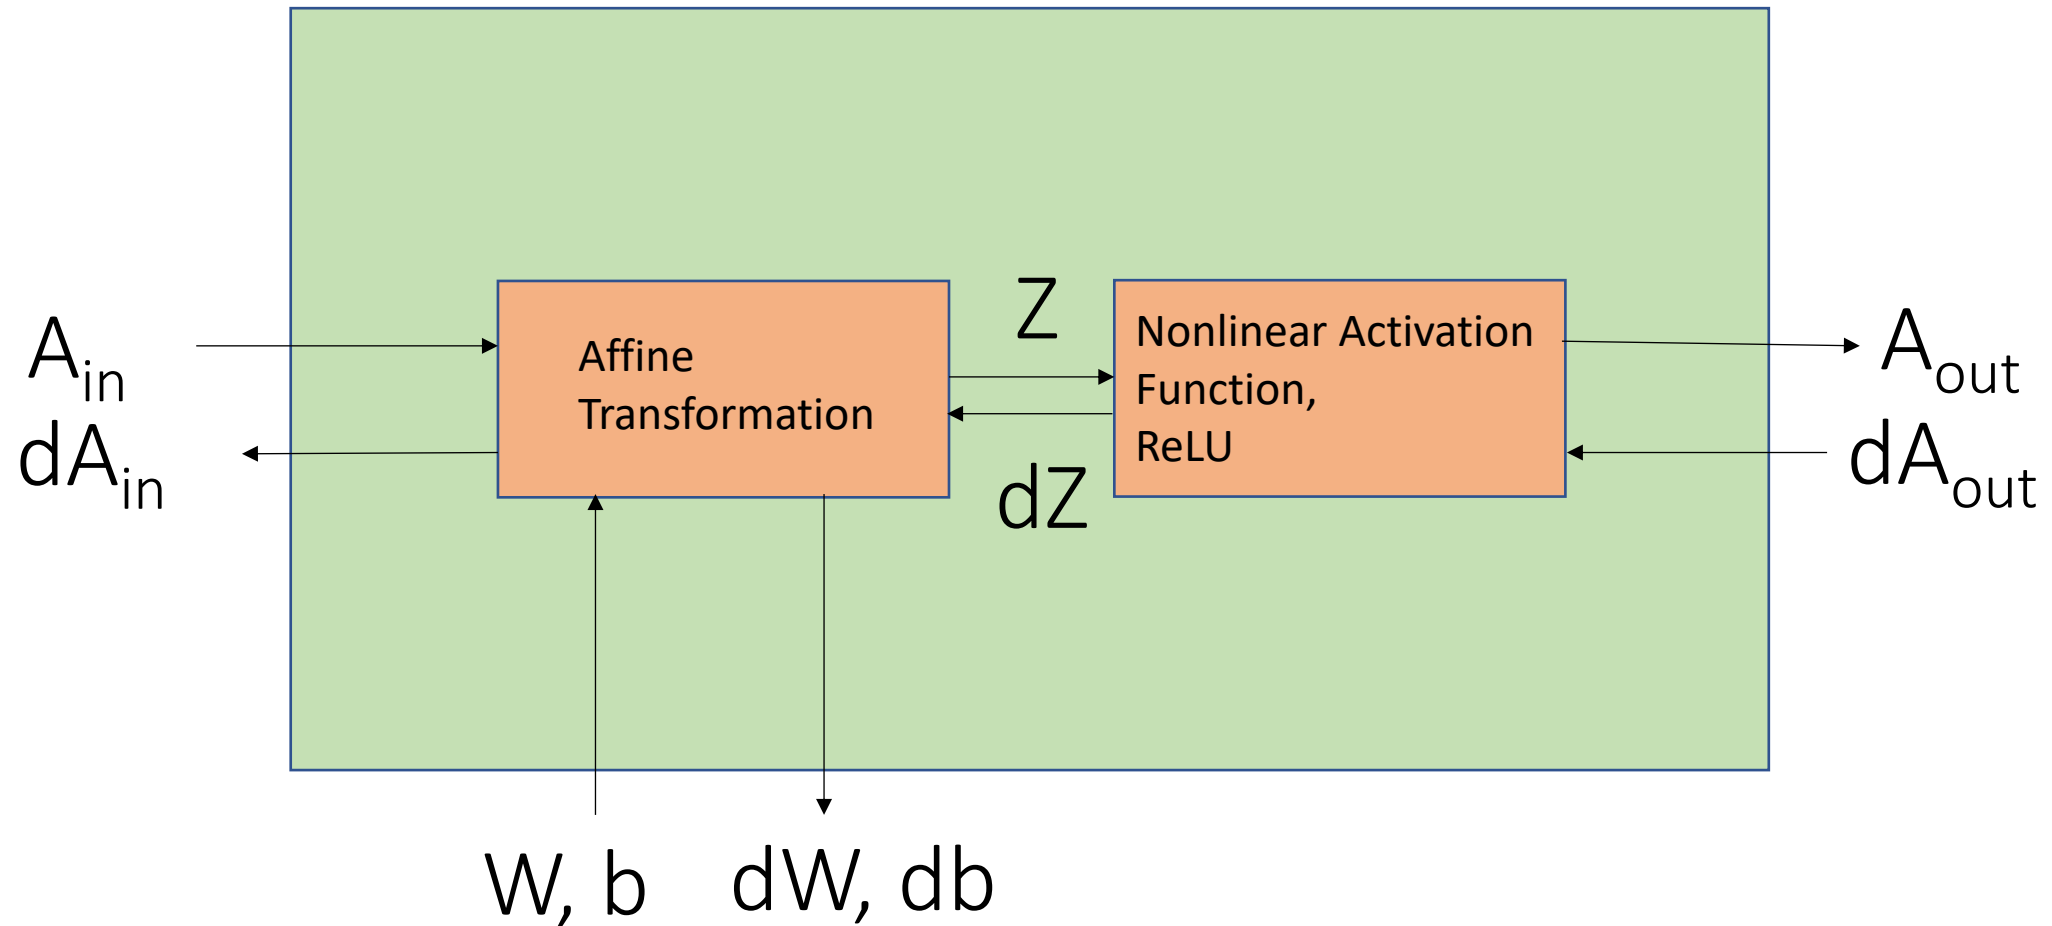


Deep Networks: Putting The Pieces Together

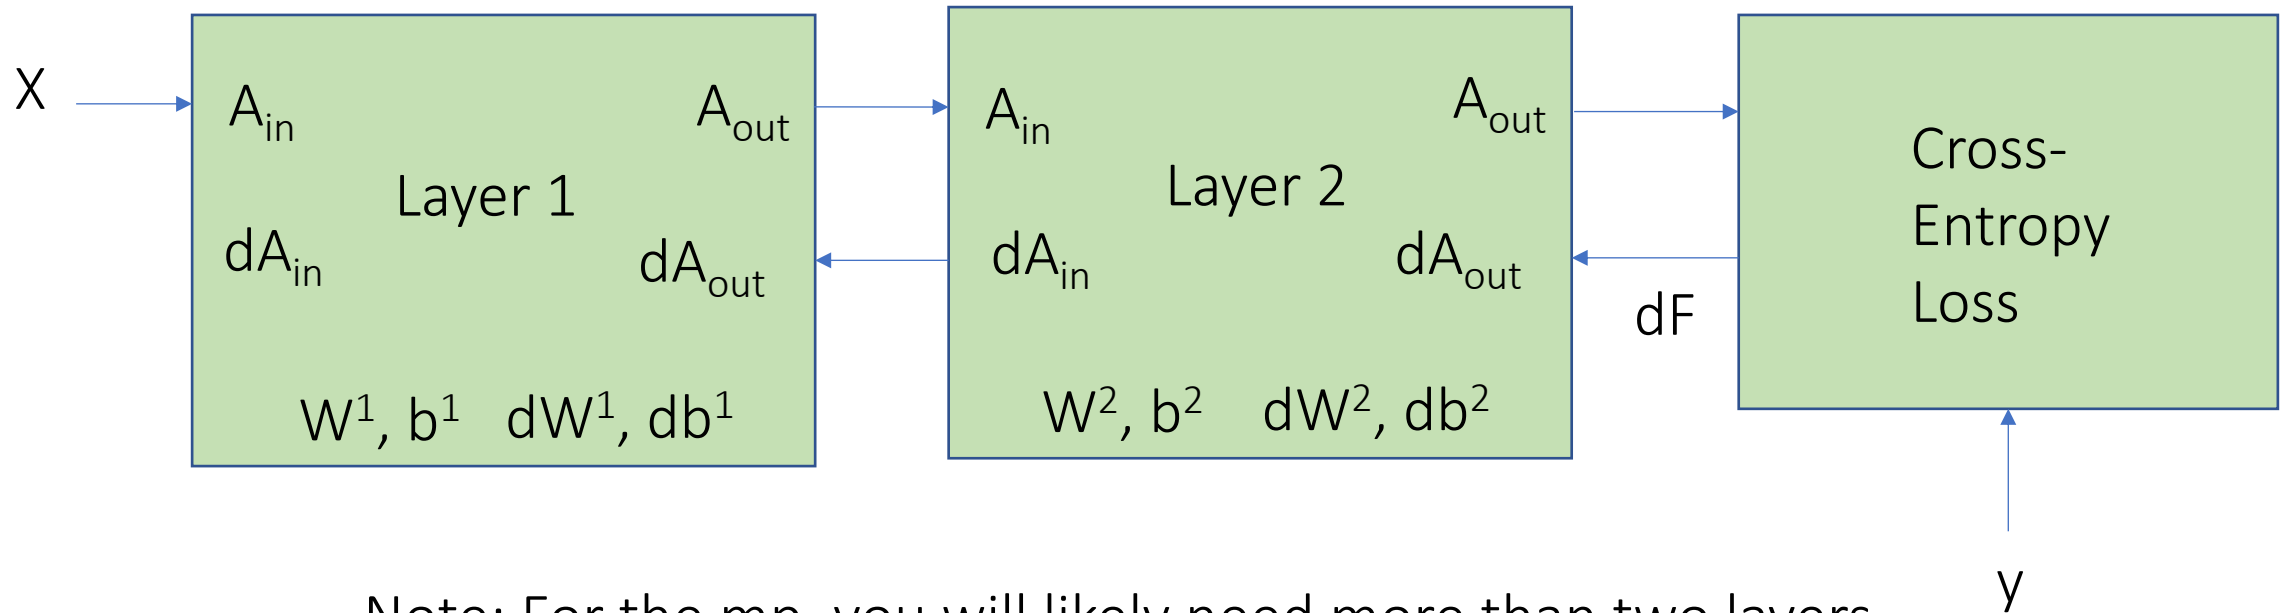
Neural Network Layer



Algorithm 1 Three Layer Network

```
1: procedure THREE-NETWORK( $X, \{W^1, W^2, W^3\}, \{b^1, b^2, b^3\}, y, \text{test}$ )
2:    $Z^1, \text{acache1} = \text{AFFINE-FORWARD}(X, W^1, b^1)$   $\triangleright$   $\text{acache} = \text{affine cache}$ 
3:    $A^1, \text{rcache1} = \text{RELU-FORWARD}(Z^1)$   $\triangleright$   $\text{rcache} = \text{relu cache}$ 
4:    $Z^2, \text{acache2} = \text{AFFINE-FORWARD}(A^1, W^2, b^2)$ 
5:    $A^2, \text{rcache2} = \text{RELU-FORWARD}(Z^2, W^2, b^2)$ 
6:    $F, \text{acache3} = \text{AFFINE-FORWARD}(A^2, W^3, b^3)$ 
7:   if  $\text{test} == \text{true}$  then
8:     classifications = argmax over all classes in logits for each example
9:     return classifications
10:  loss,  $dF = \text{CROSS-ENTROPY}(F, y)$ 
11:   $dA^2, dW^3, db^3 = \text{AFFINE-BACKWARD}(dF, \text{acache3})$ 
12:   $dZ^2 = \text{RELU-BACKWARD}(dA^2, \text{rcache2})$ 
13:   $dA^1, dW^2, db^2 = \text{AFFINE-BACKWARD}(dZ^2, \text{acache2})$ 
14:   $dZ^1 = \text{RELU-BACKWARD}(dA^1, \text{rcache1})$ 
15:   $dX, dW^1, db^1 = \text{AFFINE-BACKWARD}(dZ^1, \text{acache1})$ 
16:  Use gradient descent to update parameters i.e.  $W^1 = W^1 - \eta dW^1$ 
17:  return loss
```

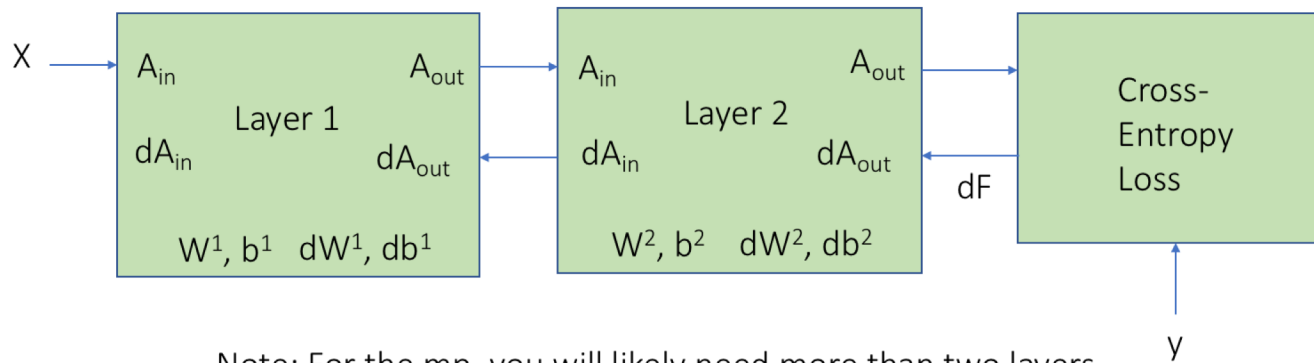
Neural Network Example and Computational Graph



Note: For the mp, you will likely need more than two layers.

Tensorflow and Autodifferentiation

- Autodifferentiation: You only have to define the forward operation. The backwards operation will be automatically computed.
- **You build the computation graph first and then run the graph in a session.**

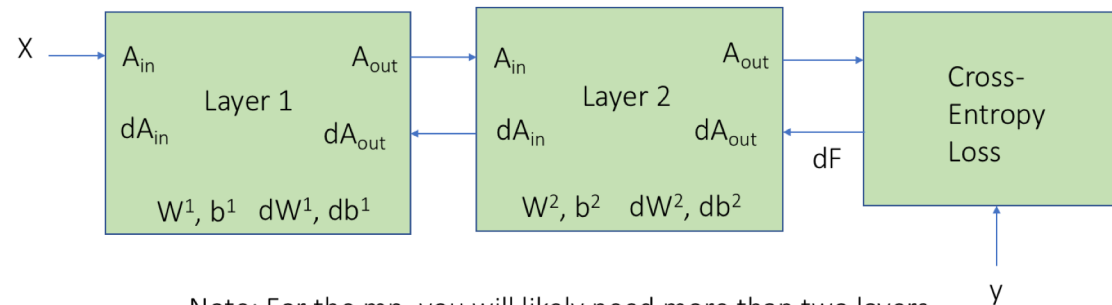


Note: For the mp, you will likely need more than two layers.

Tensorflow Example for Two Layer Net

```
class TwoLayerNet(object):
    def __init__(self, units1, units2, lr):
        self.units1 = units1
        self.units2 = units2
        self.lr = lr

    def define_forward(self, x, y):
        # define graph but does not run graph!
        layer1 = tf.contrib.layers.fully_connected(x, self.units1,
                                                    activation_fn=tf.nn.relu)
        output = tf.contrib.layers.fully_connected(layer1, self.units2,
                                                    activation_fn=None)
        loss = tf.nn.softmax_cross_entropy_with_logits(logits=output, labels=y)
        train_op = tf.train.GradientDescentOptimizer(self.lr).minimize(loss)
        return output, train_op
```



Note: For the mp, you will likely need more than two layers.

Tensorflow Example Continued

```
def main():
    # load dataset
    mnist = input_data.read_data_sets("MNIST_data/", one_hot=True)

    # Create model
    net = TwoLayerNet(128, 10, 0.01) # hidden layer has 100 units, last layer has 10
    x = tf.placeholder(tf.float32, [None, 784]) # placeholder for our data
    y = tf.placeholder(tf.float32, [None, 10])
    output, train_op = net.define_forward(x, y)

    # Run in session
    with tf.Session() as sess:
        # Train model for 1000 iterations
        sess.run(tf.global_variables_initializer())
        for _ in range(1000):
            bx, by = mnist.train.next_batch(100)
            scores, _ = sess.run([output, train_op], feed_dict={x: bx, y: by})
            accuracy = calculate_accuracy(by, scores)
            print('Training Accuracy: ', accuracy)

        # Evaluate model
        scores = sess.run(output, {x: mnist.test.images})
        accuracy = calculate_accuracy(mnist.test.labels, scores)
        print('Testing Accuracy: ', accuracy)
```

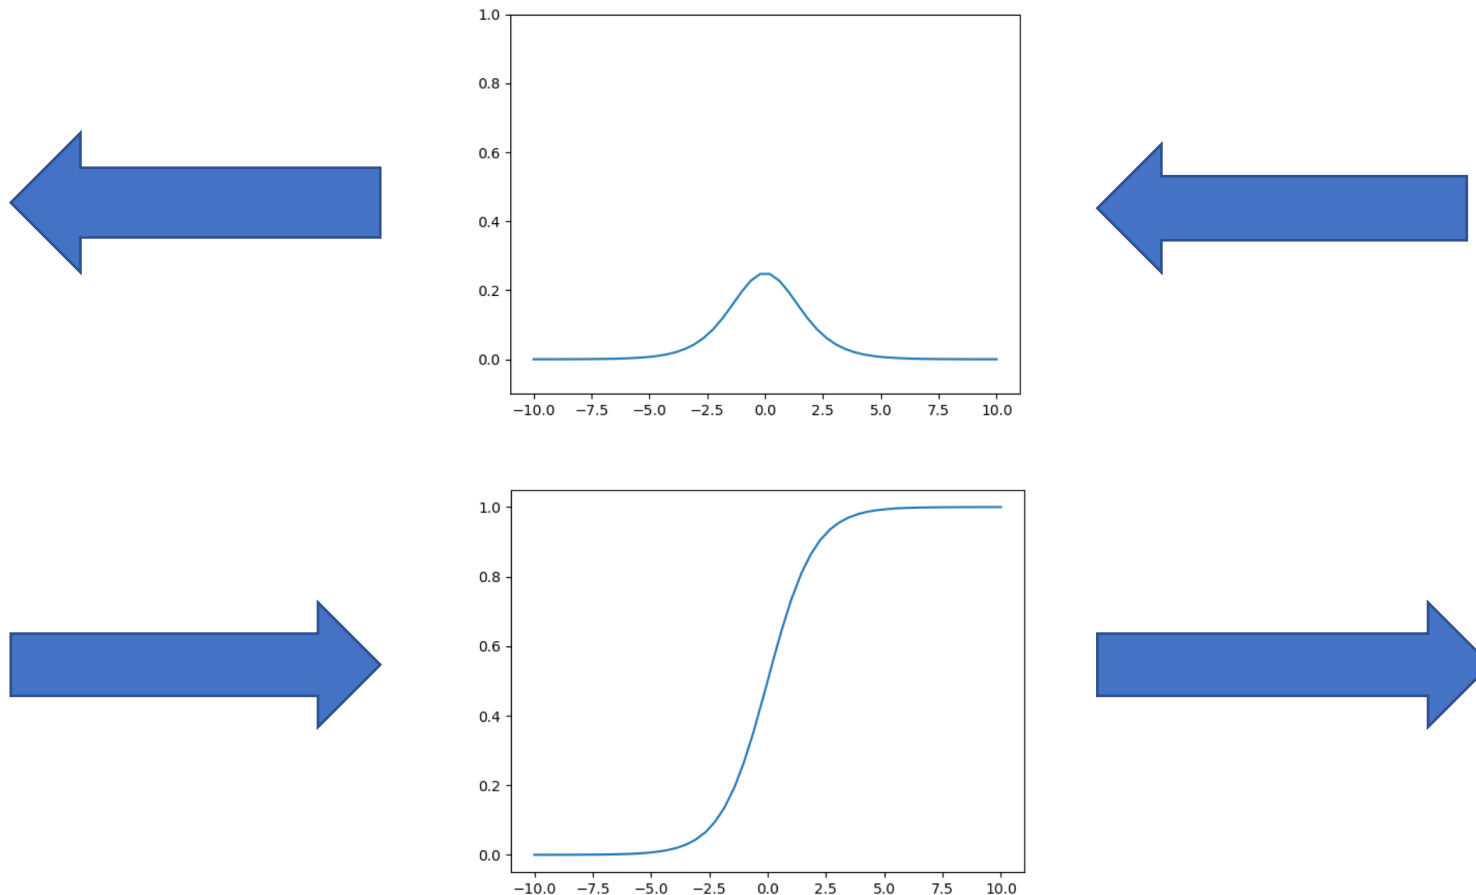
Tensorflow Ops

- Placeholders
- Variables
- Math

Operations

Why it is important to know backpropagation?

- Why do we need to know backpropagation? It helps us avoid bugs!



Terminology of Reinforcement Learning

- Policy function: function that maps a state to an action.
- Q-Function: expected future reward for states and actions.
- Value Function: expected future reward for states.
- Deep learning can learn and approximate functions.

Dataset for MP4

Batch Data X

	Ball-X	Ball-Y	Velocity-X	Velocity-Y	Paddle-Y
X_0	0.296	0.265	-0.080	0.049	0.120
X_1	0.216	0.315	-0.080	0.049	0.080
X_2	0.136	0.364	-0.080	0.049	0.120
X_3	0.056	0.413	-0.080	0.049	0.120

Batch Targets y

Action	
0	y_0
2	y_1
1	y_2
1	y_3

X_{01}

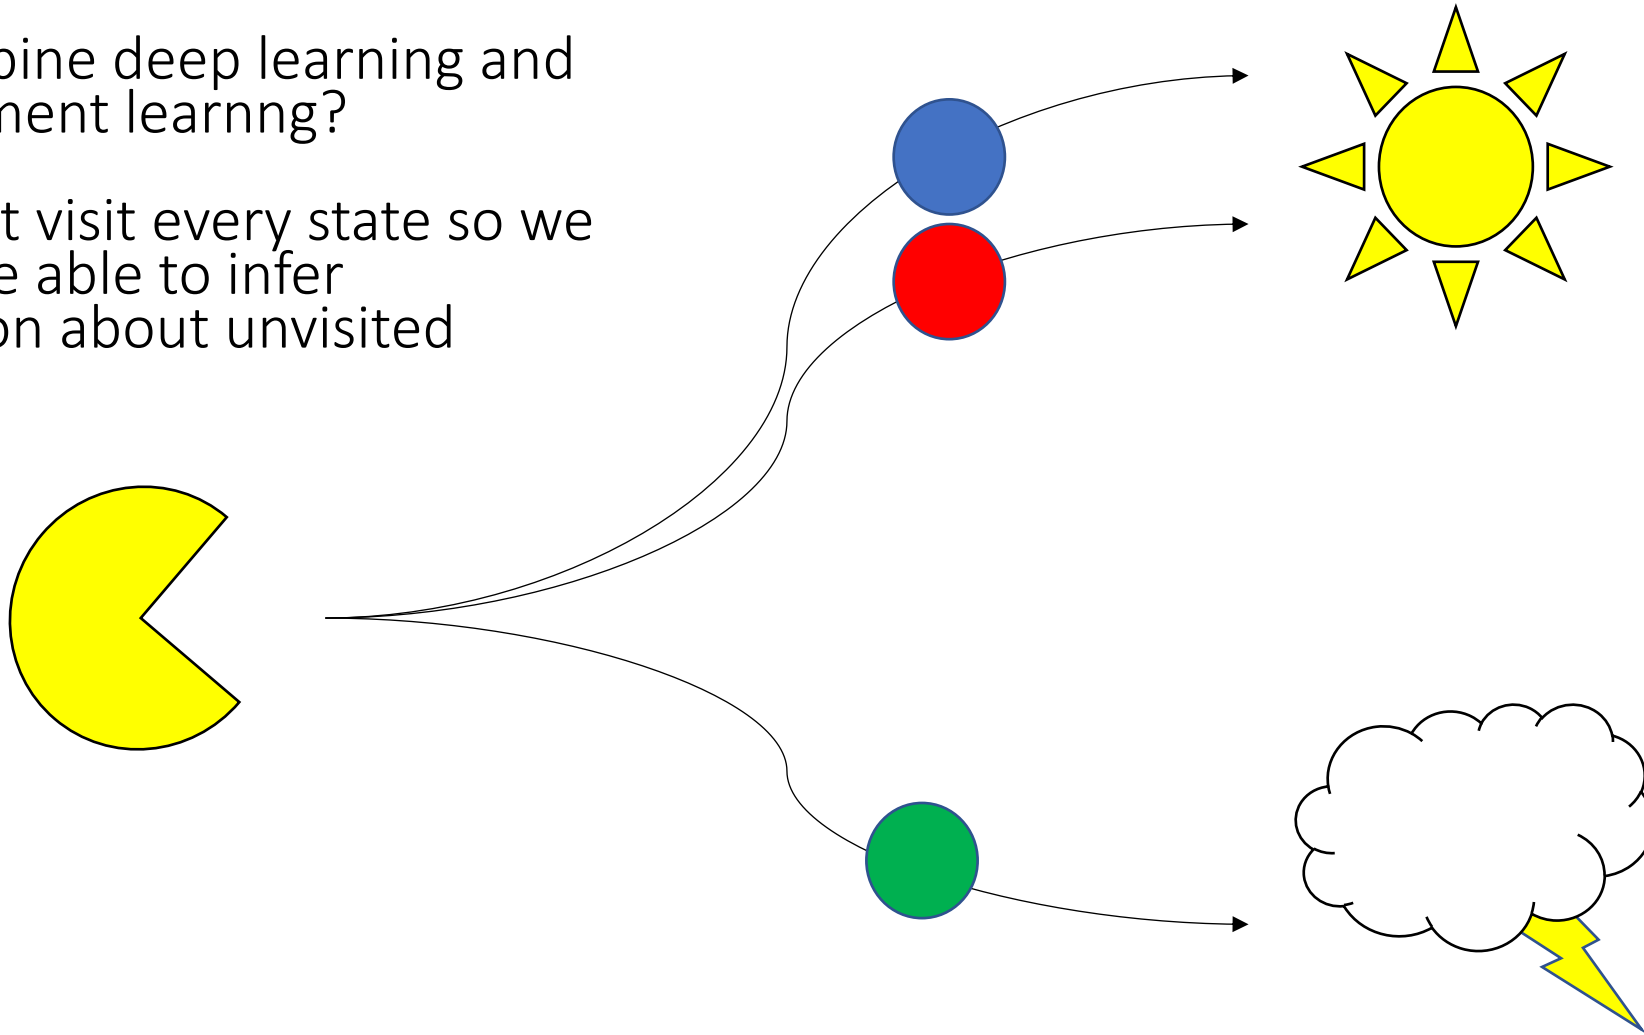
X_{32}

Behavioral Cloning / Imitation Learning
Approximate the policy function of an expert.

What are the advantages of deep learning?

Why combine deep learning and reinforcement learning?

We cannot visit every state so we need to be able to infer information about unvisited states.



Up next: Deep Q-Learning