Tetris AI

Degree programme: BSc in Computer Science | Specialisation: Computer Perception and Virtual Reality
Thesis advisor: Prof. Marcus Hudritsch
Expert: Dr. Harald Studer (Optimo Medical AG)

Artificial neural networks are here to stay, branching out into more and more areas in the pursuit of the ultimate goal, general AI. One of these areas is playing games. Games require a high level of abstraction and intuition, this provides fertile ground to hone networks in. Eventhough current top of the line research has proceeded to more challenging titles, Tetris still provides these same hurdles on top of the open problem of playing it perfectly.

Thomas Michael Voegelin

Introduction

Tetris resides in the sweetspot of difficulty and accesability that allows me to possibly achieve something profound in the given time. As playing Tetris truly perfectly is not computable, perfectly in this context means high-risk, high-reward without reaching game over. Current approaches either only clear single lines, in a feverish attempt to avoid game over, or greedily play for the most simultanious line clears, leading to game overs from time to time.

Objective

As indicated, the goal then was to train a deep neural network to play Tetris as well as possible. Although the ideal mix of risk versus saftey was what was striven for, due to the ristrictions in manpower, computational resources and time the focus lay on the why and how it learned what it did, rather than the end performance achieved. To this end a suitable architecture had to be found, implemented, tested and iterated upon.

Engine

The first step was creating or finding a way to easily interface with a Tetris game. The tought of training with an official retail version of the game was entertained for a time, however due to the emulation and limited access to the code of the game this was ultimately scrapped, in favor of a for machine learning designed version. This primarily affects the way of interfacing with the game not the mechanics, which are save for the phantom (the projected landing) block in line with the basic version of the game.

Architecture

There are an abundance of possible architectures to chose from, however as time was of the essence the first promising candidate Asynchronous Advantage Actor-Critc (A3C) has caried the project through to the end. The only architecture that had to be discarded is A3C's spiritual predecessor, Deep Q-Network (DQN).

Results

It was, given the proper architecture (A3C), rather easy to achieve some points. That is to clear a line or two, a trivial task for a human. With improvements to the way points were awarded internally the AI even became what one could call an amateur. Though super human performance was not achieved, it is thought to

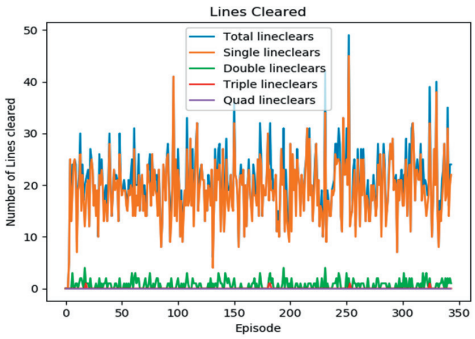

Performance after roughly 1 week of training on a very powerful system. (early training omitted)

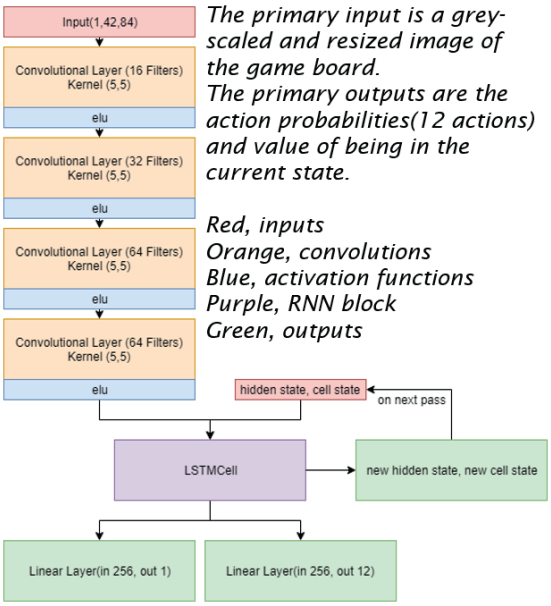

The primary input is a grey-scaled and resized image of the game board.
The primary outputs are the action probabilities(12 actions) and value of being in the current state.

Red, inputs
Orange, convolutions
Blue, activation functions
Purple, RNN block
Green, outputs

Architecture of the neural network.