Abstract

Traditional methods of controlling mechanical systems generally require a fundamental base of knowledge on the system itself to build a model out of first principles. Due to both restrictions on the model accuracy, as well as limitations of the controllers built upon those models, we explore the use of machine learning and its place in controlling mechanical systems. We study four controls problems, an inverted pendulum, simple projectile motion, a firing a trebuchet, and the rotary pendulum. Using reinforcement learning and concepts from differentiable programming, we optimize deep neural networks to map the state vectors of these problems to some possible control action. For static problems such as projectile motion and the trebuchet, we implement a gradient estimation method to manually propagate the gradients through a nondifferentiable component of our program. We test the accuracies of these learned models in simulation to accuracies of 4.3% in the projectile case, and 1.22% in the trebuchet case. For the time-varying problems such as the inverted pendulum and rotary pendulum, we implement a deep Q-network (DQN) to learn the proper control actions. These DQN are compared against the performance of an optimal controller found through LQR.