatrix} \right)$$

(2) The mean $\boldsymbol{\mu} = \boldsymbol{\mu}(X) = (m(\mathbf{x}_1), \dots, m(\mathbf{x}_n))$ for $m : \mathcal{X} \rightarrow \mathbb{R}$ where \mathcal{X} is the space your data live in (points on the line, locations on earth, molecules, etc.). Notationally, $\boldsymbol{\mu}_* = \boldsymbol{\mu}(X_*)$.

(3) The covariance between two datapoints $\mathbf{K}_{ij} = \kappa(\mathbf{x}_i, \mathbf{x}_j)$ where $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. Later in the course we'll call this $\kappa(\cdot, \cdot)$ a valid kernel. This implies that $\mathbf{K} \in \mathbb{R}^{n \times n}$, $\mathbf{K}_* \in \mathbb{R}^{n \times n_*}$, and $\mathbf{K}_{**} \in \mathbb{R}^{n_* \times n_*}$.

(a) Compute the posterior predictive distribution $p(\mathbf{f}_* | \mathbf{X}_*, \mathbf{X}, \mathbf{f})$ when we assume that the observed \mathbf{f} has no noise. *Hint:* consider conditioning a Gaussian.

(b) Compute the posterior predictive distribution $p(\mathbf{f}_* | \mathbf{X}_*, \mathbf{X}, \mathbf{y})$ where we assume our observations $\mathbf{y} = f(\mathbf{x}) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma^2)$. *Hint:* this assumption will only change \mathbf{f} to \mathbf{y} and \mathbf{K} (not \mathbf{K}_* or \mathbf{K}_{**}) into $\mathbf{K} + \sigma^2 \mathbf{I}$ from the equation in (1). For this part assume the mean $\boldsymbol{\mu} = \mathbf{0}$.

(c) Show that when predicting at only one location \mathbf{x}_* the mean of the predictive distribution from the previous part $\bar{\mathbf{f}} = \sum_{i=1}^n \alpha_i \kappa(\mathbf{x}_i, \mathbf{x}_*)$ with $\alpha = (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}$ being a constant vector depending on your training data and kernel. Armed with this knowledge, how might we restrict the predicted function $\bar{\mathbf{f}}(\mathbf{x})$ to be periodic? *Hint:* note that $\kappa(\mathbf{x}_i, \mathbf{x}_*)$ is function of \mathbf{x}_* and the sum of periodic functions is periodic.