

# Agent Based Simulation, Negotiation, and Strategy Optimization of Monopoly

TJHSST Computer Systems Lab 2007 - 2008

Nicholas Loffredo

## ABSTRACT

This Agent Based Simulation project attempts to study whether Reinforcement Learning can be an effective method for agents to improve their play towards an optimal strategy, including negotiation, for Monopoly. Monopoly provides a useful test-bed for learning algorithms in a relatively simple environment, yet is still complex enough that many of the results and methods used can be applied to more relevant real-life situations. Various policy variables were implemented, along with associated aggressiveness levels of play, for each policy. The agents attempt to learn via reinforcement learning which policy settings improve the likelihood of winning. To test if an agent is learning, it is pitted against a non-learning agent for thousands of games, and the results are statistically analyzed. The outcome of the test gives very strong evidence that reinforcement learning is working. Using this Agent Based Simulation of Monopoly as a testbed provides a fertile environment for further research.

## INTRODUCTION

Computers currently are unable to perform common human tasks such as understanding a language well enough to speak it and effectively communicate. A good example of this is negotiation. Humans are able to negotiate with one another for various goods. Computers, on the other hand, still have a long way to go in this regard. There is significant research and development ongoing towards computer-based systems being able to negotiate effectively: they could be used in many situations that currently require people, such as in diplomacy, selling/buying goods, trading goods, or just negotiating with other people in general. More importantly, it would allow people to instruct a robot/computer to negotiate using certain items and to meet certain goals, instead of having themselves or hiring other people to do it. These computers would be resistant to common human flaws, such as anger, impatience and/or forgetfulness.

Developing a computer-based system that can learn effectively in a limited environment is a first step towards being able to learn in a more complex one. The game of Monopoly is simple enough that learning should be able to be implemented within a year, yet complex enough that the method used to achieve the results may be able to be applied towards real life situations. By making a working simulation of Monopoly, a learning capability can be implemented for computer agents that will play the game.

Fundamentally, the system must simulate all the rules of Monopoly. Agents must be able to move around the board based on the dice roll, and if they so choose, be able to buy titles they land on, and buy houses on monopolies they own. Additionally, they should be able to sell houses and mortgage properties according to the rules of the game. When an agent lands on a Chance or Community Chest square, they should receive the top card from a deck which was randomly sorted before the game, and follow through on the instructions. Furthermore, to explore the research areas contained herein, agents should also be able to learn to some extent (e.g., negotiate with each other).

## BACKGROUND

According to Champanand, reinforcement learning "allows the machine or software agent to learn its behaviour based on feedback from the environment. This behaviour can be learnt once and for all, or keep on adapting as time goes by. If the problem is modelled with care, some Reinforcement Learning algorithms can converge to the global optimum; this is the ideal behaviour that maximises the reward." A good example of reinforcement learning in real life would be learning to ride a bike or how to walk. People are not given a complex set of rules to follow---they receive feedback (i.e., the bike tipping over or falling), and modify their actions to prevent the mistakes again. As described by Singh and depicted in his graph below, reinforcement learning is learning from interactions in which agents base their actions on perceptions from an environment and an associated reward system.

An expert system, on the other hand, is when an agent follows a complex set of rules written by humans. An example of this would be filling out forms or applications, as the applicant is not learning, merely following predetermined rules about how to fill out an application.

Player # 0					
Buy	Trade	Buy H	Sell H	Mort.	Unmort.
0	10	10	10	10	0
0	0	10	10	10	10
0	0	10	0	0	0
10	0	10	10	0	0
10	10	0	10	0	10
10	10	0	10	0	10
0	0	10	10	0	10
10	0	0	0	0	10

Aggressiveness Policy Vector of learning agent after 40,000 games

## CONCLUSIONS AND FURTHER RESEARCH

While my program was not able to find an optimal strategy for Monopoly or have in-depth negotiation, I believe the project was a success because the learning agent does indeed learn. Considering that I had to spend most of my time was creating a Monopoly simulation from scratch--including the complex trading and auctioning capabilities--the fact that an agent was able to learn (using what appears to be an original idea) is notable in itself. Moreover, the fact that the APV application of reinforcement learning does not take long to implement--and works--in a way fulfills a 'proof-of-concept', showing that the idea is a viable option for learning.

If I had more time to work on the project, I would have added or modified a large number of features. Based on the current results, I do not think that my APV approach contained enough flexibility to be able to correctly inform an agent how to play. I would implement a model that would allow for important factors such as how far the agent is in the game; how many monopolies he owns; how many monopolies his opponents own; relating the propensity to buy a property as a function of how much money the agent has available with respect to the board position, etc. I would also have attempted to find a better function for evaluating the merits of an agent's position and properties, such as estimating an intrinsic value for each property (as a function of how many properties are already owned). Another area of future research would be to build a complementary expert-based Monopoly agent system and pit it against the reinforced learning agent. Also, I would further develop the trading method so that agents can do more than attempt to buy properties from other players each turn. I would make it so that they also sell their own properties, and trade their properties for other agents' properties.