

Week 1 Study Questions

“Training”: (whether you “know” the material)

1. What is the difference between artificial intelligence, machine learning, and deep learning?
2. A neuron has a cell body, axons, and dendrites. What order does information flow through these parts?
3. What is a synapse?
4. What does “feed-forward” mean in the context of an artificial neural network?
5. What does “fully-connect” mean in the context of an artificial neural network?
6. A loss function takes two arguments. What are they?
7. Which of the following are *not* steps in neural network training?
 - computeing the prediction
 - computeing the loss function
 - reseting the parameters to random
 - displaying the image
8. Why do we need both a training set and a test set?
9. Gradient descent is an example of ...
 - a loss function
 - an optimizer
 - a neural network architecture
 - a type of neuron unit

Generalization: (whether you can apply the material)

1. A researcher tried to train a neural network to play the game Mario. The researcher rewarded the network based on the amount of time (in seconds) that Mario stayed alive. Would you expect the network to perform well? Why or why not?
2. How is an artificial neural network similar to a biological neural network? How are they different?
3. Consider the `forward` function of a neural network defined using PyTorch:

```
def forward(self, img):  
    return self.layer5(img.view(-1, 256 * 256))
```

- What is the expected size of the input image?
 - How many layers are in this network?
 - How many hidden units are in this network?
 - How many output units are in this network?
4. Which of the following are classification problems? Binary classification problems?

- Finding traffic lights in an image
- Determining the sex of a duck based on its height, weight, and other features
- Determining the sex of a duck based on its photograph
- Determining the maker of a car based on a photograph (out of the major car makers)
- Translating from English to French

5. Consider this training loop. Which of the following are true?

```
for i in range(50):
    for (input, actual) in data:
        out = network(input)
        loss = criterion(out, actual)
        loss.backward()
        optim.step()
        optim.zero_grad()
```

- The training data contains 50 items
- The training data contains 1 item
- Each training example is used once
- Each training example is used 50 times
- The optimizer used is gradient descent
- The network is a 3-layer neural network
- The variable name of the neural network is `pigeon`
- The variable name of the neural network is `network`
- The variable name of the neural network is `input`

6. How many “synaptic” connections are there in this network? Is this network fully connected?

