FoDA L29

Semester Review

Dec 8, 2022

supervised

input (x, y) ← labels

unsupervised

input $x \in \mathbb{R}^d$

output
real values
$\mathbb{R}$

regression

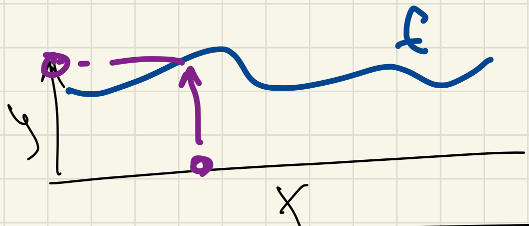

dimensionality Reduction

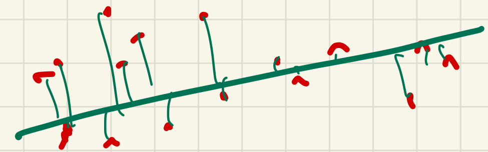

output
set

classification

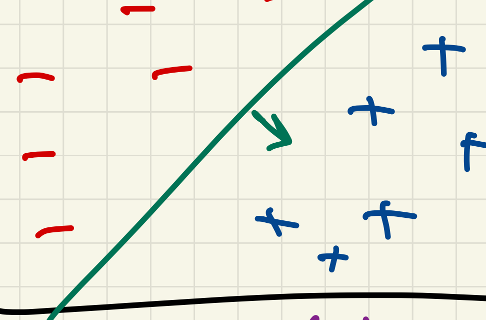

clustering

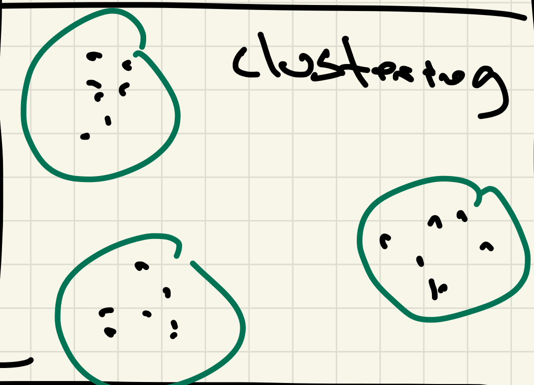

cross-validation

1. Consider the random variables X and Y described by the joint probability table.

	$X=1$	$X=2$	$X=3$
$Y=1$	0.10	0.05	0.10
$Y=2$	0.30	0.25	0.20

0.4 0.3 0.3

Derive the following values

(a) $\mathbf{Pr}(X=1)$ $0.1+0.3=0.4$

(b) $\mathbf{Pr}(X=2 \cap Y=1)$ $= 0.05$

(c) $\mathbf{Pr}(X=3 \mid Y=2)$ $= \frac{0.2}{0.3+0.25+0.2} = \frac{0.2}{0.75}$

Compute the following probability distributions.

(d) What is the marginal distribution for X?

(e) What is the conditional probability for Y, given that $X=2$?

$Y=1$ $\frac{0.05}{0.3}$

$Y=2$ $\frac{0.25}{0.3}$

Answer the following question about the joint distribution.

(f) Are random variables X and Y independent?

$\Pr(X=2)\cdot\Pr(Y=1) = \Pr(X=2, Y=1)$

$0.3 \cdot 0.25 \ ? \ 0.05$

$\frac{3}{10}\cdot\frac{1}{4} = \frac{3}{40} \neq \frac{2}{40}$

No.

(g) Is $\mathbf{Pr}(X=1)$ independent of $\mathbf{Pr}(Y=1)$?

0.4 $\quad 0.25$

$(0.4)(0.25) = \frac{1}{10} = \Pr(X=1, Y=1)$

$\frac{4}{10}\cdot\frac{1}{4} = \frac{1}{10}$

Yes.

2. Consider two models M_1 and M_2, where from prior knowledge we believe that $\mathbf{Pr}(M_1) = 0.25$ and $\mathbf{Pr}(M_2) = 0.75$. We then observe a data set D. Given each model we assess the likelihood of seeing that data given the model as $\mathbf{Pr}(D \mid M_1) = 0.5$ and $\mathbf{Pr}(D \mid M_2) = 0.01$. Now that we have the data, which model is has a higher probability of being correct?

Bayes Rule

$$\Pr(M \mid D) = \frac{\Pr(D \mid M) \cdot \Pr(M)}{\cancel{\Pr(D)}}$$

$$\Pr(M_1 \mid D) = \frac{\Pr(D \mid M_1) \cdot \Pr(M_1)}{\Pr(D)} = \frac{(0.5) \cdot (0.25)}{\Pr(D)} = \frac{0.125}{\Pr(D)}$$

$= 0.125/c$

larger

$$\Pr(M_2 \mid D) = \ldots = \frac{(0.01)(0.75)}{\Pr(D)} \quad \frac{0.0075}{\Pr(D)}$$

$= 0.0075/c$

3. Assume I observe 3 data points x_1, x_2, and x_3 drawn iid from an unknown distribution. Given a model M, I can calculate the likelihood this each data point as $\mathbf{Pr}(x_1 \mid M) = 0.5$, $\mathbf{Pr}(x_2 \mid M) = 0.1$, and $\mathbf{Pr}(x_3 \mid M) = 0.2$. What is the likelihood of seeing all of these data points, given the model M: $\mathbf{Pr}(x_1, x_2, x_3 \mid M)$?

independently identically distributed

$$\Pr(x_1, x_2, x_3 \mid M) = \underset{(0.5)}{\Pr(x_1 \mid M)} \cdot \underset{(0.1)}{\Pr(x_2 \mid M)} \cdot \underset{(0.2)}{\Pr(x_3 \mid M)}$$

$$= 0.0$$

4. Consider a pdf f so that a random variable $X \sim f$ has expected value $\mathbf{E}[X] = 3$ and variance $\mathbf{Var}[X] = 10$. Now consider $n = 10$ iid random variables $X_1, X_2, \ldots, X_{10}$ drawn from f. Let $\bar{X} = \frac{1}{10}\sum_{i=1}^{10} X_i$.

(a) What is $\mathbf{E}[\bar{X}]$? $= 3$

(b) What is $\mathbf{Var}[\bar{X}]$? $= \frac{Var[X]}{10} = 1$

(c) What is the standard deviation of $\bar{X}$? $\sqrt{Var[\bar{X}]} = \sqrt{1} = 1$

(d) Which is larger $\mathbf{Pr}[X > 4]$ or $\mathbf{Pr}[\bar{X} > 4]$?

(e) Which is larger $\mathbf{Pr}[X > 2]$ or $\mathbf{Pr}[\bar{X} > 2]$?

smaller $Pr[X \leq 2]$ or $Pr[\bar{X} \leq 2]$

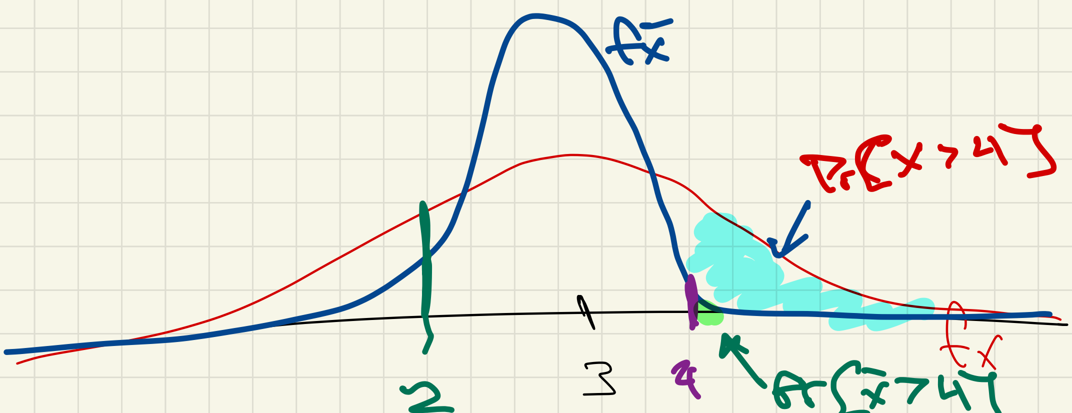

Central Limit Thm

- $E[\bar{X}] = E[X]$
- $Var[\bar{X}] = \frac{Var[X]}{10 = n}$
- in limit $n \to \infty$
  $\to$ Normal

5. Let X be a random variable that you know is in the range $[-1, 2]$ and you know has expected value of $\mathbf{E}[X] = 0$. Use the Markov Inequality to upper bound $\mathbf{Pr}[X > 1.5]$?
(Hint: you will need to use a change of variables.)

$Z' = X + 2$

$Z = X + 1$

$Z = [0, 3]$

Linearity of Expectation

$E[Z] = \boxed{E[X+1] = 0 + 1}$
$= 1$

$E[Z'] = 2$

MI

R.V. Z

- $Z \geq 0$
- $E[Z]$

$$\boxed{\Pr[Z > \alpha] < \frac{E[Z]}{\alpha}}$$

$\alpha = 2.5$

$\Pr[X > 1.5] = \Pr[Z > 2.5]$

$+1 \quad +1$

$E[Z]$

$= \Pr[Z' > 3.5] \leq \frac{1}{2.5} = 0.4$

$\leq \frac{2}{3.5} = \frac{4}{7} = 0.57$

α

6. Consider a matrix

$$A = \begin{bmatrix} 2 & 2 & 3 \\ -2 & 7 & 4 \\ -3 & -3 & -4.5 \\ -8 & 2 & 3 \end{bmatrix}.$$

$a_1 \cdot (-1.5)$

0.003
4.26
π
e

$A = USV^T$

$A^TA = US^2U^T$

(a) Add a column to A so that it is invertable. square + full rank

columns lin indep.

(b) Remove a row from A so that it is invertable. a_1 or a_3

(c) Is AA^T invertable? No

(d) Is A^TA invertable? Yes

$A \in \mathbb{R}^{4\times 3}$

$AA^T = (4\times 3)(3\times 4) \in \mathbb{R}^{4\times 4}$

rank ≤ 3

$A^TA \in \mathbb{R}^{3\times 3}$

$\hookrightarrow$ full rank

7. Consider two vectors $u = (0.5, 0.4, 0.4, 0.5, 0.1, 0.4, 0.1)$ and $v = (-1, -2, 1, -2, 3, 1, -5)$.

(a) Check if u or ~~v~~ is a unit vector. u not v $\|r\|_2 = 1$

(b) Calculate the dot product $\langle u, v \rangle$. $= 0$

(c) Are u and v orthogonal? Yes

$$1^2 = \|u\|_2^2 = 0.5^2 + 0.4^2 + \dots + 0.1^2$$

0.25 0.16

$$= 1$$

$$\langle u, v \rangle = \sum_{j=1}^{d} u_j v_j = (0.5)(-1) + (0.4)(-2) + \dots$$

$$= 0$$

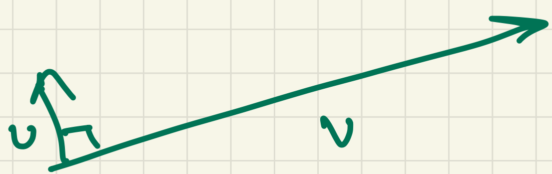

8. Consider a matrix $A \in \mathbb{R}^{n\times 4}$. Each row represents a customer (there are n customers in the database). The first column is the age of the customer in years, the second column is the number of days since the customer entered the database, the third column is the total cost of all purchases ever by the customer in dollars, and the last column is the total profit in dollars generated by the customer.

different units

For each of the following operations, decide if it is **reasonable** or **unreasonable**.

(a) Run simple linear regression using the first three columns to build a model to predict the fourth column. $\langle x, x\rangle$ reasonable

(b) Use k-means clustering to group the customers into 4 types using Euclidean distance between rows as the distance. UnR

(c) Use PCA to find the best 2-dimensional subspace, so we can draw the customers in a $\mathbb{R}^2$ in way that has the least projection error. Un Reason

(d) Use the linear classification to build a model based on the first three columns to predict if the customer will make a profit $+1$ or not -1.

ok w/ Perceptron

not ok w/ Loss fxn

miles	42
educ	1000
pop	700
	1742

9. Consider a data set (X, y) where $X \in \mathbb{R}^{n\times 3}$ we decompose into a test and a training data set $(X_{\text{train}}, y_{\text{train}})$. Assume that X_{train} is not just a subset of X, but also pads prepends a columns of all 1s. We build a linear model

$$\alpha = (X_{\text{train}}^T X_{\text{train}})^{-1} X_{\text{train}}^T y_{\text{train}}.$$

where $\alpha \in \mathbb{R}^4$. The remaining two testing data points are (x_1, y_1) and (x_2, y_2), where $x_1, x_2 \in \mathbb{R}^3$. Explain (write a mathematical expression) to use this test data to estimate the generalization error. That is, if one new data point arrives x, how much squared error would we expect the model α to have compared to the unknown true value y?

$$M_\alpha(x) = \langle \alpha, (1, x) \rangle$$

$$\frac{(M_\alpha(x_1) - y_1)^2 + (M_\alpha(x_2) - y_2)^2}{2}$$

10. Consider a function $f(x, y)$ with gradient $\nabla f(x, y) = \boxed{(x-1, 2y+x)}$. Starting at a value $(x = 1, y = 2)$, and a learning rate of $\gamma = 1$, execute one step of gradient descent.

GD

$\alpha = (x=1, y=2)$

$\alpha \leftarrow \alpha - \gamma \nabla f(x,y)$

$(1,2) - (1)(x-1, 2y+x)$

$(1,2) - (1)(0, 5)$

$(1,2) - (0,5)$

$= (1, -3)$

11. Consider running gradient descent with a fixed learning rate γ. For each of the following, we plot the function value over 10 steps (the function is different each time). Decide whether the learning rate is probably **too high**, **too low**, or **about right**.

(a) f_1: 100, 99, 98, 97, 96, 95, 94, 93, 92, 91 too low

(b) f_2: 100, 50, 75, 60, 65, 45, 75, 110, 90, 85 too high

(c) f_3: 100, 80, 65, 50, 40, 35, 31, 29, 28, 27.5, 27.3 just right

(d) f_4: 100, 80, 60, 40, 20, 0, -20, -40, -60, -80, -100 too low

12. Consider a matrix $A \in \mathbb{R}^{8\times 4}$ with squared singular values $\sigma_1^2 = 10$, $\sigma_2^2 = 5$, $\sigma_3^2 = 2$, and $\sigma_4^2 = 1$.

(a) What is the rank of A? 4

(b) What is $\|A - A_2\|_F^2$, where A_2 is the best rank-2 approximation of A. $\sigma_3^2 + \sigma_4^2 = 2+1=3$

(c) What is $\|A - A_2\|_2^2$, where A_2 is the best rank-2 approximation of A. $\sigma_3^2 = 2$

(d) What is $\|A\|_2^2$? $= \sigma_1^2 = 10$

(e) What is $\|A\|_F^2$? $= \sigma_1^2 + \sigma_2^2 + \sigma_3^2 + \sigma_4^2 = 10 + 5 + 2 + 1 = 18$

Let v_1, v_2, v_3, v_4 be the right singular vectors of A.

(f) What is $\|Av_2\|^2$? $= \sigma_2^2 = 5$

$A = USV^T$

(g) What is $\langle v_1, v_3\rangle$? $= 0$

(h) What is $\|v_4\|$? $= 1$

Let $a_1 \in \mathbb{R}^4$ be the first row of A.

(i) Write a_1 in the basis defined by the right singular vectors of A.

$$\left(\langle v_1, a_1\rangle, \langle v_2, a_1\rangle, \langle v_3, a_1\rangle, \langle v_4, a_1\rangle\right)$$

$$= V^T a_1$$

13. Draw the Voronoi diagram of the following set of points.

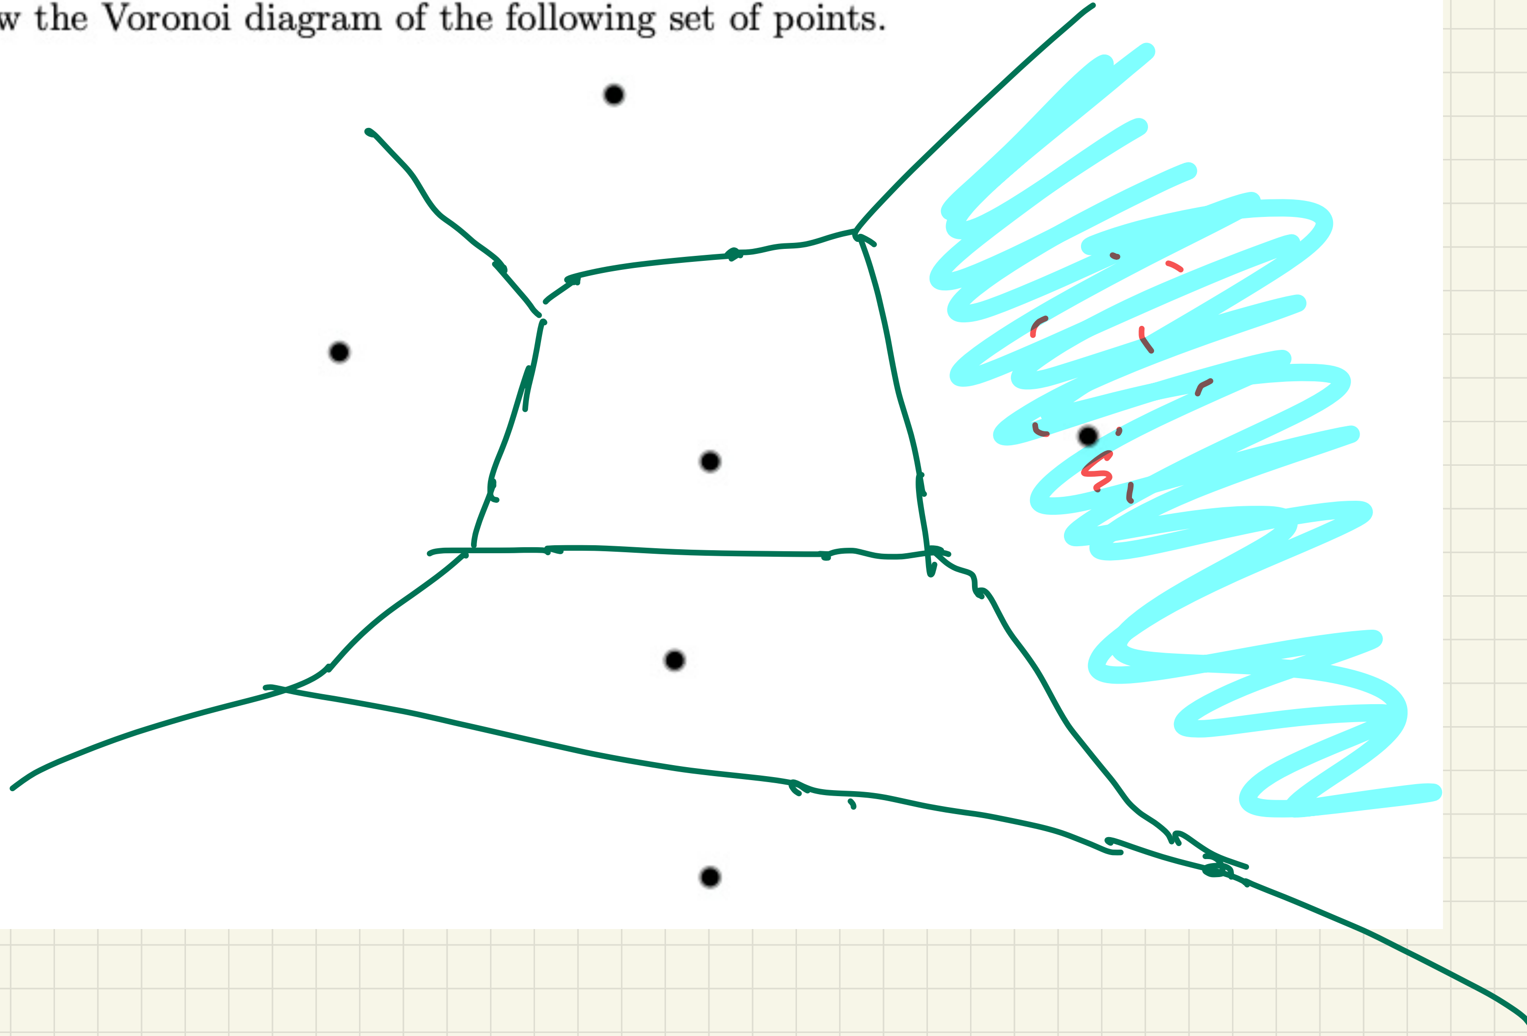

14. What should you do, if running Lloyd's algorithm for k-means clustering ($k = 2$), and you reach this scenario, where the algorithm terminates? (The black circles are data points and red stars are the centers).

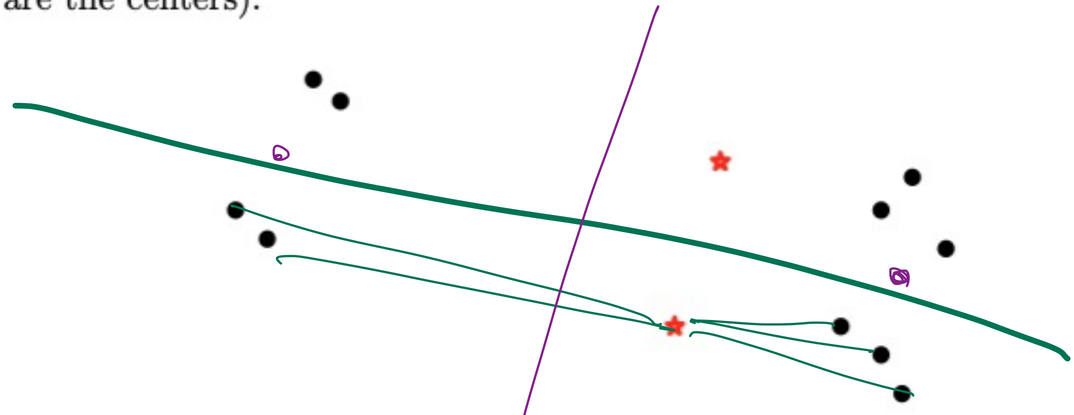

random restart

15. Consider the following "loss" function. $\ell_i(z_i) = (1 - z_i)^2/2$, where for a data point (x_i, y_i) and prediction function f, then $z_i = y_i \cdot f(x_i)$. Predict how this might work within a gradient descent algorithm for classification.

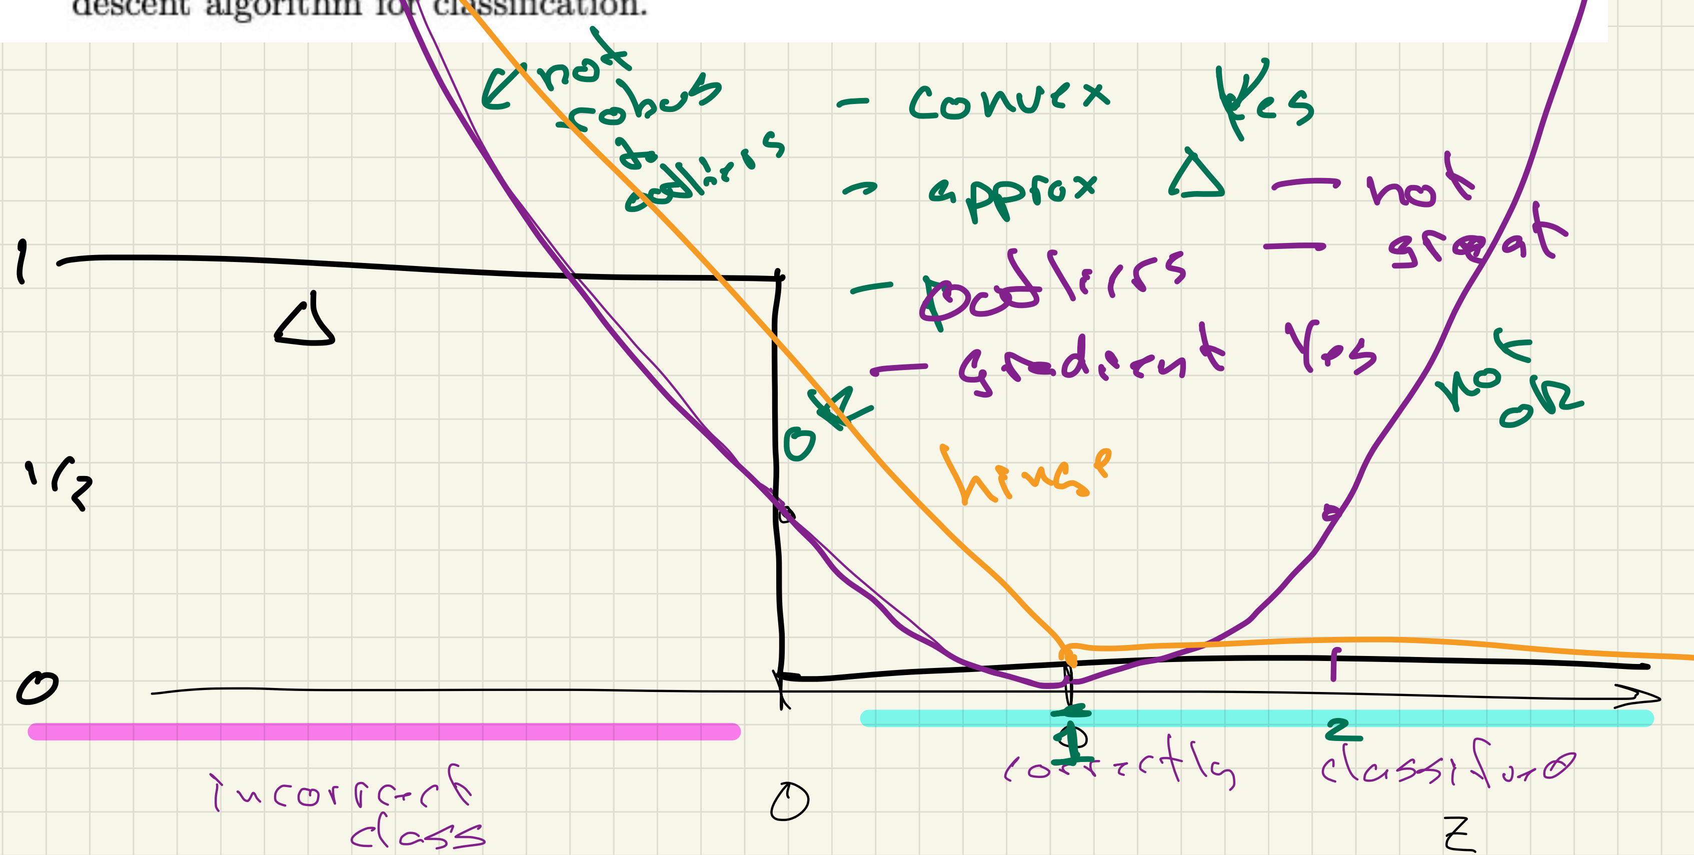

16. Consider a set of 1-dimensional data points

$$(x_1 = 0, y_1 = +1)\ (x_2 = 1, y_1 = -1)\ (x_3 = 2, y_1 = +1)\ (x_4 = 4, y_1 = +1)$$

$$(x_5 = 6, y_1 = -1)\ (x_6 = 7, y_1 = -1)\ (x_7 = 8, y_1 = +1)\ (x_8 = 9, y_1 = -1)$$

Predict **-1** or **+1** using a kNN (k-nearest neighbor) classifier with k = **3** on the following queries.

(a) $x = 3$

(b) $x = 9$

(c) $x = -1$