CS6423: Scalable Computing **Name:** _____
Spring 2020
Mid Term Practice Questions

Time Limit: Student ID _____

1. (20 points) You design the 2-layer fully connected neural network shown in Figure 1. All activations are sigmoids and your optimizer is stochastic gradient descent. The loss function is $(\hat{y} - y)^2$. in the network.

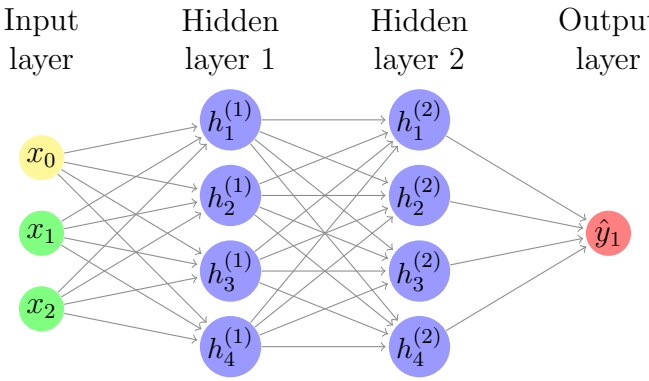

Figure 1: Simple deep network with sigmoid activations

   (a) (10 points) Draw the computation graph for this example.

   (b) (5 points) Compare how Caffe and TensorFlow would implement this example.

   (c) (5 points) Compare the assignment of CPU/GPU processors for Caffe and TensorFlow.

2. (20 points) We have a deep network with input x, target label t output y, hidden nodes z, and loss $\mathcal{L}$ defined by the following:

$$\begin{aligned} h &= Wx + b \\ \hat{y} &= ReLU(h) \\ \mathcal{L} &= 1/2(y - t)^2 \end{aligned}$$

   where ReLU(x) is $\max(0, x)$ for $x > 0$ and 0 otherwise. See the neural network shown in Figure 2

   (a) (10 points) Draw the computation graph for this example.

   (b) (5 points) If we initialise $W = 0$, $b = 0.5$ and input $x = [1\ 1\ 0]^T$ with target $t = 2$, compute the loss for this network.

   (c) (5 points) Compute how the weights W, b are updated for a learning rate $\eta = 1$ for this first iteration of backpropagation.

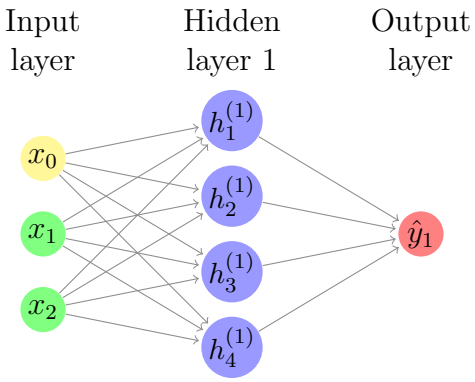

Figure 2: Simple deep network with sigmoid activations