CS 365: Introduction to Scientific Modeling
Rock-Paper-Scissors Dynamics
Assignment 3
November 4, 2014

1 Introduction

The idea of this assignment is to use the genetic algorithm (GA) to evolve strategies for playing a simple two-person game known as Rock-Paper-Scissors (RPS). In RPS, each player, or *agent*, has a choice of three alternatives: Rock (R), Paper (P), and Scissors (S). After each player has made its choice, the winner is determined as follows: Rock beats Scissors (by smashing the scissors); Paper beats Rock (by wrapping the Rock); Scissors beats Paper (by cutting the paper). For one round of the game, all choices have equal probability of winning, so no strategy is favored. However, the game becomes more interesting when there are repeated interactions between the players. In this assignment we will be as interested in the dynamics of a population of RPS players as we are in evolving a single rule that plays the game well.

There are two parts to the assignment. In Part I, you will use the GA to evolve strategies for playing RPS, and in Part II you will study the dynamics of RPS on a spatial grid using a more simplistic learning rule (imitate your most successful neighbor).

2 Assignment

You are free to choose any publicly available GA software or to write your own using the language of your choice. One possible choice is a very simple GA written in Matlab by George Bezerra, or the more sophisticated GA package that is part of the Matlab Optimization toolbox.

1. **Part I Due: Wed. Nov. 12:** Use the GA to learn strategies for RPS using the representation described in Section 2.1. Fitness will be assessed by

   (a) Test the GA against a simple Nash strategy (the strategy that chooses randomly among the three alternatives with equal probability). That is, evaluate each individual's fitness by measuring its success playing the repeated RPS against the Nash strategy (you decide how many iterations and justify your choice). Report how the mean, max, min fitness of the population changes over time (one graph), and then make a second plot showing for each time step, the frequencies of Rock, Paper, and Scissors plays (3 time plots on one graph). Finally, select your best-fitness individual and analyze its strategy, that is, how is the strategy succeeding. Is it better than Nash? If so, why?

   (b) Now, change the fitness function so that an individual i plays against x other individuals on each generation to determine i's fitness. The best value for x will depend on your population size—you will need to experiment with different values of x to learn which values produce the most interesting dynamics, and report those. Each i vs. x interaction consists of a repeated RLP game, say for 51 rounds. Redo the plots and analysis for this co-evolutionary case.

2. **Part II Due: Wed. Nov. 19:** Study a spatial version of RPS.

(a) Add space to the model, so that the population occupies a 2-dimensional array, with the state of each cell representing a single agent. It will be convenient to represent each different state using a different color. Here, the representation of a strategy will be much simpler than in the previous section—each square's strategy is to always play one of: R, S, P. Each square in the array will play a single iteration with each neighbor including itself. Assign each Win +1 points, each Loss -1 points, and each Tie 0 points. Then, add up the scores from all 9 games to compute fitness for each square. (note: you will need to do some bookkeeping to avoid replaying the same game multiple times.)

Instead of running the GA, each strategy will simply adopt (imitate) the strategy of its most successful neighbor, including itself. We recommend using the Moore neighborhood with wraparound, but if you have time, you can experiment with von Neumann neighborhoods, or with neighborhoods that have a larger radius.

Once you have the basic game implemented, begin studying the effects of different initial conditions. Two variations that you will want to start with are:

- Random: Each cell is set randomly to one of the three possible states.
- Segregated: Assign one sector of the grid to R, a second sector to S, and the third sector to P, so that there are three solid blocks, corresponding to the three states.

Try out some runs with these different initial conditions, compare their behavior, and report your results concisely using a combination of plots, representative runs, and narrative text.

2.1 Representation for Part I

You will be modifying the representation described in class for the Prisoner's Dilemma game. In that representation, a history of the past three moves is maintained and used as an index into an array that specifies what move to make. If you choose to maintain a history of two moves, then your strategy will have 3^4 possibilities rather than 2^6 (because there are three possible choices at each time step and you are remembering only two time steps).

2.2 What to hand in

On Nov, 12 you are expected to hand in a working GA (source code listing), some plots of its performance (e.g., mean fitness plotted against generation), descriptions of examples of high-fitness individuals, and a 3-5 page writeup. On Nov, 19, you should hand in an augmented writeup that describes your new representation, how the spatial algorithm works, details of your fitness function, your new results, a comparison between the various methods, and your conclusions.

The writeup should describe at least the following, but concisely. You do not get extra points for verbosity:

- The problem you are trying to solve (in your own words);
- How you set out to solve the problem, e.g., which genetic algorithm software you used, its basic algorithms (what kind of selection, crossover, mutation, etc.);
- The representation (how did you represent strategies to the GA);
- The fitness function(s) you tried;
- The basic parameters for your runs (population size, generation time, crossover rate, mutation rate, etc);
- Inputs and outputs;
- Experimental design (what experiments did you run and why);
- Experimental results, including results about individual performance (e.g., the best strategy in the population) and about population dynamics;
- Discussion and Conclusions (how well did your program work and why do you think it performed the way it did)?