

Question Paper

Exam Date & Time: 24-May-2023 (10:00 AM - 01:00 PM)



MANIPAL ACADEMY OF HIGHER EDUCATION

Manipal School of Information Sciences (MSIS), Manipal
Second Semester Master of Engineering - ME (Artificial Intelligence and Machine Learning) Degree Examination - May 2023

Deep Learning [AML 5202]

Marks: 100

Duration: 180 mins.

Wednesday, May 24, 2023

Answer all the questions.

- 1) [10 points] [TLO 1.1, CO 1] Suppose we have 10^3 samples corresponding to 3 output labels and that each sample is a $3 \times 28 \times 28$ color image. If we apply the softmax algorithm to this data, what are the dimensions of the following quantities assuming that the bias-trick pre-processing has been performed: (10)

- Data matrix;
- Weight matrix;
- Probability matrix;
- Adjusted probability matrix;
- Total average data loss;
- Regularization loss;
- Gradient of total loss w.r.t. the weight matrix?

- 2) [10 points] [TLO 1.2, CO 1] Consider the following dataset for binary classification: (10)

$$x^{(1)} = \begin{bmatrix} 2 \\ 5 \end{bmatrix}, x^{(2)} = \begin{bmatrix} 3 \\ 1 \end{bmatrix}, x^{(3)} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \\ y^{(1)} = 1, y^{(2)} = 0, y^{(3)} = 0.$$

Calculate the SVM loss using bias-trick for the following weights and bias values:

$$W = \begin{bmatrix} 0.1 & -0.4 \\ 0.5 & -0.3 \end{bmatrix}, b = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

- 3) [10 points] [TLO 3.2, CO 3] Consider the same dataset, weights, and bias values from the previous problem. Calculate the Softmax loss using a regularization strength of 0.3. (10)

- 4) [10 points] [TLO 2.2, CO 2] Using the same setup from the previous problem, perform one step of gradient descent with learning rate = 10^{-2} . Round all numbers to 2 decimal places. (10)

- 5) (10)

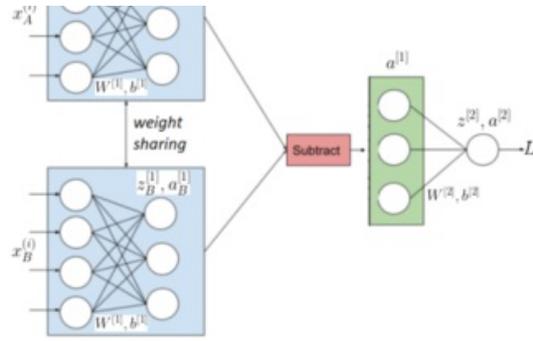
[10 points] [TLO 1.2, CO 1] For a particular binary classification task, your friend modifies the logistic regression loss function as follows:

$$L = \alpha [-y \log(a^{[l]})] - \beta [(1 - y) \log(1 - a^{[l]})],$$

where the model has one input layer and one output layer that is sigmoid-activated. Note that your friend introduces the (unknown) coefficients α and β that have to be fine-tuned like any hyperparameter. What kind of a binary classification task might your friend be solving? Give a simple problem setup and reasonable values for the two parameters.

- 6) [10 points] [TLO 3.1, CO 2] A Siamese neural network consists of twin networks which accepts distinct inputs but share the same weights. The outputs of the twin networks are usually joined later on by one or more layers. The following image shows such a network with a pair of distinct input samples $x_A^{(i)}$ and $x_B^{(i)}$ both of which are of size 4×1 : (10)





Note the following:

- $x_A^{(i)}$ and $x_B^{(i)}$ together constitute the i th input sample for the network;
- the weights for the first hidden layer is the same ($W^{[1]}$) as shown in the blue regions;
- $ReLU$ activation is used in the hidden layer and sigmoid activation is used in the output layer.

Fill in the question marks below for the forward propagation for the i th sample leading to the loss L_i :

$$\begin{aligned}
 ? &= W^{[1]}x_A^{(i)} + b^{[1]} \\
 a_A^{[1]} &= ?, \\
 z^{[2]} &= W^{[1]} \times ? + ?, \\
 ? &= ReLU(z^{[2]}) \\
 ? &= a_A^{[1]} - a_B^{[1]}, \\
 z^{[2]} &= ? \times a^{[1]} + b^{[2]}, \\
 a^{[2]} &= ?, \\
 L_i &= -(y^{(i)} \times ? + (1-?) \times \log(1 - a^{[2]})).
 \end{aligned}$$

7) [10 points] [TLO 3.1, CO 2] For the Siamese neural network in the previous problem, draw a tree diagram starting with L_i on the top and ending with the weights $W^{[1]}$ and bias $b^{[1]}$ by clearly showing the intermediate variables. (10)

8) (10)

[10 points] [TLO 3.2, CO 3] Using the tree diagram for the Siamese neural network in the previous problem, we want to compute the gradient $\nabla_{W^{[1]}}(L_i)$.

- The gradient $\nabla_{W^{[1]}}(L_i)$ ends up being a combination of quantities as follows:

$$\left[\underbrace{T}_{3 \times 4 \times 3 \text{-tensor}} \times \underbrace{D_A}_{3 \times 3 \text{-matrix}} \times \underbrace{I}_{3 \times 3 \text{-matrix}} - \underbrace{T}_{3 \times 4 \times 3 \text{-tensor}} \times \underbrace{D_B}_{3 \times 3 \text{-matrix}} \times \underbrace{I}_{3 \times 3 \text{-matrix}} \right] \times \underbrace{A}_{3 \times 1 \text{-matrix}} \times \text{constant}.$$

Match the following with the quantities in the expression above:

$$(1) \nabla_{W^{[1]}}(z_A^{[1]}) \quad (2) \nabla_{W^{[1]}}(z_B^{[1]}) \quad (3) \nabla_{a^{[1]}}(z^{[2]}) \quad (4) \nabla_{z^{[2]}}(L_i) \quad (5) \nabla_{a_A^{[1]}}(a^{[1]}) \quad (6) \nabla_{a_B^{[1]}}(a^{[1]}) \quad (7) \nabla_{z_A^{[1]}}(a_A^{[1]}) \quad (8) \nabla_{z_B^{[1]}}(a_B^{[1]}) .$$

- Compute all the gradients *except* $\nabla_{W^{[1]}}(z_A^{[1]})$ and $\nabla_{W^{[1]}}(z_B^{[1]})$, and use fact that

$$\nabla_{W^{[1]}}(z_{A \text{ or } B}^{[1]}) \times \text{vector} = \nabla_{W^{[1]}}(W^{[1]}x_{A \text{ or } B}^{(i)} + b^{[1]}) \times \text{vector} = \text{vector} \times (x_{A \text{ or } B}^{(i)})^T$$

to fill-in the missing entries below:

$$\nabla_{W^{[1]}}(L_i) = \begin{bmatrix} ? & ? & ? \\ 0 & ? & ? \\ ? & ? & I(z_{A_3}^{[1]}) - I(z_{B_3}^{[1]}) \end{bmatrix} (??)^T (a^{[2]} - ?) \left(? - (x_B^{(i)})^T \right).$$

9) [10 points] [TLO 4.1, CO 3] In the schematic given below for batch processing, fill in the question marks in the red boxes: (10)

Number of samples = 11
 Number of iterations = 10
 Batch size = 3
 Number of epochs (a.k.a. number of batches) = 4

Epoch 1:

 Iteration number 1:

The notation $L(W)_{W^0}$ means L evaluated at W^0 which represents initial weights

Loss: $L = \frac{1}{2}(L_e + L_o + L_s)$

[7 10 0]
Iteration number 2:
[6 1 3]
Iteration number 3:
[8 5 2]
Iteration number 4:
[4 9]
Epoch 2:
Iteration number 5:
[3 7 5]
Iteration number 6:
[8 4 9]
Iteration number 7:
[2 10 6]
Iteration number 8:
[0 1]
Epoch 3:
Iteration number 9:
[9 6 10]
Iteration number 10:
[4 8 3]

$$\text{and Gradient: } \nabla_W(L) = \frac{1}{3} (\nabla_W(L_2) + \nabla_W(L_{10}) + \nabla_W(L_0)) \Big|_{W^0}$$

$$\text{and Update: } W^1 = W^0 - \alpha \nabla(W) \Big|_{W^0} = W^0 - \alpha \left(\frac{1}{3} (\nabla_W(L_2) + \nabla_W(L_{10}) + \nabla_W(L_0)) \Big|_{W^0} \right)$$

$$\text{Loss: } L = \frac{1}{3} (L_0 + L_1 + L_2) \Big|_{W^1}$$

$$\text{and Gradient: } \nabla_W(L) = \frac{1}{3} (\nabla_W(L_0) + \nabla_W(L_1) + \nabla_W(L_2)) \Big|_{W^1}$$

$$\text{and Update: } W^2 = W^1 - \alpha \nabla(W) \Big|_{W^1} = W^1 - \alpha \left(\frac{1}{3} (\nabla_W(L_0) + \nabla_W(L_1) + \nabla_W(L_2)) \Big|_{W^1} \right) = W^1 - \alpha \left(\frac{1}{3} (\nabla_W(L_2) + \nabla_W(L_{10}) + \nabla_W(L_0)) \Big|_{W^1} \right) - \alpha \left(\frac{1}{3} (\nabla_W(L_0) + \nabla_W(L_1) + \nabla_W(L_2)) \Big|_{W^1} \right)$$

$$\text{Update: } W^3 = W^0 - ?$$

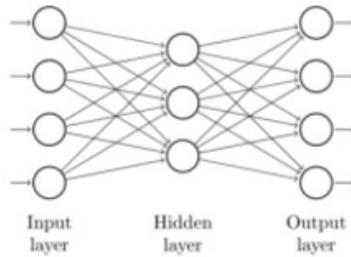
$$\text{Update: } W^4 = W^0 - ?$$

What is the update $W^1 = W^0 - ?$ if batch size = number of samples = 11 and number of iterations = 1?

10)

(10)

[10 points] [TLO 4.2, CO 4] An autoencoder is a neural network designed to learn feature representations in an unsupervised manner. Unlike a standard multi-layer network, an autoencoder has the same number of nodes in its output layer as its input layer. An autoencoder is trained to reconstruct its own input, that is, to minimize the reconstruction error. A simple autoencoder architecture is shown below:



- Assuming that the input sample is $x^{(i)}$, write down the forward propagation equations starting with $z^{[1]}$ and ending with the loss L_i . Since the loss measures how well the sample $x^{(i)}$ can be reconstructed, find a simple expression for it involving $a^{[2]}$ and $x^{(i)}$.
- The gradient descent update for $W^{[1]}$ can be written as:

$$W^{[1]} = W^{[1]} + \alpha (-\nabla_{W^{[1]}}(L_i))$$

$$= W^{[1]} - \alpha [D_1 \times (W^{[2]})^T \times D_2 \times 2(a^{[2]} - x^{(i)})] (x^{(i)})^T,$$

$$\text{where } D_1 = \begin{bmatrix} \sigma(z_1^{[1]}) (1 - \sigma(z_1^{[1]})) & 0 & 0 \\ 0 & \sigma(z_2^{[1]}) (1 - \sigma(z_2^{[1]})) & 0 \\ 0 & 0 & \sigma(z_3^{[1]}) (1 - \sigma(z_3^{[1]})) \end{bmatrix},$$

$$\text{and } D_2 = \begin{bmatrix} \sigma(z_1^{[2]}) (1 - \sigma(z_1^{[2]})) & 0 & 0 & 0 \\ 0 & \sigma(z_2^{[2]}) (1 - \sigma(z_2^{[2]})) & 0 & 0 \\ 0 & 0 & \sigma(z_3^{[2]}) (1 - \sigma(z_3^{[2]})) & 0 \\ 0 & 0 & 0 & \sigma(z_4^{[2]}) (1 - \sigma(z_4^{[2]})) \end{bmatrix}.$$

Using this, answer when can slow learning happen? That is, when does $W^{[1]}$ get updated very slowly?

-----End-----