Exercise *Computer Vision (CSCI 4220U)*

Please hand in this paper to the instructor before the end of the lecture.

Name: _____

Student number: _____ Date: _____

Q. Write down the number of parameters for an MLP network that takes in an image of size 640×480 and classifies it into one of three classes. The MLP comprises 2 hidden layers. The size of each hidden layer is 5.

Q. Say, we are solving a *regression* problem. The model makes the following predictions: $\hat{y}_i$ where $i \in [1, N]$. The corresponding ground truth labels are y_i where $i \in [1, N]$. Compute the MSE loss.

Q. Now let's change the previous problem to a *classification* problem: how should we compute the loss for this example?

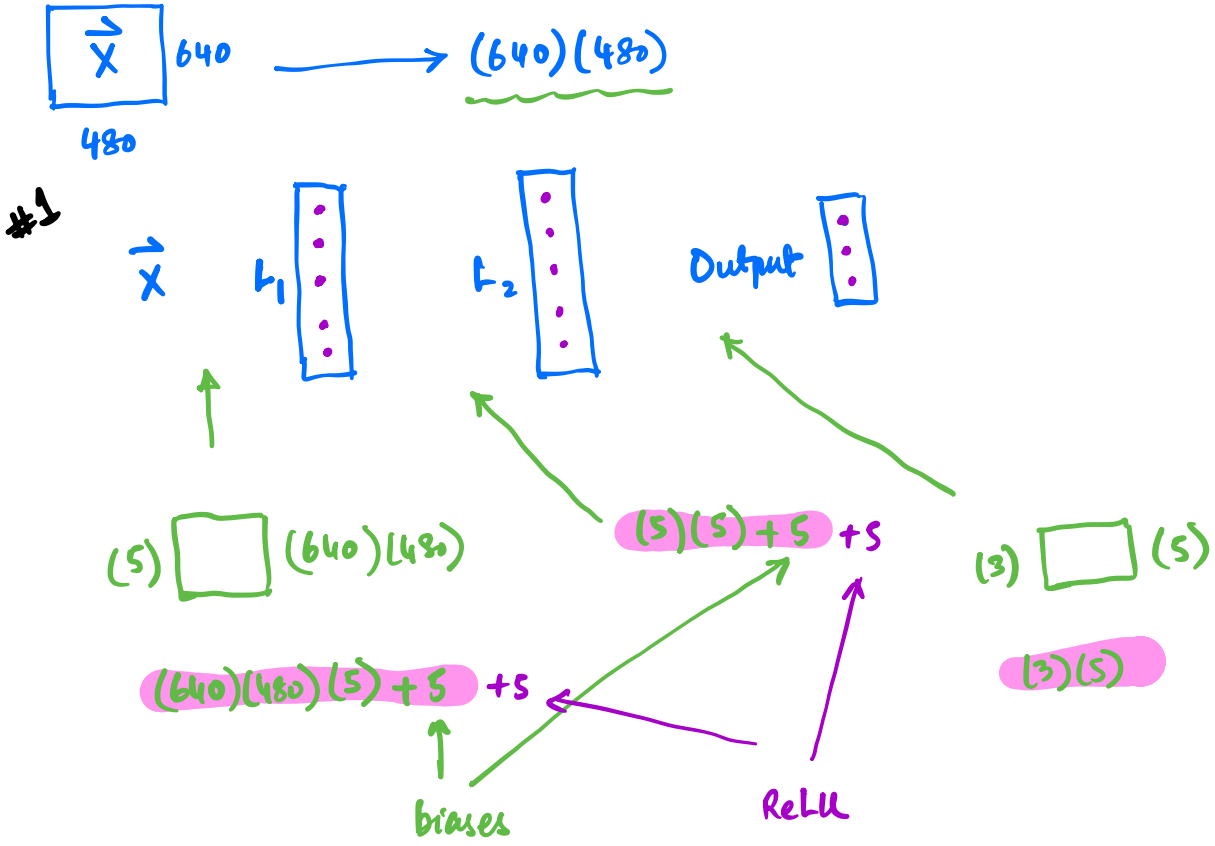

#2 $MSE = \frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2$

#3 Loss used is cross-entropy $H(p,q) = -\sum p(x_i) \log q(x_i)$

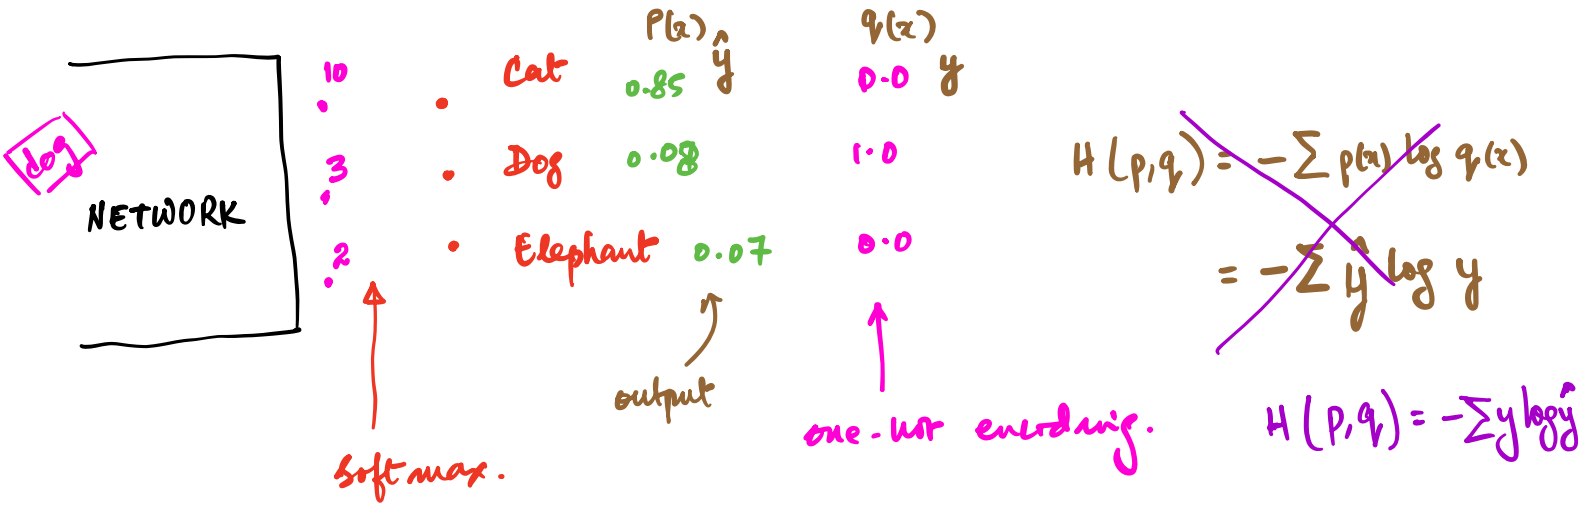
P(x)
ŷ
q(x)
y
dog
NETWORK
10
3
2
Cat
0.85
0.0
Dog
0.08
1.0
Elephant
0.07
0.0
output
Softmax.
one-hot encoding.
H(p,q) = -Σ p(x) log q(x)
= -Σ y log ŷ
H(p,q) = -Σ y log ŷ

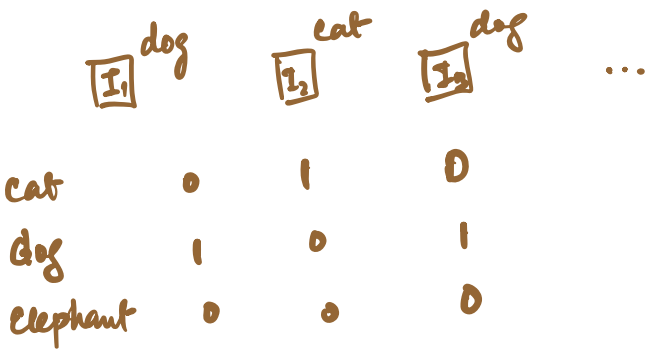
I1
dog
I2
cat
I3
dog
...
cat 0 1 0
dog 1 0 1
elephant 0 0 0