Learning to Play Kendama

CIS519: Project Proposal

Foster Collins 80470581

Tighe Costa 16238791

Daniel Orol 21768721

Project: We propose teaching a simulated robotic arm to play cup-and-ball, a game where one must try to catch a tethered ball in a cup. This problem is representative of many manipulation scenarios where robots must indirectly control objects with non-trivial dynamics, such as operating machinery, playing ping-pong, or pouring a glass of water. Accurate dynamic models are fundamental to robot control, but physics-based techniques are not adequate for sufficiently complex problems.

Approach: There are a multitude of possible learning approaches to playing cup-and-ball [1]. The three approaches we are considering are as follows: direct, indirect, and imitation learning. We will first look into indirect model learning, which learns both a planning and execution stage of a model to determine the dynamics.

In terms of the actual learning algorithm, we will likely use some form of reinforcement learning, which is a commonly used method for existing solutions to similar problems. Some specific methods we will be looking more into are Q-Learning, Policy Gradient, and Markov Decision processes. The advantage of these techniques are that they use a reward function to learn without requiring a predetermined model of the system physics; however, they are typically used in spaces with discrete actions, which is not ideal for our problem space. These methods are used for similar problems such as learning an inverted pendulum controller. [2]

Datasets: We will be generating our own datasets for this project. We plan on implementing our project in Python and simulating the robot in Gazebo, as we are familiar with both and there is extensive documentation on interfacing between the two. While most Gazebo simulation is done in real-time, it is possible to silence the graphics output of the simulator and obtain large training and testing datasets at computation speed. These tools allow us to continuously monitor the states of the ball and the cup; all the information we need for our dataset.

Project Scope: We believe that it is important to be able to shift the scope of the project based on our progress throughout the rest of the semester. Our fundamental problem is to learn to move the cup so we can catch the attached ball. As a starting point, we plan on directly manipulating the cup (with some velocity and acceleration constraints) and perceiving the ball's ground truth state in the simulator. Once we have successfully learned on this process, we can expand it in a number of ways. For example, we could simulate the cup on a robotic arm learning joint movements or learn modified versions of the game like balero (where one must catch the ball on a spike) or kendama (a version with two cups and one spike).

References

[1] Stefan Schaal. Learning robot control. *The handbook of brain theory and neural networks*, 2:983–987, 2002.

[2] Gerald Van Den Berg. *Aggregation-Based Learning in the Inverted Pendulum Problem.* PhD thesis, Princeton University, 2011.