STAT 542, Spring 2020 **Name (Print):** _____

Midterm, 04/15/2020 **Net ID (Print):** _____

Time Limit: 1PM - 5PM

This exam contains 8 pages (including this cover page) and 7 problems. Please read the following descriptions and requirements carefully.

- I suggest using one of these approaches to complete the exam. Please make sure to **include your name and NetID on the first page, and have proper question labels**. You do not need to copy the questions. You should **submit your answers/files to Compass2g**.
  - Answer all questions on a separate paper. Then take a picture of your answers. Please make sure that you have proper lighting and resolution when taking the picture.
  - Create an R markdown file (or a latex file) and type your answers. Mathematical equations should be done in Latex. Only choose this one if you are confident with your latex skill.
  - Print the exam and scan a finished version.
- **Mysterious or unsupported answers will not receive full credit**. A correct answer, unsupported by explanation, calculations, or algebraic work will receive no credit; an incorrect answer supported by substantially correct explanations and calculations might still receive partial credit.
- **Organize your work**, in a reasonably neat and coherent way.
- **You are NOT allowed to discuss the content of this exam to anyone else until the end of Apr 15.** A violation of this policy will lead to an immediate F as your final score of this course!

Section	Points	Score
1	15	
2	12	
3	15	
4	18	
5	15	
6	10	
7	15	
Total:	100	

1. Bias-Variance Trade-off. Multiple choices. Each question may have more than one correct answer.

   (a) (3 points) In general, which of the following regression models is expected to achieve the smallest variance for prediction if the model provides a consistent estimation of the truth?

   A. Linear regression
   B. Nadaraya–Watson estimator
   C. Smoothing spline
   D. k nearest neighbor

   (b) (3 points) Which of the following statement(s) is true for 1-nearest neighbor?

   A. It can achieve small variance
   B. It can achieve small bias
   C. At a target point x_0 where $P(Y = 1|X = x_0)$ is almost 1, 1-nearest neighbor classification makes low prediction error
   D. 1-nearest neighbor regression is similar to kernel estimator with small bandwidth

   (c) (3 points) Which of the following parameters can be used to tune the bias-variance trade-off

   A. λ in Lasso and Ridge
   B. k in nearest neighbor
   C. number of folds k in cross-validation
   D. α in regularized QDA for estimating the covariance matrix $\widehat{\Sigma}_k(\alpha) = \alpha\widehat{\Sigma}_k + (1-\alpha)\widehat{\Sigma}$ (or, see HW5 Q2b)
   E. number of trees T in boosting

   (d) (3 points) More bias is likely to happen at which of the following scenario?

   A. Predicting at a boundary point instead of an interior point when the Nadaraya–Watson estimator is used
   B. Knots positions are not optimized when using a quadratic spline
   C. Using EM algorithm instead of gradient descent when optimizing a likelihood function
   D. Using logistic regression instead of k nearest neighbor

   (e) (3 points) What are the consequences if both bias and variance of a model converge to 0

   A. The model has optimal training error
   B. The model has optimal prediction error on independent testing data
   C. The model is consistent
   D. The model could be inconsistent

2. Unsupervised Learning. Multiple choices. Each question may have more than one correct answer.

   (a) (3 points) Which of the following scenarios can benefit from performing PCA analyses?

   A. Reducing the dimensionality of the data
   B. Visualizing the data in a low dimensional plot
   C. Finding the most significant variable that associated with the outcome
   D. Detecting outliers
   E. Finding the variable with the largest marginal variance

   (b) (3 points) Which of the following method can guarantee to find the best clustering result that minimizes the within-cluster variation?

   A. Combinatorial algorithm
   B. k means
   C. Hierarchical clustering
   D. Multidimensional scaling
   E. Self-organizing maps

   (c) (3 points) A researcher collected a dataset with 200 observations, 1000 covariates (X) and a continuous outcome Y. PCA was performed on X and the following plot was obtained. Which of the following statements is true?

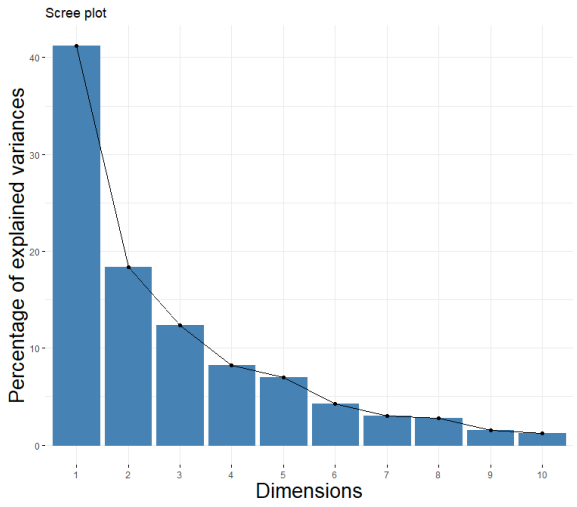

   A. The data lies approximately in a low dimensional sub-space
   B. A nearest neighbor regression will perform well on this data
   C. If Y depends only on the first several principal components, a ridge regression will perform well on this data
   D. If Y depends only on the first several principal components, a Lasso regression will perform well on this data
   E. If Y depends only on the first several principal components, a nearest neighbor regression will perform well on this data

(d) (3 points) Which of the following algorithms was covered in our lecture

A. Self-organizing map
B. UMAP
C. Spectral Clustering
D. tSNE
E. Principal Coordinates Analysis

3. Suppose we fit a linear (or penalized linear) regression using the observed data: $\mathbf{X}_{n\times p}$ and $\mathbf{y}_{n\times 1}$, with $n > p$. Suppose the true underlying model is given by $Y = X\boldsymbol{\beta} + \epsilon$ with ϵ's follows iid normal distribution, and $\mathbf{X}$ has full column rank. Let $\widehat{\boldsymbol{\beta}}^{\text{OLS}}$, $\widehat{\boldsymbol{\beta}}^{\text{Lasso}}$ and $\widehat{\boldsymbol{\beta}}^{\text{Ridge}}$ denote the OLS, lasso and ridge estimators, respectively on this same data. For ridge and Lasso, we assume that λ is a nonzero constant. Answer the following questions.

(a) (3 points) Select the smallest and the largest items from the list below ($\|\cdot\|$ denotes ℓ_2 norm)

A. $\|y - X\widehat{\boldsymbol{\beta}}^{\text{OLS}}\|^2$
B. $\|y - X\widehat{\boldsymbol{\beta}}^{\text{Lasso}}\|^2$
C. $\|y - X\widehat{\boldsymbol{\beta}}^{\text{Ridge}}\|^2$
D. $\|y\|^2$

(b) (3 points) What is true regarding the Hat-matrix?

A. The rank is p
B. The rank is n
C. $\mathbf{H}(y - X\widehat{\boldsymbol{\beta}}^{\text{OLS}}) = \mathbf{0}$
D. $\mathbf{H}(y - X\widehat{\boldsymbol{\beta}}^{\text{Lasso}}) = \mathbf{0}$
E. $\mathbf{H}(y - X\widehat{\boldsymbol{\beta}}^{\text{Ridge}}) = \mathbf{0}$

(c) (3 points) Suppose we run a ridge regression with just one column of X, e.g., regress $\mathbf{y}$ on X_1 and get coefficient β_1. We now include an exact copy $\tilde{X}_1 = X_1$, and regress $\mathbf{y}$ on both $\tilde{X}_1$ and X_1. We would expect the two coefficients to be

A. One of them is the same as β_1, the other one is 0
B. The two are exactly the same
C. Ridge regression cannot obtain a proper solution in this case
D. They will be the same as the Lasso regression solution with some other λ value
E. Any situation could happen, depends on the correlation between $\mathbf{y}$ and X_1

(d) (3 points) The Lasso regression solution path is the same as

A. Forward stepwise regression
B. Forward stagewise regression
C. Backward stepwise regression
D. LARS
E. Best subset selection

(e) (3 points) When we decrease the turning parameter λ of Lasso from infinity to 0, the residual sum of squares of the training data will

A. Steadily increase
B. Steadily decrease
C. Remain constant
D. Increase initially, and then start decreasing
E. Decrease initially, and then start increasing

4. Other topics. Multiple choices. Each question may have more than one correct answer.

(a) (3 points) Select the models with the smallest and the largest degrees of freedom.

A. Cubic spline with 4 knots
B. Quadratic spline with 6 knots
C. Linear spline with 8 knots
D. Natural cubic spline with 11 knots
E. Piecewise linear with 5 knots

(b) (3 points) Consider two curves $\hat{g}_1$ and $\hat{g}_2$, defined by

$$\hat{g}_1 = \arg\min_g \left(\sum_{i=1}^{n} (y_i - g(x_i))^2 + \lambda \int \left[g^{(3)}(x) \right]^2 dx \right)$$

$$\hat{g}_2 = \arg\min_g \left(\sum_{i=1}^{n} (y_i - g(x_i))^2 + \lambda \int \left[g^{(4)}(x) \right]^2 dx \right)$$

where $g^{(m)}$ represents the mth derivative of g.

A. As $\lambda \to \infty$, $\hat{g}_1$ have the smaller training RSS
B. As $\lambda \to \infty$, $\hat{g}_2$ have the smaller training RSS
C. As $\lambda \to 0$, $\hat{g}_1$ have the smaller testing RSS
D. As $\lambda \to 0$, $\hat{g}_2$ have the smaller testing RSS

(c) (3 points) Suppose that we take a data set, divide it into equally-sized training and test sets, and then try out two different classification procedures. First, we use logistic regression and get an error rate of 20% on the training data and 30% on the test data. Next, we use 1-nearest neighbors (i.e. $k = 1$) and get an average error rate (averaged over both training and testing data sets) of 18%. Based on these results, which method should we prefer to use for the classification of new observations?

A. Logistic regression
B. 1-nearest neighbor
C. A combination of the two
D. Cannot decide

(d) (3 points) Increasing λ in Lasso will

A. Increase the bias
B. Decrease the bias
C. Increase the variance
D. Decrease the variance

(e) (3 points) In a genetic study, with possibly millions of SNPs being measured, which of the following models may greatly suffer in terms of its performance

A. Best subset selection
B. Lasso
C. Ridge
D. Fused Lasso
E. Nadaraya–Watson kernel estimator

(f) (3 points) Which of the following is computationally the easiest (when both n and p are large)? Choose only one answer.

A. Best subset selection
B. Lasso
C. Nonliear SVM
D. Boosting with deep trees
E. Deep learning

5. Another example of the EM algorithm. In genetics, we often assume a Hardy-Weinberg equilibrium. Consider the example of ABO blood type and read the following table. This is essentially saying that among human beings, the overall probability of having a Gene "A" is p_A, the probability of having Gene "B" is P_B, and the probability of having Gene "O" is p_O. Since each person will have two copies of a gene, one from each patient, it is possible that you have six different possible "Genotypes". Of course, the probability of having "AA" would be p_A^2, meaning that both your patients contribute a type "A" gene randomly to you. And the rest of the table should be self-explanatory. On the other hand, we know that your blood type (phenotype) will only be "A", "B", "O" and "AB", because, for example, if you have one "A" gene and one "O" gene, your blood type will be "A". And the rest of the possibilities are also demonstrated in the table.

Genotype	Probability	Phenotype
AA	P_A^2	A
AO	$2P_AP_O$	A
BB	P_B^2	B
BO	$2P_BP_O$	B
OO	P_O^2	O
AB	$2P_AP_B$	AB

Here comes the problem: we observe the phenotype (your blood type), but not the genotype (the types of your two genes) if without a gene sequencing machine. Hence, from the observed phenotype, it is difficult to estimate the probabilities P_A, P_B and P_O ($= 1 - P_A - P_B$). An EM algorithm can be used. Suppose we observe n subjects and their phenotypes, we want to estimate these parameters. Then by the likelihood of a multinomial distribution, we should have the likelihood as (after removing some constants that involves combinatorial numbers)

$$L(P_A, P_B) = (P_A^2 + 2P_AP_O)^{n_A} \times (P_B^2 + 2P_BP_O)^{n_B} \times (P_O^2)^{(n_O)} \times (2P_AP_B)^{(n_{AB})}.$$

Here, n_A, n_B, n_AB, and n_O represent the number of subjects with the respective phenotype (blood type), and they should, of course, sum up to n. It is difficult to directly optimize this likelihood and solve for the parameters. However, we could introduce hidden variables n_{AO}, n_{AA}, n_{BO}, n_{BB}, which indicate the count of unobserved genotypes, respectively. And they should satisfy $n_{AO} + n_{AA} = n_A$ and $n_{BO} + n_{BB} = n_B$, hence we only need n_{AA} and n_{BB} as the hidden variables. The rest of the job is to perform the EM algorithm.

(a) (5 points) Write down the complete data log-likelihood (you can ignore unnecessary constants if you want) based on both the observed (n_A, n_B, n_{AB}, n_O) and unobserved (n_{AA}, n_{BB}) data. In other words, the analog of $\ell(\mathbf{x}, \mathbf{z}|\boldsymbol{\theta})$ (if you use our notation) of this problem.

(b) (5 points) For the E Step, calculate the conditional expectation of n_{AA} and n_{BB} given the observed data and the current parameter values. This is the analog of $E[Z|\mathbf{x}, \boldsymbol{\theta}]$. Hint: what is the conditional distribution of n_{AA} given n_A with all the other parameters known? This could be just one line of equation, but you should provide an explanation.

(c) (5 points) For the M Step, provide the formula to update the parameter estimate of P_A, given that all the counts are observed. Hint: from the relationship, we can realize that $2P_A^2 + 2P_AP_O + 2P_AP_B = 2P_A$. On the left-hand side, each quantity has its respective observed counts, on the right-hand side, it is our parameter of interest. This could be just one line of equation, but you should provide an explanation.

6. Variance of two-dimensional kernel density estimator. In HW4, we derived the bias part of a two-dimensional kernel density estimator. Let's finish the consistency proof by deriving the variance part (at a single target point). All assumptions of this question are exactly the same as HW4.

   (a) (8 points) Derive the variance of the estimator

   $$\mathrm{Var}[\widehat{f}(x)] =$$

   (b) (2 points) What is the rate of variance?

7. In the support vector machine problem, in some cases, we do not want to treating each observation equally. This means that we will assign a weight w_i to each observation. If an observation has a nonzero slack variable, we will penalize it proportional to the weights. This results in the following primal form of SVM in the non-separable case:

$$\begin{aligned} \min \quad & \frac{1}{2}\boldsymbol{\beta}^T\boldsymbol{\beta} + C\sum_{i=1}^{n} w_i\xi_i \\ \text{subject to} \quad & y_i(x_i^T\boldsymbol{\beta} + \beta_0) \geq 1 - \xi_i, \quad i = 1, \ldots, n \\ & \xi_i \geq 0 \quad i = 1, \ldots, n \end{aligned}$$

   (a) (3 points) Write down the Lagrangian $\mathcal{L}$ of this constrained optimization problem by introducing two sets of Lagrange multiplier. It is recommended that you use α_i's and γ_i's.

   (b) (5 points) Optimize the Lagrangian by minimizing over $\boldsymbol{\beta}$ and β_0 and derive their solutions.

   (c) (5 points) Plug the solution back to the Lagrangian and find the dual form of this SVM problem.

   (d) (2 points) Suppose there are three observations (their labels are represented by dot and circle). Based on their current weights, the SVM decision line is plotted below. If the weights of all circles increase slightly, where do you expect the new decision line lies? Draw the new decision line approximately on the plot.