Lecture 9.3 Exercises

1. The file exercise1.m contains code to learn a simple forward model for a single joint arm with inertia and damping. The estimated acceleration is computed as a weighted combination of the feedback torque and estimated velocity. The weights are adjusted using gradient descent with the difference between the actual and estimated acceleration as the learning signal. Training is run for 20 trials, and the actual and estimated velocity and acceleration are plotted for the first and last trials.

   a. Adjust the learning rate by changing the value of *alpha.* What effect do different values have on learning?

   b. Plot the changes in the weights over the learning trials. How does the model learn to utilize the feedback torque and estimated velocity to estimate acceleration?

   c. Change the model to use feedback error based learning, using the feedback torque as the error signal.

2. Implement reinforcement learning of a forward model for the arm model from exercise 1 using action values.