

- Please do not open the exam before you are instructed to do so.
- **Electronic devices are forbidden on your person**, including cell phones, tablets, headphones, and laptops. Leave your cell phone off and in a bag; it should not be visible during the exam.
- The exam is closed book and closed notes except for your two 8.5×11 inch cheat sheets.
- You have 2 hours and 50 minutes (unless you are in the DSP program and have a larger time allowance).
- Please write your initials at the top right of each page after this one (e.g., write “JD” if you are John Doe). Finish this by the end of your 2 hours and 50 minutes.
- Mark your answers on the exam itself in the space provided. Do **not** attach any extra sheets.
- For multiple choice questions, fill in the bubble for the single best choice.
- For short and long answer questions, write within the boxes provided. If you run out of space, you may use the last four pages to continue showing your work.

| | |
|---------------------------------------|--|
| Your Name | |
| Your SID | |
| Name and SID of student to your left | |
| Name and SID of student to your right | |
| Any winter break plans? | |
| Have a funny ML joke? | |

- CS 189
- CS 289A

This page intentionally left blank.

1 Multiple Choice

For the following questions, select the **single best response**. Each question is worth 1.5 points.

- Which of the following is a benefit of t-SNE over PCA?
 - t-SNE is robust to local optima during optimization.
 - t-SNE creates a highly interpretable mapping of the original points.
 - t-SNE can be computed in a closed-form solution.
 - t-SNE preserves rich local structures.
- Which of the following is **not true** about Transformer models?
 - They add/concatenate a positional encoding to each token in the input sequence before passing them into the first transformer layer.
 - The key, query and value vectors generated from each token in a self-attention layer must have the same dimension.
 - The attention computations for each head in a multi-head attention layer can be parallelized.
 - Transformer models use masked attention for autoregressive tasks during training time to prevent lookup of future tokens within each self-attention layer.
- In an L2-regularized linear regression model, which of the following is the most likely effect of increasing the value of λ ?
 - Decrease the model's bias and increase its variance.
 - Increase the model's bias and decrease its variance.
 - Increase both the model's bias and its variance.
 - Decrease both the model's bias and its variance.
- The general form of regularized linear regression is

$$\arg \min_w \|y - Xw\|_2^2 - \lambda \cdot r(w)$$

where $r : \mathbb{R}^d \rightarrow \mathbb{R}$ is the regularization function. Which of these regularization functions would tend to induce a sparse w ? In other words, which function would most likely cause some coefficients of w to be set to zero?

- 1
- $\sum_{i=1}^d |w_i|$
- $\sum_{i=1}^d w_i^2$
- $\max(|w_1|, \dots, |w_d|)$

5. A classic example of a kernel for edge detection is shown below. Consider $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$, defined as the convolution of this kernel onto a single-channel image. Assume that there is appropriate padding resulting in an output with the same shape. What type of operation does this function **not** exhibit equivariance with?

$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

- Permutations
 - Translations
 - Rotations in multiples of 90 degrees
 - Horizontal and vertical reflections
6. In the context of graph neural networks, which of the following is an **invalid** function for updating each nodes' value based on its neighbors? Let h_v be the embedding for neighbor $v \in \mathcal{N}(u)$, where u is the target node.
- $\sum_v h_v$
 - $\max_v h_v$
 - $\sum_v v \cdot h_v$
 - $\prod_v h_v$
7. Consider a binary classification problem in which we use an asymmetrical loss function: $L(1, 0) = 5$, $L(0, 1) = 1$, and 0 otherwise. Remember that $L(\hat{y}, y)$ denotes the loss for a class prediction of \hat{y} and a ground-truth class y . Which of the following is true of a point lying on the Bayes decision boundary?
- The posterior probability of class 0 is equal to the posterior probability of class 1.
 - The posterior probability of class 0 is equal to $\frac{1}{5}$ the posterior probability of class 1.
 - The posterior probability of class 1 is equal to $\frac{1}{5}$ the posterior probability of class 0.
 - The class-conditional probability of class 0 is equal to the class-conditional probability of class 1.
8. Which of the following is **true** about k -Nearest Neighbors?
- As $k \rightarrow \infty$, the model's error approaches the Bayes error.
 - Each data point's Voronoi cell defining the decision boundaries is a convex set for any possible distance metric.
 - The k -Nearest Neighbors algorithm is generally computationally intensive at test time.
 - k -Nearest Neighbors is an unsupervised learning technique.
9. Which of the following is **not true** about the convergence of value iteration?

- For a finite MDP and discount factor $\gamma \in (0, 1)$, value iteration is guaranteed to converge.
 - Policy iteration and value iteration do not converge to the same optimal values.
 - Value iteration can take an infinite amount of iterations to converge to the exact optimal values.
 - Value iteration will always converge to the same values, regardless of how the values are initialized.
10. When training trees in a random forest, a valid reason for why we restrict ourselves to a random feature subset at each split is:
- To encourage implicit feature selection, similar to what LASSO does in regression.
 - To decrease single tree variance, which improves random forest performance.
 - To prevent the dominance of some features to cause each tree to become very similar.
 - None of the above.
11. Which of the following is **true** about neural networks?
- Adding nonlinear activation functions like ReLU or Sigmoid after each layer of a neural network increases model expressiveness, as opposed to linear activation functions.
 - A neural network's activation functions must be monotonic for it to converge.
 - Adding layers to a neural network without any nonlinearities will increase performance.
 - The tanh activation function does not suffer from vanishing gradients.
12. Which of the following is NOT a way in which control theory differs from traditional pattern recognition?
- The controller must be able to adapt to varying conditions in an environment.
 - The controller must generalize to sensory inputs under a variety of lighting and occlusion conditions in an environment.
 - The controller must deal with a dynamic environment that changes over time, as opposed to a static, unchanging environment.
 - The controller must learn to recover from disturbances in an environment.
13. Which of the following best explains why human-centric video models lack deeper understanding and fail to achieve the performance of image models?
- There is a lack of human-centric video data online as opposed to human-centric image data.
 - Human-centric video data is currently too low-quality (in terms of resolution) for deep models to meaningfully discern objects in an image.
 - Video models have yet to understand the hierarchy of human behavior, from particular actions at a timestep to larger goals and intentions.



initial here

- Image models are trained on i.i.d. images, while the temporality of frames in a video breaks this i.i.d. assumption.
14. Consider a Hidden Markov Model (HMM) with a sequence of observations $O = (o_1, o_2, \dots, o_T)$ and states $S = (s_1, s_2, \dots, s_T)$. Which of the following is true about the Viterbi algorithm?
- The Viterbi algorithm maximizes the probability at each state and timestep t , i.e. it gets the optimal solution at each step.
 - The Viterbi algorithm performs Monte Carlo sampling over the state sequence.
 - The Viterbi probability $v_t(j)$ is computed by taking the sum over all possible paths leading up to state j at timestep t , weighted by the transition probability from the previous state to the current state.
 - The Viterbi algorithm chooses the state sequence that maximizes the likelihood of the observation sequence.



initial here

2 Short Answer

1. (4 points) Suppose you create a 2D convolutional layer with n_{in} input channels, n_{out} output channels, kernel size k , stride s and padding p .

(a) How many learnable parameters does this layer have?

(b) We pass in a 3D tensor with height H , width W and depth n_{in} as input to this layer. What will be the height, width and depth of the output?



initial here


2. (2 points) In discussion, we derived the Bellman expectation equation for state value functions, which expressed $V(s)$ for some state s as a recursive function of $V(s')$ (where s' is the state at the following time step), the reward r , transition probabilities $p(s', r | s, a)$, policy $\pi(a | s)$ and discount factor $\gamma \in [0, 1]$. Write down a Bellman equation for the state-action value function, i.e., express $Q(s, a)$ as a recursive function in terms of $Q(s', a')$ (where s' is the state at the next time step and a' the action taken from it), the reward r , transition probabilities $p(\cdot, \cdot | \cdot, \cdot)$, policy $\pi(\cdot | \cdot)$ and discount factor $\gamma \in [0, 1]$.

3. (2 points) In discussion we showed why training a random forest of stumps (trees with a single split) is generally a bad idea. However, another form of ensemble learning covered in class, boosting, is often done using stumps. Explain why stumps are good learners for boosting algorithms, in contrast to the deep trees used in random forests.



initial here

4. (3 points) Consider the following computation graph:



The computations of the intermediary/output nodes follow these equations:

$$\begin{aligned} d &= s(a) \\ f &= \frac{d}{a} \\ g &= 3e + 4d + 2f \\ e &= a^2 * c + b \end{aligned}$$

$s(z) = \frac{1}{1+e^{-z}}$, which is the sigmoid function Find an expression for $\frac{\partial g}{\partial a}$. In this problem, all variables are scalars.



initial here

5. (3 points) In lecture, we saw to how derive linear regression using maximum likelihood estimation assuming Gaussian errors. Mathematically, we saw that OLS was the MLE for β under the model

$$y \sim N(X\beta, \sigma^2 I)$$

In this problem, we relax the assumption that each y_i has the same (known) error variance σ^2 and instead has its own (known) variance σ_i^2 . Our probabilistic model is therefore:

$$y \sim N(X\beta, \Sigma)$$

where $\Sigma = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_n^2)$.

Derive the maximum likelihood estimate of β under this model. Assume $\beta \in \mathbb{R}^d$, $X \in \mathbb{R}^{n \times d}$, $y \in \mathbb{R}^n$ and that X is full rank.



initial here

6. (4 points) Consider the problem of text classification, where you want to learn a probability of whether a token is associated with one of d classes. An example use case of this model would be sentiment analysis, where you would want to classify a word as being happy, sad, or angry. More formally, let X be a random variable that captures the distribution of possible labels for some token. The possible labels themselves are one of d classes: $1, 2, \dots, d$. We collect n samples of X and sum up the counts of each class label, which we will call C_1, C_2, \dots, C_d (note these are also random variables). Note that $C_1 + C_2 + \dots + C_d = n$. With this data, we generate estimates of the class probabilities $\hat{p}_1, \hat{p}_2, \dots, \hat{p}_d$. Let the ground truth class probabilities be p_1, p_2, \dots, p_d .

One way to determine $\hat{p}_1, \hat{p}_2, \dots, \hat{p}_d$ is through **Laplace smoothing**, which pretends that we saw the token α additional times for each class. *Without* Laplace smoothing, the probabilities would be:

$$\hat{p}_i = \frac{C_i}{n}$$

With Laplace smoothing, the probabilities would be:

$$\hat{p}_i = \frac{C_i + \alpha}{n + d\alpha}$$

(a) (2 points) Provide a reason for why Laplace-smoothed probabilities may be preferable to basic counting in text classification. 1-2 sentences is sufficient. *Hint*: think about edge cases.



initial here

- (b) (2 points) For a specific class i , calculate the bias of the Laplace-smoothed probabilities, as compared to the ground truth probabilities. Simplify as much as possible. *Hint:* What is $\mathbb{E}[C_i]$, i.e. what's our expectation of the number of labels with class i ?



initial here

3 Everything is Gaussian

One activation function that has become popular in language models and modern deep learning models is GELU, or Gaussian Error Linear Units. GELU is defined as follows:

$$\text{GELU}(x) = xP(X \leq x) = x\Phi(x)$$

where $X \sim \mathcal{N}(0, 1)$. In other words, it is the product of x and the CDF $\Phi(x)$ of a standard normal distribution. In this problem, we will compare it to the ReLU activation you learned about in class.

- (a) (3 points) Prove that $\text{GELU}(x) \leq \text{ReLU}(x)$ for all $x \in \mathbb{R}$.



initial here

- (b) (2 points) Consider the value of $\text{GELU}(x)$ at $x = 0$ and at the limits $x \rightarrow \pm\infty$. Explain why GELU can be thought of as a smoothed approximation of ReLU.

- (c) (3 points) Prove or disprove that GELU is a convex function. What about ReLU?
Hint: the derivative of a distribution's CDF is its PDF.

4 Books Are All You Need

Alice works at a local bookstore and is responsible for their weekly promotion strategy. Having learned about Markov Decision Processes (MDPs), she plans to apply this method to optimize book sales. The bookstore focuses mainly on two genres: fiction and non-fiction. Each week, Alice must decide which genre to promote based on the store's current main genre, aiming to maximize sales over the semester.

Alice formulates the problem as an MDP. The state space $\mathcal{S} = \{F, N\}$ represents the main genre at the beginning of week t (F for fiction, N for non-fiction). The action space $\mathcal{A} = \{A, B\}$ represents the promotion plan for week t , with plan A promoting fiction and plan B promoting non-fiction.

The transition dynamics $P(s'|s, a)$ describe the likelihood of the promotion strategy changing the main genre. For example, if the current main genre is fiction (F), promoting fiction (action A) has an 80% chance of keeping the genre as fiction and a 20% chance of switching to non-fiction (N), whereas promoting non-fiction has an equal chance of keeping the genre as fiction and switching to non-fiction. The same dynamics apply if the current main genre is non-fiction. The transition dynamics are:

$$\begin{aligned} P(s' = F|s = F, a = A) &= 0.8; & P(s' = N|s = F, a = A) &= 0.2 \\ P(s' = F|s = F, a = B) &= 0.5; & P(s' = N|s = F, a = B) &= 0.5 \\ P(s' = F|s = N, a = A) &= 0.5; & P(s' = N|s = N, a = A) &= 0.5 \\ P(s' = F|s = N, a = B) &= 0.2; & P(s' = N|s = N, a = B) &= 0.8 \end{aligned}$$

The reward $R(s, a)$ represents the sales for week t . Sales are typically higher when the promotion strategy aligns with the current main genre of the store. The rewards are defined as:

$$\begin{aligned} R(s = F, a = A) &= 10; & R(s = F, a = B) &= 5 \\ R(s = N, a = A) &= 5; & R(s = N, a = B) &= 10 \end{aligned}$$

Alice wishes to determine the optimal promotion strategy using the MDP model. Observing that non-fiction has been selling better, she initializes the value functions as $V_0(F) = 5, V_0(N) = 10$. However, she prefers promoting fiction, so she sets the initial policy as $\pi_0(A|s) = 0.6, \pi_0(B|s) = 0.4, \forall s \in \mathcal{S}$. She also assumes a discount factor $\gamma = 0.5$.



initial here

- (a) (1 point) Both the policy iteration algorithm and the value iteration algorithm can be used to find the optimal policies given an MDP. In both algorithms, we need to compute the following expectation:

$$Q(s, a) = R(s, a) + \gamma \sum_{s'} P(s'|s, a)V(s')$$

Compute $Q(s = F, a = A)$ using $V_0(s)$.

- (b) (2 points) In part (b)-(d), assume $Q(\cdot, \cdot)$ are the following:

$$Q(s = F, a = A) = 10$$

$$Q(s = F, a = B) = 8$$

$$Q(s = N, a = A) = 6$$

$$Q(s = N, a = B) = 15$$

Perform one step of policy evaluation by computing the updated value functions $V(F)$ and $V(N)$ using the initial policy $\pi_0(\cdot|\cdot)$ and $Q(\cdot, \cdot)$.



initial here

- (c) (2 points) Perform one step of policy improvement by computing the best action $\pi(s)$ for states F and N using $Q(\cdot, \cdot)$.

- (d) (2 points) Perform one step of value iteration by computing the updated value functions $V(\cdot)$ for states F and N using $Q(\cdot, \cdot)$.

- (e) (2 points) Alice notices that the success of the current week's promotion depends on the sequence of promoted genres over the past two weeks. How might this new scenario be problematic in the original MDP model Alice formulated? How should the state space \mathcal{S} be redefined to account for this new insight?

5 Mixture of Linear Regressors

In homework 7, we employed the Expectation-Maximization algorithm to learn the parameters of a Gaussian mixture model. In this problem, we will apply the EM algorithm to learn the parameters for a mixture of linear regression models (also called a “Mixture of Experts” model in ML literature).

Suppose we have a dataset $D = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ of IID samples, where data point \mathbf{x}_i belongs to one of two hidden classes $z_i = 0$ or $z_i = 1$. If \mathbf{x}_i belongs to hidden class $z_i = 0$, then its corresponding label was generated according to $y_i = \boldsymbol{\theta}_0^T \mathbf{x}_i + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \sigma^2)$. Similarly, if \mathbf{x}_i belongs to hidden class $z_i = 1$, then its corresponding label was generated according to $y_i = \boldsymbol{\theta}_1^T \mathbf{x}_i + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \sigma^2)$. We will assume that σ^2 is known. On the other hand, the true parameters $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$ as well as the hidden labels z_i (the hidden label for \mathbf{x}_i) are unknown.

Since it is reasonable to assume that the binary-valued hidden variable z_i depends on the data point \mathbf{x}_i , we will model it using a logistic function with parameter $\boldsymbol{\phi}$, i.e.,

$$P(z_i = 1 \mid \mathbf{x}_i; \boldsymbol{\phi}) = s(\boldsymbol{\phi}^T \mathbf{x}_i)$$

$$P(z_i = 0 \mid \mathbf{x}_i; \boldsymbol{\phi}) = 1 - s(\boldsymbol{\phi}^T \mathbf{x}_i)$$

where $s(z) = \frac{1}{1 + \exp(-z)}$ is the sigmoid function.

In this problem, we will walk through how the expectation-maximization algorithm can be applied to learn the parameters $\boldsymbol{\theta}_0$, $\boldsymbol{\theta}_1$ and $\boldsymbol{\phi}$. Why might we care about such a model? Intuitively, standard OLS linear regression might not model the following dataset as well as a combination of two linear regression models over different regions of the data domain:

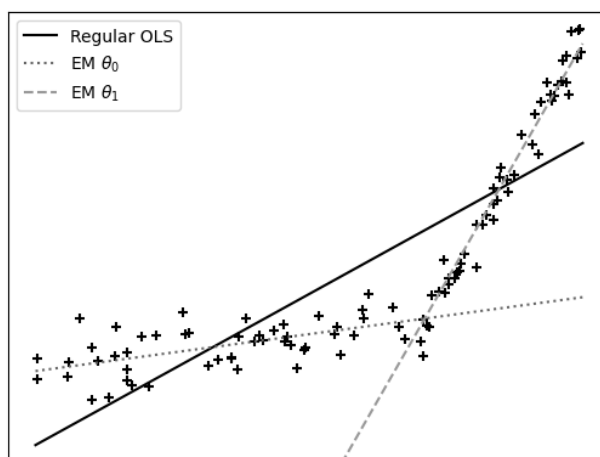


Figure 1: We can see that the mixture model parameters $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$ derived using the EM algorithm fit the dataset much better than the regular OLS solution!

A note on notation: the expression $P(a_1, \dots, a_i | b_1, \dots, b_j; \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_k)$ indicates that the distribution $P(a_1, \dots, a_i | b_1, \dots, b_j)$ is parameterized by $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_k$.

- (a) (3 points) Find the log-likelihood of the parameters given the dataset D , i.e., write down an expression for $l(\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \boldsymbol{\phi}) = \log P(y_1, \dots, y_n \mid \mathbf{x}_1, \dots, \mathbf{x}_n; \boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \boldsymbol{\phi})$. You may leave your expression as a sum of logarithms.

Hint: note that $P(y \mid x) = P(y, z = 0 \mid x) + P(y, z = 1 \mid x)$ by the law of total probability.

Our goal is to find the parameters that maximize the log-likelihood above. Unfortunately, there is no closed-form expression for the log-likelihood above so we cannot use our standard technique of setting its gradient equal to 0. Therefore, we will proceed with the iterative Expectation-Maximization algorithm.

Letting $q_i(k)$ (a shorthand for $q_i(z_i = k)$) be a possible distribution over z_i , we can use Jensen's inequality to lower bound our log-likelihood by

$$\begin{aligned} l(\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \boldsymbol{\phi}) &\geq \mathcal{F}(\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \boldsymbol{\phi}, q) = \sum_{i=1}^n \sum_{k \in \{0,1\}} q_i(k) \log \left(\frac{P(y_i, z_i = k \mid \mathbf{x}_i; \boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \boldsymbol{\phi})}{q_i(k)} \right) \\ &= \sum_{i=1}^n \sum_{k \in \{0,1\}} q_i(k) \log \left(\frac{P(y_i \mid \mathbf{x}_i, z_i = k; \boldsymbol{\theta}_k) P(z_i = k \mid \mathbf{x}_i; \boldsymbol{\phi})}{q_i(k)} \right) \end{aligned}$$

We note that $q_i(0) + q_i(1) = 1$.



initial here

Now, as seen in the homework, we can view the EM algorithm as a coordinate ascent algorithm where we perform the following steps in an alternating order:

$$q^{t+1} = \operatorname{argmax}_q \mathcal{F}(\theta_0^t, \theta_1^t, \phi^t, q) \quad \text{E-step}$$

$$\theta_0^{t+1}, \theta_1^{t+1}, \phi^{t+1} = \operatorname{argmax}_{\theta_0, \theta_1, \phi} \mathcal{F}(\theta_0, \theta_1, \phi, q^{t+1}) \quad \text{M-step}$$

where $t = 0$ represents the initial guesses for the variables above. Note that the $(\cdot)^t$ here does not represent exponentiation but it denotes the value of a variable after the t th iteration of the EM algorithm.

(b) (3 points) We claim that $q_i^{t+1}(k) = P(z_i = k \mid \mathbf{x}_i, y_i; \theta_0^t, \theta_1^t, \phi^t)$ is a valid update for the E-step (note that this is just a generalization of the expectation step from the homework). Find an expression for $q_i^{t+1}(1)$.

Hint: you may find the Bayes rule helpful: $P(z \mid x, y) = P(y \mid x, z)P(z \mid x)/P(y \mid x)$.

Hint: you may the hint for part (a) helpful as well.



(c) We will now derive our M-step updates:

$$\theta_0^{t+1}, \theta_1^{t+1}, \phi^{t+1} = \operatorname{argmax}_{\theta_0, \theta_1, \phi} \mathcal{F}(\theta_0, \theta_1, \phi, q^{t+1})$$

Note that q^{t+1} is the output of the E-step above, and is fixed (in other words, the probabilities $q_i^{t+1}(k)$ for $i = 1, \dots, n$ and $k = 0, 1$ are to be treated as fixed constants) throughout this part. Suppose the data points $\mathbf{x}_i \in \mathbb{R}^d$ are stacked row-wise into the $n \times d$ data matrix X and the labels are stacked into a column vector $\mathbf{y} \in \mathbb{R}^n$. For the sake of convenience, we will assume that X is full rank.



initial here

- (i) (4 points) Find the closed-form solution for θ_0^{t+1} by setting the gradient of \mathcal{F} to $\mathbf{0}$.
Hint: it may be helpful to define the matrix $W_0 = \text{diag}(q_1^{t+1}(0), \dots, q_n^{t+1}(0))$. Note that $q_i^{t+1}(0) \geq 0$ since it is a probability.

- (ii) (4 points) Find the closed-form solution for θ_1^{t+1} by setting the gradient of \mathcal{F} to $\mathbf{0}$.
Hint: it may be helpful to define the matrix $W_1 = \text{diag}(q_1^{t+1}(1), \dots, q_n^{t+1}(1))$. Note that $q_i^{t+1}(1) \geq 0$ since it is a probability.



initial here

- (iii) (4 points) Unlike the updates θ_0^{t+1} and θ_1^{t+1} above, the update for ϕ^{t+1} does not have a closed form solution. So, we can choose to approximate it using some iterations of gradient ascent $\phi \leftarrow \phi + \nabla_{\phi} \mathcal{F}(\theta_0, \theta_1, \phi, q^{t+1})$ instead. Write out the GA update.
Hint: you may use without proof that $s'(z) = s(z)(1 - s(z))$.



initial here

6 Machine Unlearning: Decision Trees

On the midterm, we explored how to remove data from OLS linear regression. Now, we investigate the same problem for decision trees. This problem is largely inspired by Brophy and Lowd’s 2021 ICML paper “Machine Unlearning for Random Forests.”

We have a model that has been trained on features $X \in \mathbb{R}^{n \times d}$ and their targets $y \in \mathbb{R}$. The problem machine unlearning tackles is “removing” the i th datapoint, (X_i, y_i) ; that is, modifying the model such that it was as if the datapoint was not included during training.

For this problem, we assume that we are fitting a decision tree on a binary classification problem with a single continuous feature.

- (a) (3 points) To start, let’s consider a naive strategy for unlearning that is computationally expensive: just retrain the model without the datapoint.

Consider a dataset, where $X = [1, 2, 3, 4, 5, 6, 7]$ and $y = [0, 0, 0, 1, 0, 1, 1]$. First fit a depth-1 decision tree (just one split) to the entire dataset. Then remove the fourth datapoint $(4, 1)$ and refit the tree. **Draw both trees along with their splits including the leaf datapoints.**



- (b) (4 points) Typically, when removing a point from a decision tree, we will need to recompute many of the splits above that point as that point influenced those splits, which can introduce a lot of extra overhead. If we use a less naive unlearning approach than above, we will want to efficiently edit splits. Randomness can help us achieve this goal.

We consider two types of random splits:

- (1) A split where we select the split uniformly at random on the range $[v_{\min}, v_{\max}]$, where v_{\min} and v_{\max} are the minimum and maximum feature values at a given split.
- (2) Let’s call a valid split the mean of two adjacent (if we sorted the data by the feature) datapoints. We then randomly sample (without replacement), k of these split and select the best one.



initial here

If we randomly sample a point to remove, what is the probability that we have to refit the first kind of random split? What about the second? Your answer should be in terms of n' and k .

Assume we have n' points at the split (therefore $n' - 1$ valid splits) and that all splits are equally likely to be the best split of a given subset. For this part only, assume that we don't care if we remove a point that affects a split which was not the best of the k -sample.



- (c) (2 points) Let's now consider how we could check if a split needed to be refit after removing a data point in the second kind of split. To do this, we would remove any splits affected by the removal (ie we remove one of the points used to compute the split boundary) and then resample new splits.

What statistics could we store on each of the type 2 splits that would allow us to determine if our current split became suboptimal after removing a datapoint? You should only use $O(1)$ space per split.

Hint: what do we need to compute the information gain of a split in binary classification?





initial here

(d) (2 points) Which split would you prefer at shallower depths in the tree? Be sure to explain your reasoning.

Hint: see the ‘for only this part’ note in (b).

(e) (3 points) Devise an algorithm (in words) to remove a datapoint from a tree of max depth 1 where the split is of type (2). Assume each split contains a statistic that fulfills the criterion of part (c).

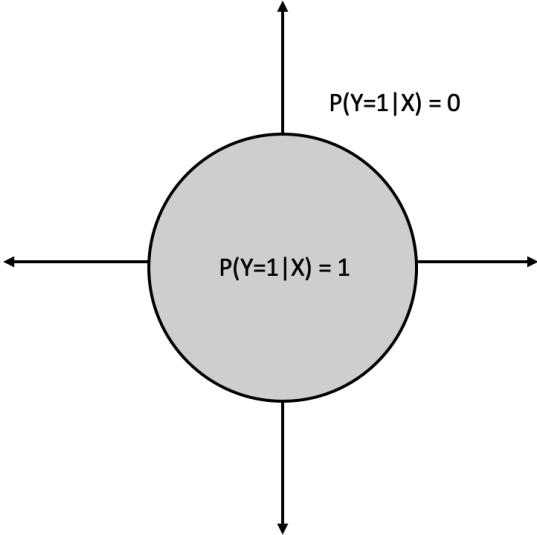
Pretend that this is the function stub you are writing the logic for:

```
def remove_point(root_node: TreeNode, split_type: Union[1, 2], X, y)
```

Hint: don’t forget about leaf nodes.

7 Nearest Neighbors (CS 289A Only)

Consider performing a case of binary classification with feature vectors $\mathbf{x} = [x_1, x_2]^T$, where $x_1, x_2 \in [-1, 1]$. Suppose that we have an estimate for the conditional probability, shown in the figure below:



where the circular region is centered at $(0,0)$ and has a radius of $r < 1$. We will investigate the performance of a nearest neighbor classifier in this problem.

- (a) (2 points) To start, we are interested in training a 1NN classifier. We first uniformly sample n points from the region $[-1, 1] \times [-1, 1]$. Assume that the sampled points' labels exactly follow the estimated conditional distribution in the figure above. Now suppose we use our 1NN classifier to classify a point at $(0, 0)$. Assuming this point comes from the same conditional distribution, what is the probability of misclassification?



initial here


(b) (2 points) What number of points n should we select so that the probability of misclassification is δ for $\delta \in (0, 1)$?

(c) (2 points) Now instead of a 1NN classifier, we will consider the generalized case of a kNN classifier, for some arbitrary odd value k . Suppose that we again sample n points uniformly from the region $[-1, 1] \times [-1, 1]$ and they follow their labels follow the conditional above. Using this kNN classifier, what will the probability of misclassification of point $(0, 0)$ be?



initial here

(d) (4 points) Suppose that instead of the estimated conditional distribution shown in the figure, we instead estimate the conditional $P(Y = 1|X) = 0.8$ inside the circular region, and $P(Y = 1|X) = 0.3$ everywhere else. Once again, we sample n points and then look to classify a point at $(0, 0)$ with a 1NN classifier. Assuming the sampled and test point labels both follow this new conditional distribution, what is the probability of misclassification of this point? Hint: consider separately all possible scenarios for misclassification.





initial here

You may use this page to show extra work. Clearly mark your work with the problem number here, and also mention in the problem-specific box that your work is continued here.



initial here

You may use this page to show extra work. Clearly mark your work with the problem number here, and also mention in the problem-specific box that your work is continued here.



initial here

You may use this page to show extra work. Clearly mark your work with the problem number here, and also mention in the problem-specific box that your work is continued here.



initial here

You may use this page to show extra work. Clearly mark your work with the problem number here, and also mention in the problem-specific box that your work is continued here.