

Game of Pong AI

Yinghuan Wang

Policy Gradient (PG) and *Reinforcement Learning* (RL) are the frontier of machine learning. PG can help RL solve problems with no immediate rewards which is a traditional puzzle for RL. PG + RL can be applied to Atari games, Go(AlphaGo), even helicopter control. I created the pipeline using numpy then I refactored the policy network using PyTorch which greatly improved learning speed. Initially I wrote the neural network using numpy with RMSProp but later I refactored the neural network using PyTorch framework with Adam and RMSProp. Moreover, the neural network using Adam and RMSProp learns significantly faster than the neural network using numpy and RMSProp.