

Agent Based Simulation, Negotiation, and Strategy Optimization of Monopoly

TJHSST Computer Systems Lab 2007 - 2008

Nicholas Loffredo

ABSTRACT

Computers have a difficult time performing common human tasks, such as learning a language well enough to be able to "talk" intelligently with someone/something. Monopoly, one of the most well known and understood board games in the United States, if not the world, provides a good environment to see whether or not a computer can "learn" to negotiate through a number of strategies. It is much simpler than negotiating in the real world, due to the simplified environment, yet complex enough that it may be useful as an example of computer negotiation. By creating a Monopoly simulation with computer agents playing the game, it can be used as a test bed for these computer negotiations. The methods used in this test bed, if it works, could then be applied to more complex computer negotiation. The agents can be given aggressiveness values for different negotiation techniques, such as price "stubbornness" when selling or buying properties from other agents. The results from running the simulation hundreds of times can then be graphed to show which strategies are the optimal strategies for agents.

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. Most humans are able to negotiate with one another for various goods. Computers, on the other hand, can't. If computer were 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 hiring people to do it. These computers would be resistant to common human flaws, such as anger or impatience.

Making a computer that can negotiate effectively in a limited environment is a first step towards being able to negotiate in a more complex one. The game of Monopoly is simple enough that negotiation 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 negotiation. By making a working simulation of Monopoly, a negotiation 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, 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. 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 'do' whatever the card says. Furthermore, to explore the research areas contained herein, agents should also be able to negotiate with players. In particular, these negotiations will be based on aggressiveness levels. For example, how far an agent is willing to drop/raise his initial price in order to complete a negotiation.

BACKGROUND

Surprisingly, not much research has been done into making agents for Monopoly that can 'learn' the optimal strategy for negotiation. One of the few existing simulations is one that determines the probability of landing on each square in Monopoly. This can be used to see if my Monopoly simulation results correlate to theirs, and determine whether or not the simulation works correctly. Fortunately, there has been research done into reinforcement learning, which is effectively how the agents will learn. Reinforcement learning is when agents take a number of actions over a course of a game, and then are basically 'told' that they did well (when they won) or they did poorly (when they lost). Based on this feedback, each agent will try to change its aggressiveness values (which may involve different strategies) to find the winning values. Also, agents may learn from the strategies other agents used, and whether they won or lost, to determine how their aggressiveness levels should change, which varies from the traditional approach slightly. There is not a state-of-the-art reinforcement learning program yet, however. Everything I Need to Know About Business I Learned from Monopoly, which discusses various strategies for Monopoly, can be used to see if agents develop the strategies the book discusses.

PROCEDURES

PRELIMINARY TESTING and ANALYSIS

To test my program, an interface was designed to help me spot errors in the game structure. The game works as it should, with agents moving around the board, buying properties, paying income or luxury tax, paying rent on other player's properties, and correctly following the instructions on Chance or Community Chest cards. These were tested by using the interface to watch each step of agents to ensure that they moved the correct number of squares, landed on the correct properties, paid rent when necessary, etc. Agents correctly performed their required actions.

I also tested my program by running it a large number of times and seeing if the learning agent actually has a higher winning rate than the random agent.

EXPECTED RESULTS

I expect to be able to find an optimal strategy for Monopoly based on agent victory rates vs. other agents (and their strategies). I also expect to find that reinforcement learning can be applied to Monopoly effectively.



Chain Go: false

Player 0 Pts: 294	Ratio: 58%
Player 1 Pts: 216	Ratio: 42%