Solutions

10-601 Machine Learning
Fall 2022
Exam 2 Practice Problems
November 4, 2022
Time Limit: N/A

Name:
AndrewID:

Instructions:

- Fill in your name and Andrew ID above. Be sure to write neatly, or you may not receive credit for your exam.
- Clearly mark your answers in the allocated space **on the front of each page.** If needed, use the back of a page for scratch space, but you will not get credit for anything written on the back of a page. If you have made a mistake, cross out the invalid parts of your solution, and circle the ones which should be graded.
- No electronic devices may be used during the exam.
- Please write all answers in pen.
- You have N/A to complete the exam. Good luck!

Instructions for Specific Problem Types

For "Select One" questions, please fill in the appropriate bubble completely:

Select One: Who taught this course?

- ● Henry Chai
- ○ Marie Curie
- ○ Noam Chomsky

If you need to change your answer, you may cross out the previous answer and bubble in the new answer:

Select One: Who taught this course?

- ● Henry Chai
- ○ Marie Curie
- ~~●~~ Noam Chomsky

For "Select all that apply" questions, please fill in all appropriate squares completely:

Select all that apply: Which are scientists?

- ■ Stephen Hawking
- ■ Albert Einstein
- ■ Isaac Newton
- □ I don't know

Again, if you need to change your answer, you may cross out the previous answer(s) and bubble in the new answer(s):

Select all that apply: Which are scientists?

- ■ Stephen Hawking
- ■ Albert Einstein
- ■ Isaac Newton
- ~~■~~ I don't know

For questions where you must fill in a blank, please make sure your final answer is fully included in the given space. You may cross out answers or parts of answers, but the final answer must still be within the given space.

Fill in the blank: What is the course number?

10-601	10-~~7~~601

1 MLE/MAP

1. For the following questions, answer True or False and provide a brief justification of your answer.

   1. **True or False:** Consider the linear regression model $y = w^T x + \epsilon$. Assuming $\epsilon \sim \mathcal{N}(0, \sigma^2)$ and maximizing the conditional log-likelihood is equivalent to minimizing the sum of squared errors $\|y - w^T x\|_2^2$.

      True. The squared error term comes from the squared term in the Gaussian distribution.

   2. **True or False:** Consider n data points, each with one feature x_i and an output y_i. In linear regression, we assume $y_i \sim \mathcal{N}(wx_i, \sigma^2)$ and compute $\hat{w}$ through MLE.

      Suppose $y_i \sim \mathcal{N}(\log(wx_i), 1)$ instead. Then the maximum likelihood estimate $\hat{w}$ is the solution to the following equality:

      $$\sum_{i=1}^{n} x_i y_i = \sum_{i=1}^{n} x_i \log(wx_i)$$

      .

      False. The likelihood function can be written as

      $$\prod_{i=1}^{n} \frac{\exp(-(y_i - \log(wx_i))^2/2)}{\sqrt{2\pi}} = \frac{\exp(-\sum_{i=1}^{n}(y_i - \log(wx_i))^2/2)}{\sqrt{2\pi}}$$

      Differentiating wrt w and setting to zero gives us

      $$\sum_{i=1}^{n} 2(y_i - \log(wx_i))\frac{x_i}{wx_i} = 0 \implies \sum_{i=1}^{n} y_i = \sum_{i=1}^{n} \log(wx_i)$$

2. **Math:** Let $X_1, X_2, ..., X_N$ be data drawn independently from a uniform distribution over a diamond-shaped area with edge length $\sqrt{2}\theta$ in $\mathbb{R}^2$, where $\theta \in \mathbb{R}^+$ (see Figure 1). Thus, $X_i \in \mathbb{R}^2$ and the distribution is

   $$p(x|\theta) = \begin{cases} \frac{1}{2\theta^2} & \text{if } \|x\| \leq \theta \\ 0 & \text{otherwise} \end{cases}$$

where $\|x\| = |x_1| + |x_2|$ is the L_1 norm. Find the maximum likelihood estimate of θ.

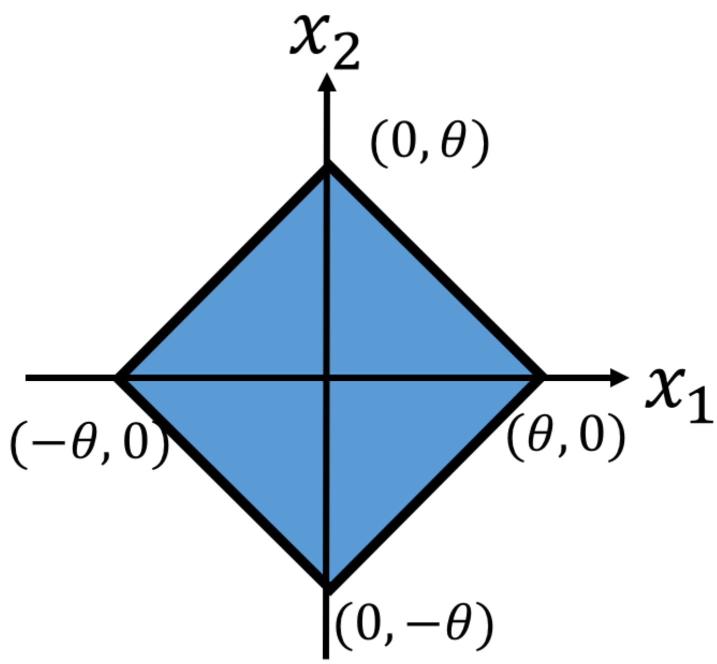

Figure 1: Area of $\|x\| \leq \theta$

Analysis:

The likelihood function is

$$L(X_1, X_2, ..., X_N; \theta) = \frac{1}{(2\theta^2)^N} 1 \left\{ \max_{1 \leq i \leq N} \|X_i\| \leq \theta \right\}$$

To maximize likelihood, we want θ to be as small as possible with the constraint of

$\underset{1\leq i\leq N}{max} \|X_i\| \leq \theta$, otherwise the likelihood drops to 0. So the MLE of θ is

$$\widehat{\theta} = \underset{1\leq i\leq N}{max} \|X_i\|$$

3. **Math:** Suppose we want to model a 1-dimensional dataset of N real valued features $\left(x^{(i)}\right)$ and targets $\left(y^{(i)}\right)$ by:

$$y^{(i)} \sim \mathcal{N}\left(\exp(wx^{(i)}), 1\right),$$

where w is our unknown (scalar) parameter and $\mathcal{N}$ is the normal distribution with probability density function:

$$f(a)_{\mathcal{N}(\mu,\sigma^2)} = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(a-\mu)^2}{2\sigma^2}\right)$$

Can the maximum conditional negative log likelihood estimator of w be solved analytically? If so, find the expression for w_{MLE}. If not, say so and write down the update rule for w using gradient descent.

Cannot be found analytically.

Taking the derivative of the negative log likelihood with respect to w yields:

$$\frac{\partial \text{NLL}}{\partial w} = \sum_{i}^{N} -x^{(i)}y^{(i)}\exp(wx^{(i)}) + x^{(i)}\exp(2wx^{(i)})$$

Update rule is thus

$$w \leftarrow w - \eta\frac{\partial \text{NLL}}{\partial w}$$

4. Assume we have n random variables $x_i, i \in [1, n]$, each drawn independently from a Normal distribution with mean μ and variance σ^2.

$$p(x_1, x_2, ..., x_n|\mu, \sigma^2) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(\frac{-(x_i-\mu)^2}{2\sigma^2}\right)$$

a) Write the log-likelihood function $\ell(x_1, x_2 ... x_n | \mu, \sigma^2)$.

$$\log(\prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}\sigma} \exp\{\frac{-(x_i - \mu)^2}{2\sigma^2}\}) = \sum_{i=1}^{n}[\log(\frac{1}{\sqrt{2\pi}\sigma}) - \frac{(x_i - \mu)^2}{2\sigma^2}] \tag{1}$$

$$= -n\log(\sqrt{2\pi}\sigma) - \frac{1}{2\sigma^2}\sum_{i=1}^{n}(x_i - \mu)^2 \tag{2}$$

b) Derive an expression for the Maximum Likelihood Estimate for the variance (σ^2).

We can find the estimator by solving $\nabla_\sigma l(x_1, x_2 ... x_n | \mu, \sigma^2) = 0$.

$$-n\frac{1}{\sqrt{2\pi}\sigma}\sqrt{2\pi} + \frac{1}{\sigma^3}\sum_{i=1}^{n}(x_i - \mu)^2 = 0 \tag{3}$$

$$\frac{1}{\sigma^3}\sum_{i=1}^{n}(x_i - \mu)^2 = \frac{n}{\sigma} \tag{4}$$

$$\frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^2 = \sigma^2 \tag{5}$$

5. Assume we have n random variables $x_i, i \in [1, n]$, each drawn independently from a Bernoulli distribution with mean θ. Recall that in a Bernoulli distribution $X \in \{0, 1\}$ and the pdf is:

$$p(X|\theta) = \theta^x (1-\theta)^{1-x}$$

a) Derive the likelihood $L(\theta|X_1, \ldots, X_n)$.

$L(\theta; X_1, \ldots, X_n) = \prod_{i=1}^n p(X_i; \theta)$
$L(\theta; X_1, \ldots, X_n) = \prod_{i=1}^n \theta^{x_i}(1-\theta)^{1-x_i}$
$L(\theta; X_1, \ldots, X_n) = \theta^{\sum_{i=1}^n x_i}(1-\theta)^{n-\sum_{i=1}^n x_i}$
Either of the final two steps are acceptable.

b) Show that the log-likelihood is:

$$l(\theta|X_1, \ldots, X_n) = \left(\sum_{i=1}^n X_i\right) \log(\theta) + \left(n - \sum_{i=1}^n X_i\right) \log(1-\theta)$$

$l(\theta; X_1, \ldots, X_n) = \log L(\theta; X_1, \ldots, X_n)$
$l(\theta; X_1, \ldots, X_n) = \log \left[\theta^{\sum_{i=1}^n x_i}(1-\theta)^{n-\sum_{i=1}^n x_i}\right]$
$l(\theta; X_1, \ldots, X_n) = (\sum_{i=1}^n x_i) \log(\theta) + (n - \sum_{i=1}^n x_i) \log(1-\theta)$

c) Show that the MLE is $\hat{\theta} = \frac{1}{n}\sum_{i=1}^{n} X_i$.

Take the derivative of the log likelihood and set it to zero

$$\frac{dl}{d\theta} = \frac{d}{d\theta}\left[(\textstyle\sum_{i=1}^n x_i)\log(\theta) + (n - \sum_{i=1}^n x_i)\log(1-\theta)\right] = 0$$

$$\frac{\sum_{i=1}^n x_i}{\theta} - \frac{n - \sum_{i=1}^n x_i}{1-\theta} = 0$$

$$\left(\textstyle\sum_{i=1}^n x_i\right)(1-\theta) - \left(n - \sum_{i=1}^n x_i\right)\theta = 0$$

$$\textstyle\sum_{i=1}^n x_i - n\theta = 0$$

$$\hat{\theta} = \frac{1}{n}(\textstyle\sum_{i=1}^n X_i)$$

(1) Renato (2) Derive MLE starting with a probability representation. Derive the likelihood function, convert to log likelihood, and then derive the MLE estimate. (3) 10601 Spring 2016 Practice Midterm

6. Magnetic Resonance Imaging (MRI) scans are commonly used to generate detailed images of patients' internal anatomy at hospitals. The scanner returns an image with N pixels. For each pixel we extract the noise from that pixel to obtain a vector of noise terms $\mathbf{x} \in \mathbb{R}^N$ s.t. $\forall\ i \in \{1...N\}$, $x_i \geq 0$ and x_i is independent and identically distributed and follows a Rayleigh distribution. The probability density function of a Rayleigh distribution is given by:

$$f(x \mid \sigma) = \frac{x}{\sigma^2}\exp\left(\frac{-x^2}{2\sigma^2}\right)$$

for scale parameter $\sigma \geq 0$ and $x \geq 0$.

i. (2 points) Write the log-likelihood $\ell(\sigma)$ of a noise vector $\mathbf{x}$ obtained from one image. Report your answer in terms of the variables x_i, i, N, σ, the function $\exp(\cdot)$, and any constants you may need. For full credit you must push the log through to remove as many multiplications/divisions as possible.

$$\ell(\sigma) = \sum_{i=1}^{N} \left[\log x_i - 2 \log \sigma - \frac{x_i^2}{2\sigma^2} \right]$$

ii. (2 points) Report the maximum likelihood estimator of the scale parameter, σ, for a single image's noise vector $\mathbf{x}$.

$$0 = \frac{\partial}{\partial \sigma} \sum_{i=1}^{N} \log p(x_i \mid \sigma) = \sum_{i=1}^{N} \frac{-2}{\sigma} + \frac{x_i^2}{\sigma^3}$$

$$\implies \hat{\sigma} = \left[\frac{1}{2N} \sum_{i=1}^{N} x_i^2 \right]^{\frac{1}{2}}$$

2 Probability and Naive Bayes

2.1 Probability

1. For each question, choose the correct option.

   1. **Select one:** Which of the following expressions is equivalent to $p(A|B, C, D)$?

      ◯ $\frac{p(A,B,C,D)}{p(C|B,D)p(B|D)p(D)}$

      ◯ $\frac{p(A,B,C,D)}{p(B,C)p(D)}$

      ◯ $\frac{p(A,B,C,D)}{p(B,C|D)p(B)p(C)}$

      Answer is (a). $p(A|BCD) = \frac{p(A,B,C,D)}{p(B,C,D)} = \frac{p(A,B,C,D)}{p(C|B,D)p(B,D)} = \frac{p(A,B,C,D)}{p(C|B,D)p(B|D)p(D)}$

   2. **True or False:** Let μ be the mean of some probability distribution. $p(\mu)$ is always non-zero.

      ◯ True

      ◯ False

      No, taking the example of a distribution that put 0.5 probability on +1 and 0.5 probability on -1. Though their mean is 0, p(0) is zero.

2. **True or False:** Assume we have a sample space Ω. For each question, choose True or False; no justification needed.

   1. If events A, B, and C are disjoint then they are independent.

      ◯ True

      ◯ False

      False. If they are disjoint, i.e. mutually exclusive, they are very dependent! (what does disjoint mean in terms of the probabilities of A, B, and C? What about independent?)

   2. $P(A|B) \propto \frac{P(A)P(B|A)}{P(A|B)}$.

      ◯ True

      ◯ False

      False $P(A|B) \propto \frac{P(A)P(B|A)}{P(B)}$

   3. $P(A \cup B) \leq P(A)$.

      ◯ True

      ◯ False

False $P(A \cup B) \geq P(A)$

4. $P(A \cap B) \geq P(A)$.

   ◯ True

   ◯ False

False $P(A \cap B) \leq P(A)$

2.2 Naive Bayes

1. Consider the following data. It has 4 features $\mathbf{X} = (x_1, x_2, x_3, x_4)$ and 3 labels $y \in \{+1, 0, -1\}$. Assume that the probabilities $p(\mathbf{X}|y)$ and $p(y)$ are both Bernoulli distributions. Answer the questions that follow under the Naive Bayes assumption.

x_1	x_2	x_3	x_4	y
1	1	0	1	+1
0	1	1	0	+1
1	0	1	1	0
0	1	1	1	0
0	1	0	0	-1
1	0	0	1	-1
0	0	1	1	-1

1. Compute the Maximum Likelihood Estimates for $p(x_i = 1|y), \forall i \in \{1, 2, 3, 4\}$ and $\forall y \in \{+1, 0, -1\}$.

	$y = +1$	$y = 0$	$y = -1$
$x_1 = 1$			
$x_2 = 1$			
$x_3 = 1$			
$x_4 = 1$			

	$y = +1$	$y = 0$	$y = -1$
$x_1 = 1$	0.5	0.5	1/3
$x_2 = 1$	1	0.5	1/3
$x_3 = 1$	0.5	1	1/3
$x_4 = 1$	0.5	1	2/3

2. Compute the Maximum Likelihood Estimates for the prior probabilities $p(y = +1), p(y = 0), p(y = -1)$.

$p(y = +1) = \frac{2}{7}$, $p(y = 0) = \frac{2}{7}$ and $p(y = -1) = \frac{3}{7}$.

3. Use the values computed in the above two parts to classify the data point $(x_1 = 1, x_2 = 1, x_3 = 1, x_4 = 1)$ as belonging to class $+1, 0$ or -1.

   According to Naïve Bayes assumption, features are independent given y, thus we can write the conditional joint probability as

$$p(x_1 = 1, x_2 = 1, x_3 = 1, x_4 = 1) = p(x_1 = 1, x_2 = 1, x_3 = 1, x_4 = 1|y)p(y) \tag{6}$$

$$= p(y) \prod_{i=1}^{4} p(x_i = 1|y). \tag{7}$$

   We calculate the probability given different values of y and pick the one with highest probability:

$$p(y = +1) \prod_{i=1}^{4} p(x_i = 1|y = +1) = \frac{1}{2} \cdot 1 \cdot \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{2}{7} = \frac{1}{28} \tag{8}$$

$$p(y = 0) \prod_{i=1}^{4} p(x_i = 1|y = 0) = \frac{1}{2} \cdot \frac{1}{2} \cdot 1 \cdot 1 \cdot \frac{2}{7} = \frac{1}{14} \tag{9}$$

$$p(y = -1) \prod_{i=1}^{4} p(x_i = 1|y = -1) = \frac{1}{3} \cdot \frac{1}{3} \cdot \frac{1}{3} \cdot \frac{2}{3} \cdot \frac{3}{7} = \frac{2}{189} \tag{10}$$

   Since $y = 0$ yields the largest value, we classify the data as $\hat{y} = 0$.

2. You are given a dataset of 10,000 students with their sex, height, and hair color. You are trying to build a machine learning classifier to predict the sex of a student, so you randomly split the data into a training set and a testing set. Here are the specifications of the data set:

   - sex ∈ {male,female}
   - height ∈ [0,300] centimeters
   - hair ∈ {brown, black, blond, red, green}
   - 3240 men in the data set
   - 6760 women in the data set

   Under only the assumptions necessary for Naïve Bayes (not the distributional assumptions you might naturally or intuitively make about the data set), answer True or False and provide a one sentence justification of your answer.

   1. **True or False:** Height is a continuous valued variable. Therefore, Naïve Bayes is not appropriate since it cannot handle continuous valued variables.

False. Naïve Bayes can handle both continuous and discrete values as long as the appropriate distributions are used for conditional probabilities. For example, Gaussian for continuous and Bernoulli for discrete

2. **True or False:** Since there aren't similar numbers of men and women in the data set, Naïve Bayes will have high test error.

False. Since the data was randomly split, the same proportion of male and female will be in the training and testing sets. Thus this discrepancy will not affect testing error.

3. **True or False:** $p(\texttt{height}|\texttt{sex}, \texttt{hair}) = p(\texttt{height}|\texttt{sex})$.

True. This results from the conditional independence assumption required for Naïve Bayes.

4. **True or False:** $p(\texttt{height}, \texttt{hair}|\texttt{sex}) = p(\texttt{height}|\texttt{sex}) * p(\texttt{hair}|\texttt{sex})$.

True. This results from the conditional independence assumption required for Naïve Bayes.

2.3 Naive Bayes, Logistic Regression

1. Suppose you wish to learn $P(Y|X_1, X_2, X_3)$, where Y, X_1, X_2 and X_3 are all boolean-valued random variables. You consider both Naïve Bayes and Logistic Regression as possible approaches.

   For questions 1-5, answer True or False and provide a one sentence justification for your answer.

   1. **True or False:** In this case, a good choice for Naïve Bayes would be to implement a Gaussian Naïve Bayes classifier.

   Answer: False. Given the X_i are boolean, it is better to model $P(X_i|Y)$ with a Bernoulli rather than Gaussian distribution

2. **True or False:** To learn $P(Y|X_1, X_2, X_3)$ using Naïve Bayes, you must make conditional independence assumptions, including the assumption that Y is conditionally independent of X_1 given X_2.

   Answer: False. Naïve Bayes assume $(\forall i \neq j) X_i$ is conditionally independent of X_j given Y.

3. **True or False:** Logistic regression is certain to be the better choice in this case.

   Answer: False. It will depend on the number of training examples available, and whether the Naïve Bayes assumptions are satisfied.

4. **True or False:** We can train Naïve Bayes using maximum likelihood estimates for each parameter, but not MAP estimates.

   False: MAP estimates are just MLE spiced up with priors on the parameters of $P(X_i|Y_j)$ (prior knowledge that we can inject into the model), so there's no reason we can't add it in.

5. **True or False:** We can train Logistic Regression using maximum likelihood estimates for each parameter, but not MAP estimates.

   False: We can assign priors on the regression model by assuming $y = w^T x + \epsilon$ where ϵ is "noise" from a distribution (e.g. Gaussian)

6. How many parameters must be estimated for your Bernoulli Naïve Bayes classifier? List the parameters.

   7 - We need $P(Y = 1)$ and $P(Y = 0) = 1 - P(Y = 1)$. Then we need $P(X_i = 0|Y = 0), P(X_i = 1|Y = 0), P(X_i = 0|Y = 1), P(X_i = 1|Y = 1)$ for $i \in \{1, 2, 3\}$.

But given the binary parameters, $P(X_i = 1|Y = 0) + P(X_i = 0|Y = 0) = 1$, so we only need to estimate one of these conditional probabilities ($P(X_i = 0|Y = 0)$, for example). So in total we need only 2*3 + 1 = 7

7. How many parameters must be estimated for your Logistic Regression classifier? List the parameters.

4 - Weights for: Bias, X_1, X_2, and X_3

2. Suppose we add a numeric, real-valued variable X_4 to our problem. Note we now have a mix of some discrete-valued X_i and one continuous X_i.

1. Explain why we can no longer use Naïve Bayes, or if we can, how we would modify our original solution.

We can't use our original model because a Bernoulli distribution can't model the new data.

We must modify our solution so that a different Naïve Bayes model is trained on the continuous variables, using a different distribution than a Bernoulli one (i.e. a Gaussian), and then the result is a multiplication of the two.

Since our assumption is that each parameter is conditionally independent anyhow, we can multiply the results of the two models together safely.

2. Explain why we can no longer use Logistic Regression, or if we can, how we would modify our original solution.

Assuming the discrete variables were already transformed into one-hot features, we can simply add one weight per continuous feature to our model.

3 Logistic Regression and Regularization

1. A generalization of logistic regression to a multiclass settings involves expressing the per-class probabilities $P(y = c|x)$ as the softmax function $\frac{\exp(w_c^T x)}{\sum_{d \in C} \exp(w_d^T x)}$, where c is some class from the set of all classes C.

   Consider a 2-class problem (labels 0 or 1). Rewrite the above expression for this situation to end up with expressions for $P(Y = 1|x)$ and $P(Y = 0|x)$ that we have already come across in class for binary logistic regression.

   $P(y = 1|x) = \frac{\exp(w_1^T x)}{\exp(w_0^T x)+\exp(w_1^T x)} = \frac{\exp((w_1 - w_0)^T x)}{1+\exp((w_1 - w_0)^T x)} = \frac{\exp(w^T x)}{1+\exp(w^T x)} = p$

   Therefore, $1 - p = \frac{1}{1+\exp(w^T x)}$

2. Considering a Gaussian prior, write out the MAP objective function $J(w)_{MAP}$ in terms of the MLE objective $J(w)_{MLE}$. Name the variant of logistic regression this results in.

   $J_{MAP}(\mathbf{w}) = J_{MLE}(\mathbf{w}) - \lambda \|\mathbf{w}\|_2^2$. This is L2 regularized logistic regression.

3. Given a training set $\mathcal{D} = \{(\mathbf{x}^{(1)}, y^{(1)}, \ldots, (\mathbf{x}^{(N)}, y^{(N)})\}$ where $\mathbf{x}^{(i)} \in \mathbb{R}^d$ is a feature vector and $y_i \in \{0, 1\}$ is a binary label, we want to find the parameters $\hat{w}$ that maximize the likelihood for the training set, assuming a parametric model of the form

$$p(y = 1|x; w) = \frac{1}{1 + \exp(-w^T x)}.$$

The conditional log likelihood of the training set is

$$\ell(w) = \sum_{i=1}^{n} y_i \log p(y_i, |x_i; w) + (1 - y_i) \log(1 - p(y_i, |x_i; w)),$$

and the gradient is

$$\nabla \ell(w) = \sum_{i=1}^{n} (y_i - p(y_i|x_i; w))x_i.$$

a) Is it possible to get a closed form for the parameters $\hat{w}$ that maximize the conditional log likelihood? How would you compute $\hat{w}$ in practice?

There is no closed form expression for maximizing the conditional log likelihood. One has to consider iterative optimization methods, such as gradient descent, to compute $\hat{w}$.

b) For a binary logistic regression model, we predict $y = 1$ when $p(y = 1|x) \geq 0.5$. Show that this is a linear classifier.

Using the parametric form for $p(y = 1|x)$:

$$\begin{aligned} p(y=1|x) \geq \frac{1}{2} &\implies \frac{1}{1+\exp(-w^T x)} \geq \frac{1}{2} \\ &\implies 1+\exp(-w^T x) \leq 2 \\ &\implies \exp(-w^T x) \leq 1 \\ &\implies -w^T x \leq 0 \\ &\implies w^T x \geq 0, \end{aligned}$$

so we predict $\hat{y} = 1$ if $w^T x \geq 0$.

c) Consider the case with binary features, i.e, $x \in \{0,1\}^d$, where feature x_1 is rare and happens to appear in the training set with only label 1. What is $\hat{w}_1$? Is the gradient ever zero for any finite w? Why is it important to include a regularization term to control the norm of $\hat{w}$?

If a binary feature fired for only label 1 in the training set then, by maximizing the conditional log likelihood, we will make the weight associated to that feature be infinite. This is because, when this feature is observed in the training set, we will want to predict predict 1 irrespective of everything else. This is an undesired behaviour from the point of view of generalization performance, as most likely we do not believe this rare feature to have that much information about class 1. Most likely, it is spurious co-occurrence. Controlling the norm of the weight vector will prevent these pathological cases.

4. Given the following dataset, $\mathcal{D}$, and a fixed parameter vector, $\boldsymbol{\theta}$, write an expression for the binary logistic regression conditional likelihood.

$\mathcal{D} = \{(\mathbf{x}^{(1)}, y^{(1)} = 0), (\mathbf{x}^{(2)}, y^{(2)} = 0), (\mathbf{x}^{(3)}, y^{(3)} = 1), (\mathbf{x}^{(4)}, y^{(4)} = 1)\}$

- Write your answer in terms of $\boldsymbol{\theta}$, $\mathbf{x}^{(1)}$, $\mathbf{x}^{(2)}$, $\mathbf{x}^{(3)}$, and $\mathbf{x}^{(4)}$.

- Do not include $y^{(1)}$, $y^{(2)}$, $y^{(3)}$, or $y^{(4)}$ in your answer.
- Don't try to simplify your expression.

Conditional likelihood:

$$\left(1 - \frac{1}{1+e^{-\theta^T x^1}}\right)\left(1 - \frac{1}{1+e^{-\theta^T x^2}}\right)\frac{1}{1+e^{-\theta^T x^3}}\frac{1}{1+e^{-\theta^T x^4}}$$

5. Write an expression for the decision boundary of binary logistic regression with a bias term for two-dimensional input features $x_1 \in \mathbb{R}$ and $x_2 \in \mathbb{R}$ and parameters b (the intercept parameter), w_1, and w_2. Assume that the decision boundary occurs when $P(Y = 1 \mid \mathbf{x}, b, w_1, w_2) = P(Y = 0 \mid \mathbf{x}, b, w_1, w_2)$.

   (a) Write your answer in terms of x_1, x_2, b, w_1, and w_2.

   Decision boundary equation:

   $0 = b + w_1 x_1 + w_2 x_2$

   (b) What is the geometric shape defined by this equation?

   A line.

6. We have now feature engineered the two-dimensional input, $x_1 \in \mathbb{R}$ and $x_2 \in \mathbb{R}$, mapping it to a new input vector: $\mathbf{x} = \begin{bmatrix} 1 \\ {x_1}^2 \\ {x_2}^2 \end{bmatrix}$

   (a) Write an expression for the decision boundary of binary logistic regression with this feature vector $\mathbf{x}$ and the corresponding parameter vector $\boldsymbol{\theta} = [b, w_1, w_2]^T$. Assume that the decision boundary occurs when $P(Y = 1 \mid x, \boldsymbol{\theta}) = P(Y = 0 \mid x, \boldsymbol{\theta})$. Write your answer in terms of x_1, x_2, b, w_1, and w_2.

   Decision boundary expression:

   $0 = b + w_1 x_1^2 + w_2 x_2^2$.

   (b) What is the geometric shape defined by this equation?

An ellipse

(c) If we add an L2 regularization term when learning $[w_1, w_2]^T$, what happens to the **parameters** as we increase the λ that scales this regularization term?

The magnitude of the parameters will decrease.

(d) If we add an L2 regularization term when learning $[w_1, w_2]^T$, what happens to the **decision boundary shape** as we increase the λ that scales this regularization term?

The parameters shrink, so the ellipse will get bigger.

7. **Model Complexity:** In this question we will consider the effect of increasing model complexity, while keeping the size of the training set fixed. Consider a classification task on the real line $\mathbb{R}$ with distribution D and target function $c^* : \mathbb{R} \to \{\pm 1\}$, and suppose we have a random sample S of size n drawn iid from D. For each degree d, let ϕ_d be the feature map given by $\phi_d(x) = (1, x, x^2, \ldots, x^d)$ that maps points on the real line to $(d+1)$-dimensional space.

   Now consider the learning algorithm that first applies the feature map ϕ_d to all the training examples and then runs logistic regression. A new example is classified by first applying the feature map ϕ_d and then using the learned classifier.

   a) For a given dataset S, is it possible for the training error to increase when we increase the degree d of the feature map? **Please explain your answer in 1 to 2 sentences.**

No. Every linear separator using the feature map ϕ_d can also be expressed using the feature map ϕ_{d+1}, since we are only adding new features. It follows that the training error will not increase cvv for any given sample S.

b) Briefly **explain in 1 to 2 sentences** why the true error first drops and then increases as we increase the degree d.

When the dimension d is small, the true error is high because it is not possible to the target function is not well approximated by any linear separator in the ϕ_d feature space. As we increase d, our ability to approximate c^* improves, so the true error drops. But, as we continue to increase d, we begin to overfit the data and the true error increases again.

8. **Short Answer:** Your friend is training a logistic regression model with ridge regularization, where λ is the regularization constant. They run cross-validation for $\lambda = [0.01, 0.1, 1, 10]$ and compare train, validation and test errors. They choose $\lambda = 0.01$ because that had the lowest *test* error.

   However, you observe that the test error linearly increases from $\lambda = 0.01$ to 10 and thus, there exists a value of $\lambda < 0.01$ that gives a lower test error. You tell your friend that they should run the cross-validation for $\lambda = [0.0001, 0.001, 0.01]$ to get the optimal model.

   Do you think you did the right thing by giving your friend this suggestion? Briefly justify your answer in 1-2 concise sentences.

   No. because we should not be using test error at all in making any model selection decisions.

4 Feature Engineering and Regularization

1. **Model Complexity:** In this question we will consider the effect of increasing the model complexity, while keeping the size of the training set fixed. To be concrete, consider a classification task on the real line $\mathbb{R}$ with distribution D and target function $c^* : \mathbb{R} \to \{\pm 1\}$, and suppose we have a random sample S of size n drawn iid from D. For each degree d, let ϕ_d be the feature map given by $\phi_d(x) = (1, x, x^2, \ldots, x^d)$ that maps points on the real line to $(d+1)$-dimensional space.

   Now consider the learning algorithm that first applies the feature map ϕ_d to all the training examples and then runs logistic regression. A new example is classified by first applying the feature map ϕ_d and then using the learned classifier.

   a) For a given dataset S, is it possible for the training error to increase when we increase the degree d of the feature map? **Please explain your answer in 1 to 2 sentences.**

      No. Every linear separator using the feature map ϕ_d can also be expressed using the feature map ϕ_{d+1}, since we are only adding new features. It follows that the training error will not increase cvv for any given sample S.

   b) Briefly **explain in 1 to 2 sentences** why the true error first drops and then increases as we increase the degree d. When the dimension d is small, the true error is high because it is not possible to the target function is not well approximated by any linear separator in the ϕ_d feature space. As we increase d, our ability to approximate c^* improves, so the true error drops. But, as we continue to increase d, we begin to overfit the data and the true error increases again.

5 Neural Networks

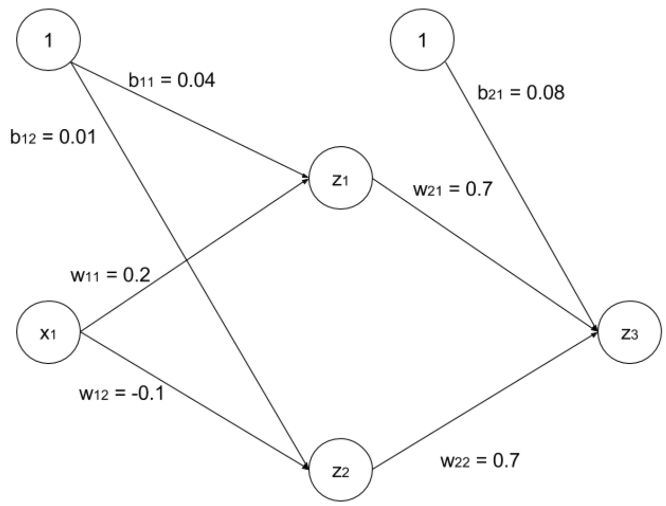

Figure 2: neural network

1. Consider the neural network architecture shown above for a binary classification problem. The values for weights and biases are shown in the figure. We define:

   $a_1 = w_{11}x_1 + b_{11}$

   $a_2 = w_{12}x_1 + b_{12}$

   $a_3 = w_{21}z_1 + w_{22}z_2 + b_{21}$

   $z_1 = \text{ReLU}(a_1)$

   $z_2 = \text{ReLU}(a_2)$

   $z_3 = \sigma(a_3)$, $\sigma(x) = \frac{1}{1+e^{-x}}$

   (i) For $x_1 = 0.3$, compute z_3 in terms of e.

   $z_3 = \frac{1}{1+e^{-0.15}}$

   (ii) Which class does the network predict for the data point ($x_1 = 0.3$)? Note that $\hat{y} = 1$ if $z_3 > \frac{1}{2}$, else $\hat{y} = 0$.

$\hat{y}(x_1 = 0.3) = 1$

(iii) Perform backpropagation on the bias term b_{21} by deriving the expression for the gradient of the loss function $L(y, z_3)$ with respect to the bias term b_{21}, $\frac{\partial L}{\partial b_{21}}$, in terms of the partial derivatives $\frac{\partial \alpha}{\partial \beta}$, where α and β can be any of $L, z_i, a_i, b_{ij}, w_{ij}, x_1$ for all valid values of i, j. Your backpropagation algorithm should be as explicit as possible — that is, make sure each partial derivative $\frac{\partial \alpha}{\partial \beta}$ cannot be decomposed further into simpler partial derivatives. Do *not* evaluate the partial derivatives.

$$\frac{\partial L}{\partial b_{21}} = \frac{\partial L}{\partial z_3} \frac{\partial z_3}{\partial a_3} \frac{\partial a_3}{\partial b_{21}}$$

(iv) Perform backpropagation on the bias term b_{12} by deriving the expression for the gradient of the loss function $L(y, z_3)$ with respect to the bias term b_{12}, $\frac{\partial L}{\partial b_{12}}$, in terms of the partial derivatives $\frac{\partial \alpha}{\partial \beta}$, where α and β can be any of $L, z_i, a_i, b_{ij}, w_{ij}, x_1$ for all valid values of i, j. Your backpropagation algorithm should be as explicit as possible — that is, make sure each partial derivative $\frac{\partial \alpha}{\partial \beta}$ cannot be decomposed further into simpler partial derivatives. Do *not* evaluate the partial derivatives.

$$\frac{\partial L}{\partial b_{12}} = \frac{\partial L}{\partial z_3} \frac{\partial z_3}{\partial a_3} \frac{\partial a_3}{\partial z_2} \frac{\partial z_2}{\partial a_2} \frac{\partial a_2}{\partial b_{12}}$$

2. In this problem we will use a neural network to distinguish the crosses ($\times$) from the circles ($\circ$) in the simple data set shown in Figure 3a. Even though the crosses and circles are not linearly separable, we can break the examples into three groups, S_1, S_2, and S_3 (shown in Figure 3a) so that S_1 is linearly separable from S_2 and S_2 is linearly separable from S_3. We will exploit this fact to design weights for the neural network shown in Figure 3b in order to correctly classify this training set. For all nodes, we will use the threshold activation function

$$\phi(z) = \begin{cases} 1 & z > 0 \\ 0 & z \le 0. \end{cases}$$

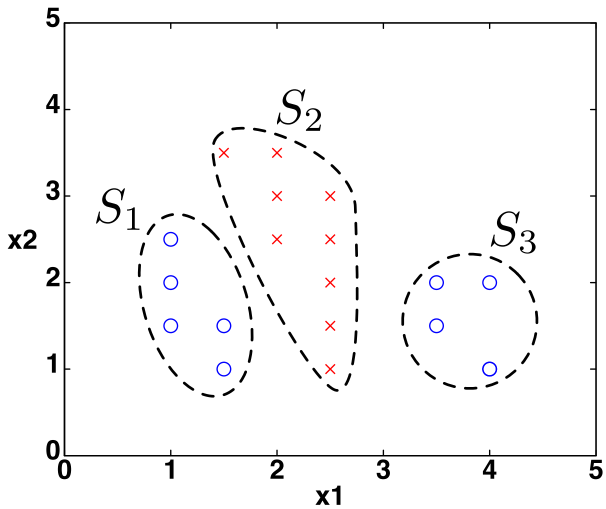

(a) The data set with groups S_1, S_2, and S_3.

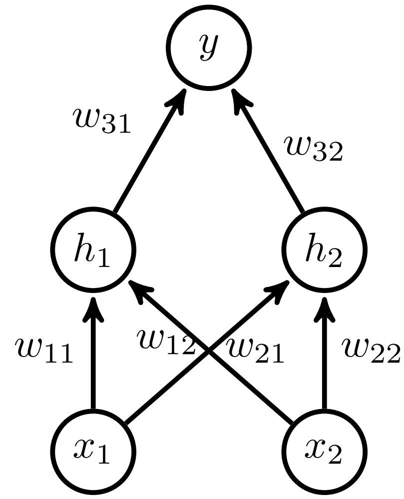

(b) The neural network architecture

Figure 3

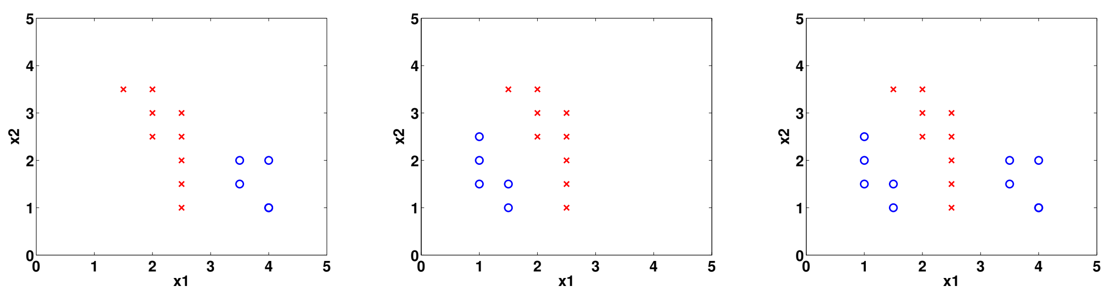

(a) Set S2 and S3 (b) Set S1 and S2 (c) Set S1, S2 and S3

Figure 4: NN classification.

(i) First we will set the parameters w_{11}, w_{12} and b_1 of the neuron labeled h_1 so that its output $h_1(x) = \phi(w_{11}x_1 + w_{12}x_2 + b_1)$ forms a linear separator between the sets S_2 and S_3.

(a) On Fig 4a, draw a linear decision boundary that separates S_2 and S_3.

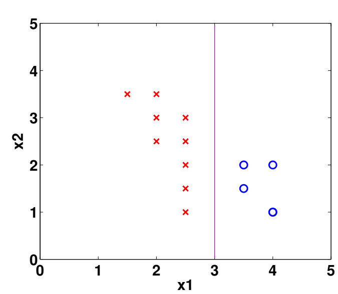

(b) Write down the corresponding weights w_{11}, w_{12}, and b_1 so that $h_1(x) = 0$ for all points in S_3 and $h_1(x) = 1$ for all points in S_2. One solution suffices and the same applies to (ii) and (iii).

$w_{11} = -1, w_{12} = 0, b_1 = 3$

(ii) Next we set the parameters w_{21}, w_{22} and b_2 of the neuron labeled h_2 so that its output $h_2(x) = \phi(w_{21}x_1 + w_{22}x_2 + b_2)$ forms a linear separator between the sets S_1 and S_2.

(a) On Fig 4b, draw a linear decision boundary that separates S_1 and S_2.

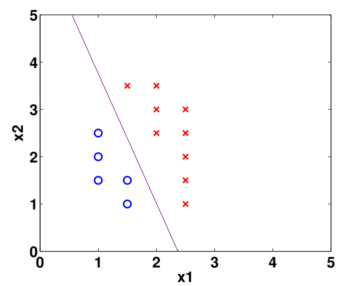

(b) Write down the corresponding weights w_{21}, w_{22}, and b_2 so that $h_2(x) = 0$ for all points in S_1 and $h_2(x) = 1$ for all points in S_2.

$w_{21} = 3, w_{22} = 1, b_2 = -7$

(iii) Now we have two classifiers h_1 (to classify S_2 from S_3) and h_2 (to classify S_1 from S_2). We will set the weights of the final neuron of the neural network based on the results from h_1 and h_2 to classify the crosses from the circles. Let $h_3(x) = \phi\big(w_{31}h_1(x) + w_{32}h_2(x) + b_3\big)$.

(a) Compute w_{31}, w_{32}, b_3 such that $h_3(x)$ correctly classifies the entire data set.

$w_{31} = 1, w_{32} = 1, b_3 = -1.5$

(b) Draw your decision boundary in Fig 4c.

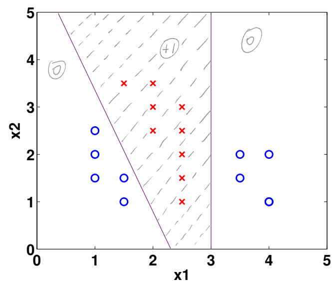

3. One part of learning parameters in a neural network is getting the gradients of the parameters.

Suppose we have a dataset $\mathcal{D}$ with N data points x_i with label y_i, where $i \in [1, N]$. x_i is a $d \times 1$ vector and $y_i \in \{0, 1\}$. We use the data to train a neural network with one hidden layer:

$$h(x) = \sigma(W_1 x + b_1)$$
$$p(x) = \sigma(W_2 h(x) + b_2),$$

where $\sigma(x) = \frac{1}{1+\exp(-x)}$ is the sigmoid function, W_1 is a n by d matrix, b_1 is a n by 1 vector, W_2 is a 1 by n matrix, and b_1 is a 1 by 1 vector.

We use cross entropy loss and minimize the negative log likelihood to train the neural network:

$$\ell_{\mathcal{D}}(W) = \frac{1}{N}\sum_{i=1}^{N} \ell_i(W) = \frac{1}{m}\sum_{i=1}^{N} -(y_i \log p_i + (1 - y_i)\log(1 - p_i)),$$

where $p_i = p(x_i), h_i = h(x_i)$.

(a) Describe how you would derive the gradients w.r.t the parameters W_1, W_2 and b_1, b_2. You do not need to write out the actual mathematical expression.

Use the chain rule.

(b) When N is large, we typically use a small subset of the dataset to estimate the gradient — stochastic gradient descent (SGD). Explain why we use SGD instead of gradient descent.

SGD converges faster than gradient descent.

(c) Derive expressions for the following gradients: $\frac{\partial l}{\partial p_i}, \frac{\partial l}{\partial W_2}, \frac{\partial l}{\partial b_2}, \frac{\partial l}{\partial h_i}, \frac{\partial l}{\partial W_1}, \frac{\partial l}{\partial b_1}$. When deriving the gradient w.r.t. the parameters in lower layers, you may assume the gradient in upper layers are available to you (i.e., you can use them in your equation). For example, when calculating $\frac{\partial l}{\partial W_1}$, you can assume $\frac{\partial l}{\partial p_i}, \frac{\partial l}{\partial W_2}, \frac{\partial l}{\partial b_2}, \frac{\partial l}{\partial h_i}$ are known.

$$
\begin{aligned}
\frac{\partial l}{\partial p_i} &= \frac{1}{m}\left(-\frac{y_i}{p_i} + \frac{1-y_i}{1-p_i}\right) \\
\frac{\partial l}{\partial W_2} &= \frac{1}{m}\sum_i \frac{\partial l_i}{\partial p_i}\frac{\partial p_i}{\partial W_2} = \frac{1}{m}\sum_i \frac{\partial l_i}{\partial p_i} p_i(1-p_i)h_i^T \\
\frac{\partial l}{\partial b_2} &= \frac{1}{m}\sum_i \frac{\partial l_i}{\partial p_i} p_i(1-p_i) \\
\frac{\partial l}{\partial h_i} &= \frac{\partial l}{\partial p_i}\frac{\partial p_i}{\partial h_i} = \frac{\partial l}{\partial p_i} p_i(1-p_i)W_2^T \\
\frac{\partial l}{\partial W_1} &= \frac{1}{m}\sum_i \frac{\partial l_i}{\partial h_i}\frac{\partial h_i}{\partial W_1} = \frac{1}{m}\sum_i \left[\frac{\partial l_i}{\partial h_i} \circ h_i \circ (1-h_i)\right] x_i^T \\
\frac{\partial l}{\partial b_1} &= \frac{1}{m}\sum_i \frac{\partial l_i}{\partial h_i}\frac{\partial h_i}{\partial b_1} = \frac{1}{m}\sum_i \frac{\partial l_i}{\partial h_i} \circ h_i \circ (1-h_i)
\end{aligned}
$$

4. Consider the following neural network for a 2-D input, $x_1 \in \mathbb{R}$ and $x_2 \in \mathbb{R}$ where:

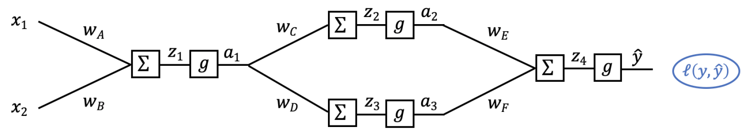

Figure 8: Neural Network

- All g functions are the same arbitrary non-linear activation function with no parameters
- $\ell(y, \hat{y})$ is an arbitrary loss function with no parameters, and:

$$z_1 = w_A x_1 + w_B x_2 \;\; a_1 = g(z_1)$$

$$z_2 = w_C a_1 \;\; a_2 = g(z_2)$$

$$z_3 = w_D a_1 \;\; a_3 = g(z_3)$$

$$z_4 = w_E a_2 + w_F a_3 \;\; \hat{y} = g(z_4)$$

Note: There are no bias terms in this network.

(a) What is the chain of partial derivatives needed to calculate the derivative $\frac{\partial \ell}{\partial w_E}$?

Your answer should be in the form: $\frac{\partial \ell}{\partial w_E} = \frac{\partial ?}{\partial ?}\frac{\partial ?}{\partial ?}\ldots$ Make sure each partial derivative $\frac{\partial ?}{\partial ?}$ in your answer cannot be decomposed further into simpler partial derivatives. **Do not evaluate the derivatives.** Be sure to specify the correct subscripts in your answer.

$\frac{\partial \ell}{\partial w_E} =$

$\frac{\partial \ell}{\partial w_E} = \frac{\partial \ell}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial z_4} \frac{\partial z_4}{\partial w_E}$

(b) The network diagram from above is repeated here for convenience: What is the

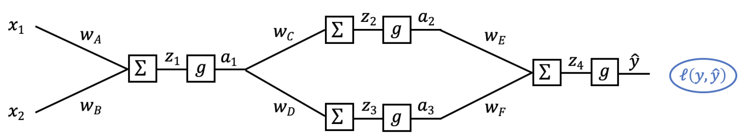

Figure 9: Neural Network

chain of partial deriviatives needed to calculate the derivative $\frac{\partial \ell}{\partial w_C}$?
Your answer should be in the form:

$$\frac{\partial \ell}{\partial w_C} = \frac{\partial ?}{\partial ?} \frac{\partial ?}{\partial ?} \dots$$

Make sure each partial derivative $\frac{\partial ?}{\partial ?}$ in your answer cannot be decomposed further into simpler partial derivatives. **Do not evaluate the derivatives**. Be sure to specify the correct superscripts in your answer.

$\frac{\partial \ell}{\partial w_C} =$

$\frac{\partial \ell}{\partial w_C} = \frac{\partial \ell}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial z_4} \frac{\partial z_4}{\partial a_2} \frac{\partial a_2}{\partial z_2} \frac{\partial z_2}{\partial w_C}$

(c) We want to modify our neural network objective function to add an L2 regularization term on the weights. The new objective is:

$$\ell(y, \hat{y}) + \lambda \frac{1}{2} \|w\|_2^2$$

where λ (lambda) is the regularization hyperparamter and $\mathbf{w}$ is all of the weights in the neural network stacked into a single vector, $\mathbf{w} = [w_A, w_B, w_C, w_D, w_E, w_F]^T$.

Write the right-hand side of the new gradient descent update step for weight w_C given this new objective function. You may use $\frac{\partial \ell}{\partial w_C}$ in your answer.

Update: $w_C \leftarrow \dots$

Update for w_C: $w_C \leftarrow w_C - \alpha \left(\frac{\partial \ell}{\partial w_C} + \lambda w_C \right)$

5. Backpropagation in neural networks can lead to slow or unstable learning because of the vanishing or exploding gradients problem. Understandably, Neural the Narwhal does not believe this. To convince Neural, Lamar Jackson uses the example of an N layer neural network that takes in a scalar input x, and where each layer consists of a single neuron. More formally, $x = o_0$, and for each layer $i \in \{1, 2, ..., N\}$, we have

$$s_i = w_i o_{i-1} + b_i$$
$$o_i = \sigma(s_i)$$

where σ is the sigmoid activation function. Note that w_i, b_i, o_i, s_i are all scalars.

i. (1 point) Give an expression for $\frac{\partial o_N}{\partial w_1}$. Your expression should be in terms of the s_i's, the w_i's, N, x_i, and $\sigma'(\cdot)$, the derivative of the sigmoid function.

$$\begin{aligned}\frac{\partial o_N}{\partial w_1} &= \frac{\partial o_N}{\partial o_{N-1}} \frac{\partial o_{N-1}}{\partial o_{N-2}} \cdots \frac{\partial o_1}{\partial w_1} \\ &= \frac{\partial o_1}{\partial w_1} \prod_{i=2}^{N} \frac{\partial o_i}{\partial o_{i-1}} \\ &= \sigma'(s_1) x \prod_{i=2}^{N} \sigma'(s_i) w_i\end{aligned}$$

EA Comments: I think we should also accept the second line of this for full credit.

ii. (1 point) Knowing that $\sigma'(\cdot)$ is at most $\frac{1}{4}$ and supposing that all the weights are 1 (i.e. $w_i = 1$ for all i), give an upper bound for $\frac{\partial o_N}{\partial w_1}$. Your answer should be in terms of x and N.

$\frac{\partial o_N}{\partial w_1} \leq x \left(\frac{1}{4}\right)^N$

6 Learning Theory

1. **True and Sample Errors:** Consider a classification problem with distribution D and target function $c^* : \mathcal{R}^d \mapsto \pm 1$. For any sample S drawn from D, answer whether the following statements are true or false, along with a brief explanation.

   a) **True or False:** For a given hypothesis space $\mathcal{H}$, it is possible to define a sufficient number of examples in S such that the true error is within a margin of ϵ of the sample error for all hypotheses $h \in H$ with a given probability.

   True. This is exactly what the PAC bound allows us to do although depending on what case we're in (realizable vs. agnostic, infinite vs. finite) the exact bound will vary.

   b) **True or False:** The true error of any hypothesis h is an upper bound on its training error on the sample S.

   False. We said true error is close to training error, but it might be smaller than training error, so it is not an upper bound.

2. Let X be the feature space and D be a distribution over X. We have a training data set

$$\mathcal{D} = \{(x_1, c^\star(x_1)), \cdots, ((x_N, c^\star(x_N)))\},$$

   x_i i.i.d from D. We assume labels $c^\star(x_i) \in \{-1, 1\}$.

   Let $\mathcal{H}$ be a hypothesis class and let $h \in \mathcal{H}$ be a hypothesis. In this question we restrict ourselves to $\mathcal{H}$. We use

$$err_S(h) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}(h(x_i) \neq c^\star(x_i))$$

   to denote the training error and

$$err_D(h) = P_{x \sim D}(h(x) \neq c^\star(x))$$

   to denote the true error. Recall that if the concept class is finite, in the realizable case

$$m \geq \frac{1}{\epsilon}\left[\ln(|\mathcal{H}|) + \ln\left(\frac{1}{\delta}\right)\right]$$

   labeled examples are sufficient so that with probability at least $1 - \delta$, all $h \in \mathcal{H}$ with $err_D(h) \geq \epsilon$ have $err_S(h) > 0$; in the agnostic case,

$$m \geq \frac{1}{2\epsilon^2}\left[\ln(|\mathcal{H}|) + \ln\left(\frac{2}{\delta}\right)\right]$$

labeled examples are sufficient such that with probability at least $1-\delta$, all $h \in \mathcal{H}$ have $|err_D(h) - err_S(h)| < \epsilon$.

a) Briefly describe the difference between the realizable case and agnostic case.

Realizable- the true classifier c^* is in $\mathcal{H}$.

Agnostic- we don't know whether c^* is in $\mathcal{H}$. It may or may not be.

b) What is the full name of PAC learning? How do ϵ and δ tie into the name?

"Probably approximately correct." The hypotheses we find with m examples are *probably* (with probability $p \geq 1-\delta$) *approximately* correct, with $err_D(h) \leq \epsilon$

c) **True or False:** Consider two finite hypothesis sets $\mathcal{H}_1$ and $\mathcal{H}_2$ such that $\mathcal{H}_1 \subset \mathcal{H}_2$. Let $h_1 = \arg\min_{h\in\mathcal{H}_1} err_S(h)$ and $h_2 = \arg\min_{h\in\mathcal{H}_2} err_S(h)$. Because $|\mathcal{H}_2| \geq |\mathcal{H}_1|$, $err_D(h_2) \geq err_D(h_1)$.

False. Since there are more hypotheses in $\mathcal{H}_2$ there might be one that better fits the data than those in $\mathcal{H}_1$.

3. **Fill in the Blanks:** Complete the following sentence by circling one option in each square (options are separated by "/"s):

In order to prove that the VC-dimension of a hypothesis set $\mathcal{H}$ is D, you must show that $\mathcal{H}$ [can / cannot] shatter [any set / some set / multiple sets] of D data points and [can / cannot] shatter [any set / some set / multiple sets] of $D+1$ data points.

In order to prove that the VC-dimension of a hypothesis set $\mathcal{H}$ is D, you must show that $\mathcal{H}$ can shatter some set of D data points and cannot shatter any set of $D+1$ data points.

4. Consider the hypothesis set $\mathcal{H}$ consisting of all positive intervals in $\mathbb{R}$, i.e. all hypotheses of the form $h(x; a, b) = \begin{cases} +1 & \text{if } x \in [a, b] \\ -1 & \text{if } x \notin [a, b] \end{cases}$

   a) **Short Answer:** In 1-2 sentences, briefly justify why the VC dimension of $\mathcal{H}$ is less than 3.

   We only need to show 3 points cannot be shattered. Consider 3 points where the two outer points have label +1 and the middle point has label -1.

   b) **Select one:** What is the VC dimension of $\mathcal{H}$?

   ◯ 0

   ◯ 1

   ◯ 2

   C

   c) **Numerical Answer:** Now, consider hypothesis sets $\mathcal{H}_k$ indexed by k, such that $\mathcal{H}_k$ consists of all hypotheses formed by k **non-overlapping** positive intervals in $\mathbb{R}$. Give an expression for the VC dimension of $\mathcal{H}_k$ in terms of k.

   Hint: Think about how to repeatedly apply the result you found in Part (b).

   $2k$

5. **Select one:** Your friend, who is taking an introductory ML course, is preparing to train a model for binary classification. Having just learned about PAC Learning, she informs you that the model is in the finite, agnostic case.

   Now she wants to know how changing certain values will change the number of labelled training data points required to satisfy the PAC criterion. For each of the following changes, determine whether the sample complexity will increase, decrease, or stay the same.

   i. (1 point) Using a simpler model (decreasing $|\mathcal{H}|$)

   ◯ Sample complexity will increase

- ◯ Sample complexity will decrease
- ◯ Sample complexity will stay the same

B

ii. (1 point) Choosing a new hypothesis set $\mathcal{H}^*$, such that $|\mathcal{H}^*| = |\mathcal{H}|$

- ◯ Sample complexity will increase
- ◯ Sample complexity will decrease
- ◯ Sample complexity will stay the same

C

iii. (1 point) Decreasing δ

- ◯ Sample complexity will increase
- ◯ Sample complexity will decrease
- ◯ Sample complexity will stay the same

A

iv. (1 point) Decreasing ϵ

- ◯ Sample complexity will increase
- ◯ Sample complexity will decrease
- ◯ Sample complexity will stay the same

A