PRACTICE FINAL EXAM

STA414/2104 Winter 2021

University of Toronto

Name:

Student #:

Enrolled in course (circle one): STA414 STA2104

Enrolled in section (circle one): Monday Tuesday

Exam duration: **150 minutes**

Please check that your exam has **10 pages**, including this one. The total possible number of points is 100.

Read the following instructions carefully:

1. Exam is closed book and internet. You can use an optional A4 aid sheet - double-sided.
2. If a question asks you to do some calculations, you must *show your work* to receive full credit.
3. Conceptual questions do not require long answers.
4. You will submit your answer to each question separately on crowdmark. If you run into any technical difficulties, you can send an email to sta414-2021-tas@cs.toronto.edu attaching your solutions. Any email sent after the exam is over will not be considered.
5. Do not share the exam with anyone or in any platform!
6. Lastly, enjoy the problems!!!

1. Academic Integrity Statement. Academic integrity is a fundamental value of learning and scholarship at the UofT. Participating honestly, respectfully, responsibly, and fairly in this academic community ensures that your UofT degree is valued and respected as a true signifier of your individual academic achievement.

The University of Toronto's Code of Behaviour on Academic Matters outlines the behaviours that constitute academic misconduct, the processes for addressing academic offences, and the penalties that may be imposed. You are expected to be familiar with the contents of this document.

Potential offences include, but are not limited to:

- Working together to answer questions.
- Looking at someone else's answers.
- Letting someone else look at your answers.
- Sharing or posting the exam questions.
- Discussing answers or the exam questions with anyone else in or outside the course.
- Misrepresenting your identity or having someone else complete your exam.

Prior to beginning this exam, you must attest that you will follow the Code of Behaviour on Academic Matters and will not commit academic misconduct in the completion of this online exam. Affirm your agreement to this by rewriting the following statement:

I, [name] (type your full name here), [stnum] (type your student number here), agree to fully abide to the Code of Behaviour on Academic Matters. I will not commit academic misconduct, and am aware of the penalties that may be imposed if I commit an academic offence.

After you write the above statement by hand, with your name and student number, you should submit it as answer to the first question. All suspected cases of academic dishonesty will be investigated following the procedures outlined in the Code of Behaviour on Academic Matters.

2. True/False. For each statement below, say whether it is true or false, and give a **one or two sentence** justification of your answer.

a) Adding more training data always reduces overfitting.

b) For k=2, the k-means algorithm is equivalent to the Fisher's Linear Discriminant.

c) An ensemble of models always reduces the bias.

d) A linear SVM will find the same decision boundary as logistic regression.

3. Reinforcement Learning.

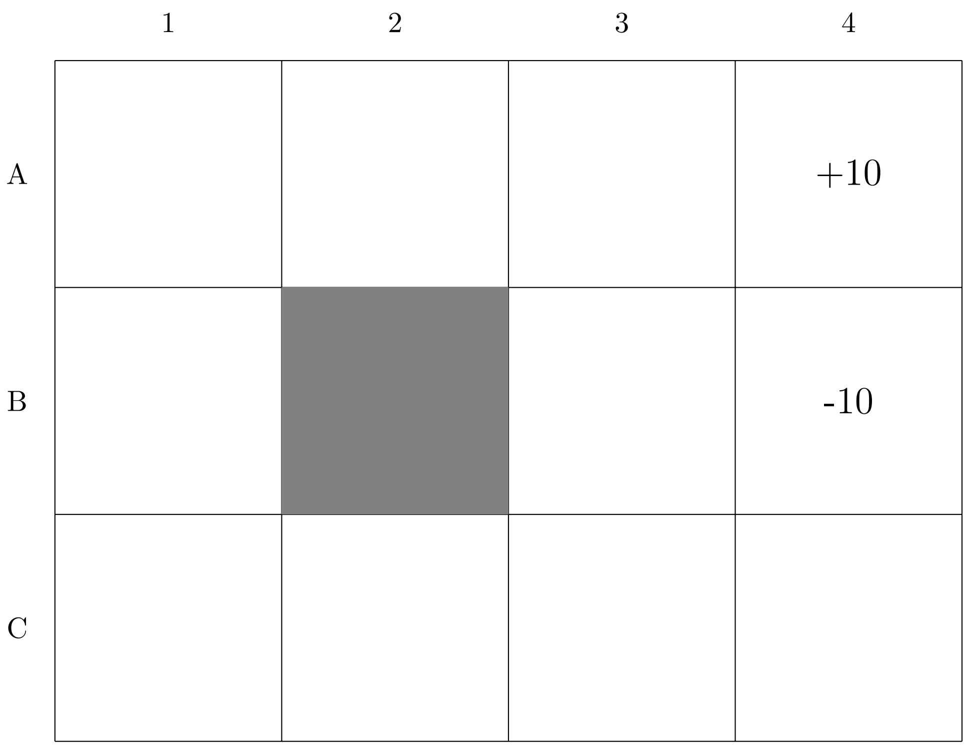

Consider the familiar robot navigation task within the gridworld shown above. You can move in any of the four directions (left/right/up/down) unless blocked by one of the gray obstacles at B2. The rewards are +10 for state A4, and -10 for state B4. The reward for every other state is 0.

a) Assume that the state transitions are deterministic. Consider applying Q-learning algorithm when all the Q values (action-value function) are initialized to zero, step size $\alpha = 1$ and discount factor $\gamma = 0.8$. Write the Q values of Q-learning algorithm after the robot has executed the following state sequences:

- B1 $\rightarrow$ A2 $\rightarrow$ A2 $\rightarrow$ A3 $\rightarrow$ B3 $\rightarrow$ B4
- A2 $\rightarrow$ A3 $\rightarrow$ A4
- C1 $\rightarrow$ C2 $\rightarrow$ C3 $\rightarrow$ B3 $\rightarrow$ A3 $\rightarrow$ A4

.

b) Assume the robot will now use the policy of always performing the action having the greatest Q value. Is this the optimal policy? Why or why not?

c) Suppose state A3 also has a reward of -10. How can we ensure that our agent is still able to find the +10 reward without going through -10 states in this new environment?

4. Backpropagation. Consider a L_2 regularized single layer neural network model that predicts continuous 1d targets $y = \sigma(z) \in \mathbb{R}$ where $z = wh + b$ and $h = \sigma(w'x + b')$, and σ is an activation function. To train, we use mean squared error from the targets $t \in \mathbb{R}$ with L_2 penalty on w, b' given by $\mathcal{L} = (y - t)^2/2 + w^2 + b'^2$

a) Write the loss as a function of the parameters w, b and compute directly $\frac{\partial \mathcal{L}}{\partial w}$, $\frac{\partial \mathcal{L}}{\partial b'}$.

b) Now compute $\frac{\partial \mathcal{L}}{\partial w}$, $\frac{\partial \mathcal{L}}{\partial b'}$ using the backprogation algorithm.

c) What are the disadvantages of doing a) versus backpropagation? Why do we use backpropagation in machine learning as opposed to direct differentiation?

5. Principal Component Analysis. Recall that the optimal PCA subspace can be determined from the eigendecomposition of the empirical covariance matrix $\hat{\mathbf{\Sigma}}$. Also recall that the eigendecomposition can be expressed as:

$$\hat{\mathbf{\Sigma}} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^\top,$$

where $\mathbf{Q}$ is an orthogonal matrix and $\mathbf{\Lambda}$ is a diagonal matrix. Assume the eigenvalues are sorted from largest to smallest. You may assume all of the eigenvalues are distinct.

1. If you've already computed the eigendecomposition (i.e. $\mathbf{Q}$ and $\mathbf{\Lambda}$), how do you obtain the orthogonal basis $\mathbf{U}$ for the optimal PCA subspace? (You do not need to justify your answer.)

2. The PCA code vector for a data point $\mathbf{x}$ is given by $\mathbf{z} = \mathbf{U}^\top(\mathbf{x} - \hat{\boldsymbol{\mu}})$ where $\hat{\boldsymbol{\mu}}$ is the data mean. Show that the entries of $\mathbf{z}$ are uncorrelated.

6. Support Vector Machines.

Support vector machines learn a decision boundary leading to the largest margin from both classes. You are training SVM on a tiny dataset with 4 points shown below. This dataset consists of two examples with class label -1 (denoted with plus), and two examples with class label +1 (denoted with triangles).

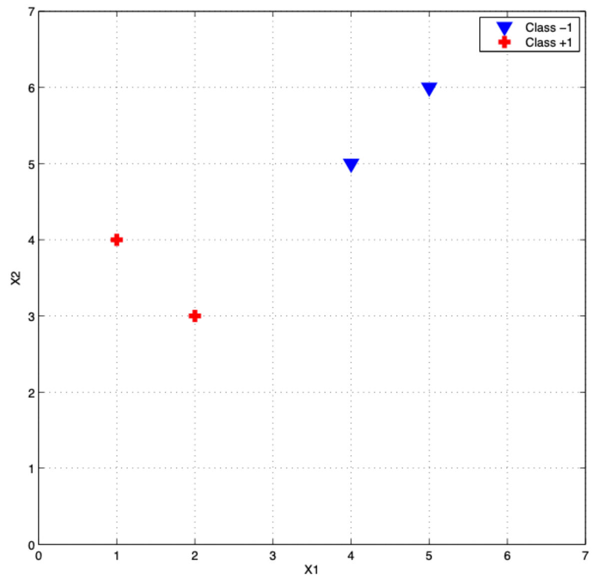

a) Write down the SVM loss function for this data and state how to find the weight vector $\mathbf{w}$ and bias b.

b) Draw the (approximate) decision boundary.

7. Maximum A posteriori Probability .

The Laplace distribution, parameterized by μ and β, is defined as follows:

$$\text{Laplace}(w; \mu, \beta) = \frac{1}{2\beta} \exp\left(-\frac{|w-\mu|}{\beta}\right).$$

Consider a variant of the homework2 question where we assume that the prior over the weights $\mathbf{w}$ consists of an independent zero-centered Laplace distribution for each dimension, with shared parameter β:

$$\begin{aligned} w_j &\sim \text{Laplace}(0, \beta) \\ t \mid \mathbf{w} &\sim \mathcal{N}(t; \mathbf{w}^\top \mathbf{x}, \sigma^2) \end{aligned}$$

For reference, the Gaussian PDF is:

$$\mathcal{N}(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right).$$

1. Suppose you have a labeled training set $\{(\mathbf{x}_i, t_i)\}_{i=1}^N$. Give the cost function you would minimize to find the MAP estimate of $\mathbf{w}$.

2. Based on your answer to part (a), how might the MAP solution for a Laplace prior differ from the MAP solution if you use a Gaussian prior?

8. EM Algorithm.

1. Is EM algorithm a supervised or an unsupervised learning method? Explain your answer.

2. How does EM algorithm and k-means compare? Write 3 similarities and 3 differences.

3. Explain why we call these steps expectation and maximization steps. What is it that we take expectation of and what is it that we maximize?