Unit 8 · Neural Networks

WORKSHEET

Full marks 31

Name _____ Class _____ Roll no. ____ Date _____ Marks _____ / 31

1 Key words [4]

Fill in the blanks with words from the box. Two words in the box are not needed. Each blank carries 0.5 mark.

activation function	backpropagation	bias	epoch	forward pass	hidden layer
learning rate	overfitting	pixel	weight		

A neuron multiplies each input by its _____, adds them up, and adds a _____, a number that sets how easily it says yes. The rule that turns this weighted sum into the neuron's output, such as a step or a sigmoid, is its _____. A layer of neurons whose outputs are not the network's answer is a _____, and working out the answer layer by layer, from the inputs to the output, is the _____. In gradient descent, the _____ sets how big each step is. One pass through all the training examples is an _____. A network that does very well on its training examples but much worse on new ones shows _____.

2 Neurons as gates [2]

Match the weights and bias of each step neuron, in Column A, with the gate it acts as, in Column B. Two items in Column B are not needed. Each carries 0.5 mark.

Column A		Column B
1. weights 1 and 1, bias −0.5	☐	**a.** AND
2. weights 1 and 1, bias −1.5	☐	**b.** NOT
3. weights −1 and −1, bias 1.5	☐	**c.** OR
4. weight −1, bias 0.5 (one input)	☐	**d.** XOR
		e. NOR
		f. NAND

3 Multiple choice [5]

Choose the correct option for each of the following questions.

i. What is the output of a step neuron with inputs 1 and 1, weights 0.5 and −2, and bias 1? [1]

a. 1 **b.** 0
c. −0.5 **d.** 1.5

ii. In a forward pass, what does a hidden layer take as its inputs? [1]

a. the right answers **b.** the outputs of the layer before it
c. the outputs of the layer after it **d.** random numbers

iii. In gradient descent, the slope of the error for a weight is +30 and the learning rate is 0.1. What happens to the weight? **[1]**

a. it grows by 3 **b.** it shrinks by 3

c. it grows by 30 **d.** it stays the same

iv. Which of these is the clearest sign of overfitting? **[1]**

a. training accuracy 60%, test accuracy 58% **b.** training accuracy 100%, test accuracy 70%

c. training accuracy 90%, test accuracy 89% **d.** training accuracy 10%, test accuracy 10%

v. A network has 100 inputs, one hidden layer of 20 neurons and 5 outputs. How many weights (not counting biases) does it have? **[1]**

a. 125 **b.** 2,000

c. 2,100 **d.** 10,000

4 What does it print? [2]

Write exactly what the program prints, in the box under it.

A reading club's neuron decides whom to invite to the next meeting. The inputs are: read the book, came last week, brought a friend (1 for yes, 0 for no).

```
1  def neuron(inputs, weights, bias):
2      total = bias
3      for x, w in zip(inputs, weights):
4          total = total + x * w
5      return 1 if total > 0 else 0
6
7  weights = [2, -1, 3]
8  bias = -2.5
9  for person in [[1, 0, 0], [1, 1, 1], [0, 0, 1], [1, 1, 0]]:
10     print(person, neuron(person, weights, bias))
```

```
OUTPUT
```

5 The forward pass [4]

Read the question, then answer.

A different network of step neurons (step(z) is 1 if z is above 0, else 0):
$h_1 = \text{step}(x_1 + x_2 - 1.5)$, $h_2 = \text{step}(x_1 + x_2 - 0.5)$, $\text{output} = \text{step}(-2\,h_1 + h_2 - 0.5)$.
Fill in the table, writing each weighted sum and its step. Which gate is the whole network?

x_1	x_2	h_1	h_2	output
0	0			
0	1			
1	0			
1	1			

6 Trace the program [5]

Fill in the trace table: one row for each time round the loop.

a. The error of a model with one weight is $(w - 5)^2$, so its slope is $2 \times (w - 5)$. The program takes four steps of gradient descent. [3]

```
1  w = 1
2  rate = 0.25
3  for step in range(1, 5):
4      slope = 2 * (w - 5)
5      w = w - rate * slope
6      print(step, slope, w)
```

step	slope	w after line 5	Output

b. The same program is run with `rate = 1.25`. Work out `w` after the first two steps. What is happening, and what should be done? [2]

7 Write a program [4]

Write the program in Python on the lines. Any program that gives the right result gets the marks.

A two-input step neuron has weights 1 and 1. Write a program that tries every bias from −3 to 3 in steps of 0.5, works out the neuron's output for the four inputs (0, 0), (0, 1), (1, 0) and (1, 1), and prints each bias that makes the neuron an AND gate (outputs 0, 0, 0, 1). It should print `-1.5` and `-1.0`.

8 Short answer [3]

Answer the following question.

A student trains a network on 100 handwritten digits for 50 epochs and reports “my model is 100% accurate”. Give two things wrong with this claim and say how to measure the model properly.

9 Think [2]

Think, then answer.

A school wants a program to work out each student’s grade (A+, A, B+, …) from their total marks, using the board’s published grade boundaries. A student suggests training a neural network on last year’s marks and grades. Is that a good idea?

10 Activity: Is it fair to your handwriting? (not marked)

At a computer or phone, with three or four friends, open **Read my handwriting** in this unit online. Each person draws every digit 0–9 twice, in their usual style, with Tidy on. Tick each one read correctly, then do the same for ०–९.

Person	0–9 read right (out of 20)	०–९ read right (out of 20)	Digits it got wrong