

- Please do not open the exam before you are instructed to do so.
- The exam is closed book, closed notes except your two page cheat sheet.
- **Electronic devices are forbidden on your person**, including cell phones, iPods, headphones, and laptops. Turn your cell phone off and **leave all electronics at the front of the room**, or **risk getting a zero** on the exam.
- You have 3 hours.
- Please write your initials at the top right of each odd-numbered page (e.g., write “JS” if you are Jonathan Shewchuk). Finish this by the end of your 3 hours.
- Mark your answers on the exam itself in the space provided. Do **not** attach any extra sheets.
- The total number of points is 150. There are 26 multiple choice questions worth 3 points each, and 7 written questions worth a total of 72 points.
- For multiple answer questions, fill in the bubbles for **ALL correct choices**: there may be more than one correct choice, but there is always at least one correct choice. **NO partial credit** on multiple answer questions: the set of all correct answers must be checked.

First name	
Last name	
SID	
First and last name of student to your left	
First and last name of student to your right	

Q1. [78 pts] Multiple Answer

Fill in the bubbles for **ALL correct choices**: there may be more than one correct choice, but there is always at least one correct choice. **NO partial credit**: the set of all correct answers must be checked.

(1) [3 pts] Which of the following are NP-hard problems? Let $X \in \mathbb{R}^{n \times d}$ be a design matrix, let $y \in \mathbb{R}^n$ be a vector of labels, let L be the Laplacian matrix of some n -vertex graph, and let $\mathbf{1} = [1 \ 1 \ \dots \ 1]^\top$.

- $\min_{\mu, y} \sum_{i=1}^k \sum_{y_j=i} |X_j - \mu_i|^2$ where each μ_i is the mean of sample points assigned class i
- $\min_{y \in \mathbb{R}^n} \frac{1}{4} y^\top L y$ subject to $|y|^2 = n; \mathbf{1}^\top y = 0$
- $\min_y \sum_{i=1}^k \sum_{y_j=i} |X_j - \mu_i|^2$ with each μ_i fixed
- $\min_{y \in \mathbb{R}^n} \frac{1}{4} y^\top L y$ subject to $\forall j, y_j \in \{-1, +1\}; \mathbf{1}^\top y = 0$

(2) [3 pts] Which clustering algorithms permit you to decide the number of clusters after the clustering is done?

- k -means clustering
- a k -d tree used for divisive clustering
- agglomerative clustering with single linkage
- spectral graph clustering with 3 eigenvectors

(3) [3 pts] For which of the following does normalizing your input features influence the predictions?

- decision tree (with usual splitting method)
- neural network
- Lasso
- soft-margin support vector machine

(4) [3 pts] With the SVD, we write $X = UDV^\top$. For which of the following matrices are the eigenvectors the columns of U ?

- $X^\top X$
- XX^\top
- $X^\top XX^\top X$
- $XX^\top XX^\top$

(5) [3 pts] Why is PCA sometimes used as a preprocessing step before regression?

- To reduce overfitting by removing poorly predictive dimensions.
- To make computation faster by reducing the dimensionality of the data.
- To expose information missing from the input data.
- For inference and scientific discovery, we prefer features that are not axis-aligned.

(6) [3 pts] Consider the matrix $X = \sum_{i=1}^r \alpha_i u_i v_i^\top$ where each α_i is a scalar and each u_i and v_i is a vector. It is possible that the rank of X might be

- $r + 1$
- $r - 1$
- r
- 0

(7) [3 pts] Why would we use a random forest instead of a decision tree?

- For lower training error.
- To better approximate posterior probabilities.
- To reduce the variance of the model.
- For a model that is easier for a human to interpret.

(8) [3 pts] What tends to be true about increasing the k in k -nearest neighbors?

- The decision boundary tends to get smoother.
- The bias tends to increase.
- The variance tends to increase.
- As the number of sample points approaches infinity (with $n/k \rightarrow \infty$), the error rate approaches less than twice the Bayes risk (assuming training and test points are drawn independently from the same distribution).

(9) [3 pts] Which of the following statements are true about the entropy of a discrete probability distribution?

- It is a useful criterion for picking splits in decision trees.
- It is maximized when the probability distribution is uniform.
- It is a convex function of the class probabilities.
- It is minimized when the probability distribution is uniform.

(10) [3 pts] A low-rank approximation of a matrix can be useful for

- removing noise.
- filling in unknown values.
- discovering latent categories in the data.
- matrix compression.

(11) [3 pts] Let L be the Laplacian matrix of a graph with n vertices. Let

$$\beta = \min_{\substack{y \in \mathbb{R}^n \\ \forall i, y_i \in \{-1, +1\} \\ \mathbf{1}^\top y = 0}} y^\top L y \quad \text{and} \quad \gamma = \min_{\substack{y \in \mathbb{R}^n \\ |y|^2 = n \\ \mathbf{1}^\top y = 0}} y^\top L y.$$

Which of the following are true for every Laplacian matrix L ?

- $\beta \geq \gamma$
- $\beta \leq \gamma$
- $\beta > \gamma$
- $\beta < \gamma$

(12) [3 pts] Which of the following are true about decision trees?

- They can be used only for classification.
- All the leaves must be pure.
- The tree depth never exceeds $O(\log n)$ for n sample points.
- Pruning usually achieves better test accuracy than stopping early.

(13) [3 pts] Which of the following is an effective way of reducing overfitting in neural networks?

- Augmenting the training data with similar synthetic examples
- Increasing the number of layers
- Weight decay (i.e., ℓ_2 regularization)
- Dropout

(14) [3 pts] If the VC dimension of a hypothesis class \mathcal{H} is an integer $D < \infty$ (i.e., $\text{VC}(\mathcal{H}) = D$), this means

- there exists some set of D points shattered by \mathcal{H} .
- no set of $D + 1$ points is shattered by \mathcal{H} .
- all sets of D points are shattered by \mathcal{H} .
- $\Pi_{\mathcal{H}}(D) = 2^D$.

- (15) [3 pts] Consider the minimizer w^* of the ℓ_2 -regularized least squares objective $J(w) = \|Xw - y\|^2 + \lambda \|w\|^2$ with $\lambda > 0$. Which of the following are true?
- $Xw^* = y$
 - w^* exists if and only if $X^T X$ is nonsingular
 - $w^* = X^+ y$, where X^+ is the pseudoinverse of X
 - The minimizer w^* is unique
- (16) [3 pts] You are training a neural network, but the training error is high. Which of the following, if done in isolation, has a better-than-tiny chance of reducing the training error?
- Adding another hidden layer
 - Adding more units to hidden layers
 - Normalizing the input data
 - Training on more data
- (17) [3 pts] Filters in the **late** layers of a convolutional neural network designed to classify objects in photographs likely represent
- edge detectors.
 - concepts such as “this image contains wheels.”
 - concepts such as “there is an animal.”
 - concepts such as “Jen is flirting with Dan.”
- (18) [3 pts] Which of the following techniques usually speeds up the training of a sigmoid-based neural network on a classification task?
- Using batch descent instead of stochastic
 - Having a good initialization of the weights
 - Increasing the learning rate with every iteration
 - Using the cross-entropy loss instead of the mean squared error
- (19) [3 pts] In a soft-margin support vector machine, decreasing the slack penalty term C causes
- more overfitting.
 - a smaller margin.
 - less overfitting.
 - less sensitivity to outliers.
- (20) [3 pts] The shortest distance from a point z to a hyperplane $w^T x = 0$ is
- $w^T z$
 - $\frac{w^T z}{|w|^2}$
 - $\frac{w^T z}{|w|}$
 - $|w| \cdot |z|$
- (21) [3 pts] The Bayes decision rule
- does the best a classifier can do, in expectation
 - chooses the class with the greatest posterior probability, if we use the 0-1 risk function
 - can be computed exactly from a large sample
 - minimizes the risk functional
- (22) [3 pts] Which of the following are techniques commonly used in training neural nets?
- linear programming
 - Newton’s method
 - backpropagation
 - cross-validation





(23) [3 pts] Which of these statements about learning theory are correct?

- The VC dimension of halfplanes is 3.
- For a fixed set of training points, the more dichotomies Π we have, the higher the probability that the training error is close to the true risk.
- The VC dimension of halfspaces in 3D is ∞ .
- For a fixed hypothesis class \mathcal{H} , the more training points we have, the higher the probability that the training error is close to the true risk.

(24) [3 pts] Which of the following statements are true for a design matrix $X \in \mathbb{R}^{n \times d}$ with $d > n$? (The rows are n sample points and the columns represent d features.)

- Least-squares linear regression computes the weights $w = (X^T X)^{-1} X^T y$.
- X has exactly $d - n$ eigenvectors with eigenvalue zero.
- The sample points are linearly separable.
- At least one principal component direction is orthogonal to a hyperplane that contains all the sample points.

(25) [3 pts] Which of the following visuals accurately represent the clustering produced by greedy agglomerative hierarchical clustering with centroid linkage on the set of feature vectors $\{(-2, -2), (-2, 0), (1, 3), (2, 2), (3, 4)\}$?



(26) [3 pts] Which of the following statements is true about the standard k -means clustering algorithm?

- The random partition initialization method usually outperforms the Forgy method.
- After a sufficiently large number of iterations, the clusters will stop changing.
- It is computationally infeasible to find the optimal clustering of $n = 15$ points in $k = 3$ clusters.
- You can use the metric $d(x, y) = \frac{x \cdot y}{|x| \cdot |y|}$.

Q2. [9 pts] A Miscellany

- (a) [3 pts] Consider a convolutional neural network for reading the handwritten MNIST letters, which are 28×28 images. Suppose the first hidden layer is a convolutional layer with 20 different 5×5 filters, applied to the input image with a stride of 1 (i.e., every filter is applied to *every* 5×5 patch of the image, with patches allowed to overlap). Each filter has a bias weight. How many weights (parameters) does this layer use?

$$20 \times (5 \times 5 + 1) = 520.$$

- (b) [3 pts] Let X be an $n \times d$ design matrix representing n sample points in \mathbb{R}^d . Let $X = UDV^T$ be the singular value decomposition of X . We stated in lecture that row i of the matrix UD gives the coordinates of sample point X_i in principal coordinates space, i.e., $X_i \cdot v_j$ for each j , where X_i is the i th row of X and v_j is the j th column of V . Show that this is true.

As V is an orthogonal matrix, we can write $XV = UDV^T V = UD$.

By the definition of matrix multiplication, $(UD)_{ij} = X_i \cdot v_j$.

- (c) [3 pts] Let $x, y \in \mathbb{R}^d$ be two points (e.g., sample or test points). Consider the function $k(x, y) = x^T \text{rev}(y)$ where $\text{rev}(y)$ reverses the order of the components in y . For example, $\text{rev} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} = \begin{pmatrix} 3 \\ 2 \\ 1 \end{pmatrix}$. Show that k cannot be a valid kernel function.

Hint: remember how the kernel function is defined, and show a simple two-dimensional counterexample.

We have that $k((-1, 1), (-1, 1)) = -2$, but this is impossible as, if k is a valid kernel, then there is some function Φ such that $k(x, x) = \Phi(x)^T \Phi(x) \geq 0$.

Q3. [10 pts] Maximum Likelihood Estimation for Reliability Testing

Suppose we are reliability testing n units taken randomly from a population of identical appliances. We want to estimate the mean failure time of the population. We assume the failure times come from an exponential distribution with parameter $\lambda > 0$, whose probability density function is $f(x) = \lambda e^{-\lambda x}$ (on the domain $x \geq 0$) and whose cumulative distribution function is $F(x) = \int_0^x f(x) dx = 1 - e^{-\lambda x}$.

- (a) [6 pts] In an ideal (but impractical) scenario, we run the units until they all fail. The failure times are t_1, t_2, \dots, t_n .

Formulate the likelihood function $\mathcal{L}(\lambda; t_1, \dots, t_n)$ for our data. Then find the maximum likelihood estimate $\hat{\lambda}$ for the distribution's parameter.

$$\begin{aligned} \mathcal{L}(\lambda; t_1, \dots, t_n) &= \prod_{i=1}^n f(t_i) = \prod_{i=1}^n \lambda e^{-\lambda t_i} = \lambda^n e^{-\lambda \sum_{i=1}^n t_i} \\ \ln \mathcal{L}(\lambda) &= n \ln \lambda - \lambda \sum_{i=1}^n t_i \\ \frac{\partial}{\partial \lambda} \ln \mathcal{L}(\lambda) &= \frac{n}{\lambda} - \sum_{i=1}^n t_i = 0 \\ \hat{\lambda} &= \frac{n}{\sum_{i=1}^n t_i} \end{aligned}$$

- (b) [4 pts] In a more realistic scenario, we run the units for a fixed time T . We observe r unit failures, where $0 \leq r \leq n$, and there are $n - r$ units that survive the entire time T without failing. The failure times are t_1, t_2, \dots, t_r .

Formulate the likelihood function $\mathcal{L}(\lambda; n, r, t_1, \dots, t_r)$ for our data. Then find the maximum likelihood estimate $\hat{\lambda}$ for the distribution's parameter.

Hint 1: What is the probability that a unit will not fail during time T ? *Hint 2:* It is okay to define $\mathcal{L}(\lambda)$ in a way that includes contributions (densities and probability masses) that are not commensurate with each other. Then the constant of proportionality of $\mathcal{L}(\lambda)$ is meaningless, but that constant is irrelevant for finding the best-fit parameter $\hat{\lambda}$. *Hint 3:* If you're confused, for part marks write down the likelihood that r units fail and $n - r$ units survive; then try the full problem. *Hint 4:* If you do it right, $\hat{\lambda}$ will be the number of observed failures divided by the sum of unit test times.

$$\begin{aligned} \mathcal{L}(\lambda; n, r, t_1, \dots, t_r) &\propto \left(\prod_{i=1}^r f(t_i) \right) (1 - F(T))^{n-r} \\ &= \left(\prod_{i=1}^r \lambda e^{-\lambda t_i} \right) (e^{-\lambda T})^{n-r} \\ &= \lambda^r e^{-\lambda \sum_{i=1}^r t_i} e^{-\lambda(n-r)T} \\ \ln \mathcal{L}(\lambda) &= r \ln \lambda - \lambda \sum_{i=1}^r t_i - \lambda(n-r)T + \text{constant} \\ \frac{\partial}{\partial \lambda} \ln \mathcal{L}(\lambda) &= \frac{r}{\lambda} - \sum_{i=1}^r t_i - (n-r)T = 0 \\ \hat{\lambda} &= \frac{r}{\sum_{i=1}^r t_i + (n-r)T} \end{aligned}$$

Q4. [10 pts] Decision Trees

Consider the design matrix

$$\begin{bmatrix} 4 & 6 & 9 & 1 & 7 & 5 \\ 1 & 6 & 5 & 2 & 3 & 4 \end{bmatrix}^T$$

representing 6 sample points, each with two features f_1 and f_2 .

The labels for the data are

$$\begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 0 \end{bmatrix}^T$$

In this question, we build a decision tree of depth 2 by hand to classify the data.

- (a) [2 pts] What is the entropy at the root of the tree?

$$-0.5 \log_2 0.5 - 0.5 \log_2 0.5 = 1$$

- (b) [3 pts] What is the rule for the first split? Write your answer in a form like $f_1 \geq 4$ or $f_2 \geq 3$. Hint: you should be able to eyeball the best split without calculating the entropies.

If we sort by f_1 , the features and the corresponding labels are

$$\begin{bmatrix} 1 & 4 & 5 & 6 & 7 & 9 \\ 0 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}.$$

If we sort by f_2 , we have

$$\begin{bmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 1 & 0 & 1 & 0 & 1 & 0 \end{bmatrix}.$$

The best split is $f_1 \geq 7$.

- (c) [3 pts] For each of the two treenodes after the first split, what is the rule for the second split?

For the treenode with labels (1, 1), there's no need to split again.

For the treenode with labels (0, 1, 0, 0), if we sort by f_1 , we have

$$\begin{bmatrix} 1 & 4 & 5 & 6 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$

If we sort using f_2 , we get

$$\begin{bmatrix} 1 & 2 & 4 & 6 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

We easily see we should choose $f_2 \geq 2$.

- (d) [2 pts] Let's return to the root of the tree, and suppose we're incompetent tree builders. Is there a (not trivial) split at the root that would have given us an information gain of zero? Explain your answer.

Yes. The rules $f_1 \geq 5$, $f_2 \geq 3$, or $f_2 \geq 5$ would all fail to reduce the weighted average entropy below 1.

Q5. [11 pts] Bagging and Random Forests

We are building a random forest for a 2-class classification problem with t decision trees and bagging. The input is a $n \times d$ design matrix X representing n sample points in \mathbb{R}^d (quantitative real-valued features only). For the i th decision tree we create an n -point training set $X^{(i)}$ through standard bagging. At each node of each tree, we randomly select k of the features (this random subset is selected independently for each treenode) and choose the single-feature split that maximizes the information gain, compared to all possible single-feature splits on those k features. Assume that we can radix sort real numbers in linear time, and we can randomly select an item from a set in constant time.

- (a) [3 pts] Remind us how bagging works. How do we generate the data sets $X^{(i)}$? What do we do with duplicate points?

For each training set $X^{(i)}$, we select n sample points from X uniformly at random **with replacement**. Duplicate points have proportionally greater weight in the entropy (or other cost function) calculations. (We will accept an answer that states that duplicate points are treated as if they were separate points infinitesimally close together.)

- (b) [3 pts] Fill in the blanks to derive the overall running time to construct a random forest with bagging and random subset selection. Let h be the height/depth (they're the same thing) of the tallest/deepest tree in the forest. You must use the tightest bounds possible with respect to n , d , t , k , h , and n' .

Consider choosing the split at a treenode whose box contains n' sample points. We can choose the best split for these n' sample points in $O(\underline{\hspace{2cm}})$ time. Therefore, the running time **per sample point in that node** is $O(\underline{\hspace{2cm}})$.

Each sample point in $X^{(i)}$ participates in at most $O(\underline{\hspace{2cm}})$ treenodes, so each sample point contributes at most $O(\underline{\hspace{2cm}})$ to the time. Therefore, the total running time for one tree is $O(\underline{\hspace{2cm}})$.

We have t trees, so the total running time to create the random forest is $O(\underline{\hspace{2cm}})$.

The blanks in order: $O(n'k)$, $O(k)$, $O(h)$, $O(kh)$, $O(nkh)$, $O(nkht)$. They are each worth half a point.

- (c) [2 pts] If we instead use a support vector machine to choose the split in each treenode, how does that change the asymptotic **query** time to classify a test point?


It slows queries down by a factor of $\Theta(d)$ or $\Theta(k)$ (depending whether you run the SVM on k features or all d features—we'll accept either interpretation), because it is necessary to inspect all d (or k) features of the query point at each treenode.

- (d) [3 pts] Why does bagging by itself (without random subset selection) tend **not** to improve the performance of decision trees as much as we might expect?

It is common that the same few features tend to dominate in all of the subsets, so almost all the trees will tend to have very similar early splits, and therefore all the trees will produce very similar estimates. The models are not decorrelated enough.

Q6. [11 pts] One-Dimensional ConvNet Backprop

Consider this convolutional neural network architecture.



In the first layer, we have a one-dimensional convolution with a single filter of size 3 such that $h_i = s\left(\sum_{j=1}^3 v_j x_{i+j-1}\right)$. The second layer is fully connected, such that $z = \sum_{i=1}^4 w_i h_i$. The hidden units' activation function $s(x)$ is the logistic (sigmoid) function with derivative $s'(x) = s(x)(1 - s(x))$. The output unit is linear (no activation function). We perform gradient descent on the loss function $R = (y - z)^2$, where y is the training label for x .

- (a) [1 pt] What is the total number of parameters in this neural network? Recall that convolutional layers share weights. There are no bias terms.

The answer is 7. There are 3 parameters in layer 1 and 4 parameters in layer 2.

- (b) [4 pts] Compute $\partial R / \partial w_i$.

$$\frac{\partial R}{\partial w_i} = -2(y - z)h_i$$

- (c) [1 pt] Vectorize the previous expression—that is, write $\partial R / \partial w$.


$$\frac{\partial R}{\partial w} = -2(y - z)h$$

- (d) [5 pts] Compute $\partial R / \partial v_j$.

$$\frac{\partial R}{\partial v_j} = -2(y - z) \frac{\partial z}{\partial v_j} = -2(y - z) \sum_{i=1}^4 \frac{\partial z}{\partial h_i} \frac{\partial h_i}{\partial v_j} = -2(y - z) \sum_{i=1}^4 w_i h_i (1 - h_i) x_{i+j-1}$$

Q7. [11 pts] Spectral Graph Partitioning

(a) [3 pts] Write down the Laplacian matrix L_G of the following graph G . Every edge has weight 1.



$$\begin{bmatrix} 2 & -1 & -1 & 0 & 0 & 0 & 0 \\ -1 & 2 & -1 & 0 & 0 & 0 & 0 \\ -1 & -1 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 1 \end{bmatrix}$$

(b) [2 pts] Find three orthogonal eigenvectors of L_G , all having eigenvalue 0.

$$x = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, y = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 1 \end{bmatrix}, z = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}$$

(c) [2 pts] Use two of those three eigenvectors (it doesn't matter which two) to assign each vertex of G a *spectral vector* in \mathbb{R}^2 . Draw these vectors in the plane, and explain how they partition G into three clusters. (Optional alternative: if you can draw 3D figures well, you are welcome to use all three eigenvectors and assign each vertex a spectral vector in \mathbb{R}^3 .)

The eigenvectors x and y give the embedding

$$1, 2, 3 \mapsto \begin{bmatrix} 1 \\ 0 \end{bmatrix}, 4, 7 \mapsto \begin{bmatrix} 0 \\ 1 \end{bmatrix}, 5, 6 \mapsto \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

Each of the three clusters is mapped to a single point in \mathbb{R}^2 .

(d) [3 pts] Let K_n be the complete graph on n vertices (every pair of vertices is connected by an edge of weight 1) and let L_{K_n} be its Laplacian matrix. The eigenvectors of L_{K_n} are $v_1 = \mathbf{1} = [1 \ \dots \ 1]^T$ and every vector that is orthogonal to $\mathbf{1}$. What are the eigenvalues of L_{K_n} ?

$$\lambda_1 = 0 \text{ and } \lambda_2 = \dots = \lambda_n = n.$$

(e) [1 pt] What property of these eigenvalues gives us a hint that the complete graph does not have any good partitions?

λ_2 is large, so there is no low-sparsity cut. (The optimal cut has sparsity $\geq \lambda_2/2$.)

Q8. [10 pts] We Hope You Learned This

Consider learning closed intervals on the real line. Our hypothesis class \mathcal{H} consists of all intervals of the form $[a, b]$ where $a < b$ and $a, b \in \mathbb{R}$. We interpret an interval (hypothesis) $[a, b] \in \mathcal{H}$ as a classifier that identifies a point x as being in class C if $a \leq x \leq b$, and identifies x as not being in class C if $x < a$ or $x > b$.

- (a) [2 pts] Consider a set containing two distinct points on the real line. Which such sets can be shattered by \mathcal{H} ?

All sets of two distinct points can be shattered.

- (b) [2 pts] Show that no three points can be shattered by \mathcal{H} .

Let $X = \{x_1, x_2, x_3\}$ with $x_1 \leq x_2 \leq x_3$. No interval can contain x_1 and x_3 without containing x_2 . (I.e., suppose x_1 and x_3 are in class C, but x_2 is not.)

- (c) [2 pts] Write down the shatter function $\Pi_{\mathcal{H}}(n)$. Explain your answer.

$$\Pi_{\mathcal{H}}(n) = \binom{n}{2} + n + 1 = \frac{n^2 + n}{2} + 1.$$

There are $\binom{n}{2}$ dichotomies with two or more points (imagine choosing the first and the last sample point in class C), n dichotomies with one point, and one dichotomy with no points.

- (d) [2 pts] Consider another hypothesis class \mathcal{H}_2 . Each hypothesis in \mathcal{H}_2 is a union of two intervals. \mathcal{H}_2 is the set of all such hypotheses (i.e., every union of two intervals on the number line). For example, $[3, 7] \cup [8.5, 10] \in \mathcal{H}_2$; that's the set of all points x such that $3 \leq x \leq 7$ or $8.5 \leq x \leq 10$.

What is the largest number of distinct points that \mathcal{H}_2 can shatter? Explain why no larger number can be shattered.

Four. If you have five distinct points, \mathcal{H}_2 cannot include the first, third, and fifth points while excluding the second and fourth.

- (e) [2 pts] Which hypothesis class has a greater sample complexity, \mathcal{H} or \mathcal{H}_2 ? Explain why.

\mathcal{H}_2 , because its shatter function grows faster (quartic rather than quadratic) and its VC dimension is greater.