PLAYING VIDEO GAMES WITH DEEP REINFORCED LEARNING

Masare Akshay Sunil *Varsha Lalwani*

CS365: Artificial Intelligence Programming

March 20, 2015

Abstract

In this project, we attempt to learn control policies directly from high-dimensional sensory input using reinforcement learning. The model would be a convolutional neural network, trained with a variant of Q-learning, whose input would be raw pixels and whose output would be a value function estimating future rewards. We would apply our method to play multiple Atari games from the Arcade Learning Environment[1], with no adjustment of the architecture or learning algorithm.

Motivation

General Game Playing is the branch of Artificial Intelligence that deals with playing multiple games using a single agent. For many years, it has been possible for a computer to play a single game by using some specially designed algorithm for that particular game. But these algorithms were useless outside their context. For example, an algorithm for chess cannot play checkers. Hence, we need General Game Playing agents to play multiple games. In this project we are trying to implement a deep reinforced learning based agent to play multiple video games. Also, if a General Game Playing system is well designed, it can be used in other areas as well, like search and rescue.

Previous work

There have been many attempts in past few years to design general game players using several techniques. The first successful Deep Reinforcement Learning based general Game Player[3] was implemented by Mnih *et. al.* of DeepMind Technologies which was motivated by the success of model-free reinforcement learning approach in a backgammon playing program. Since then, there have been various similar attempts to implement the algorithm. One such example is available here.

Approach

We will try to play the ATARI 2600 video games with Deep Reinforcement Learning. The crux of this type of learning is that it is reward based. We will use Q-Learning method to model this into our agent. Our agent will interact with the video games using the Arcade Learning Environment (ALE), which is based on the Stella Emulator[2] for ATARI 2600, removing the need of having an actual ATARI 2600 to train our agent. The ALE will provide the frames of the game in lower dimensions, easing out the processing. This data will represent a single state in the game. Then the feature extracting is done from sequence of such states using a series of Convolutional Neural Networks. Using these, an Action is chosen, which is performed, then the result of this action is modelled as reward, positive or negative, using Q-Learning.

References

[1] The Arcade Learning Environment: An Evaluation Platform for General Agents by Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling *Journal of Artificial Intelligence Research 47, pp. 253-279, 2013.*

[2] Stella Emulator: http://stella.sourceforge.net/

[3] Playing Atari with Deep Reinforcement Learning by Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou**,** Daan Wierstra, Martin Riedmiller *NIPS Deep Learning Workshop*, 2013.