Investigation of an Adaptive Cribbage Player

Corey Hebert

CS 152 - Neural Nets

Fall 2007

Graham Kendall and Stephen Shaw

- Introduction
- Evolutionary Strategies
  - Genetic Algorithms
- Experiments
- Discussion

Cribbage

- Game has several phases
  - Each player is dealt 6 cards
  - Each player discards 2 of those cards
  - Players alternate playing cards
  - Players score their hands
- Winner is first to 121 points

Objective

- Create a cribbage player that learns rather than is told explicitly how to play
- Focus on choosing which cards to discard to maximize score when you count the hand
- Investigate an Evolutionary Strategy
  - Feasibility
  - Specific parameters

Genetic Algorithms (GA)

- Method to find minima/maxima of functions
  - Initial set of points is selected
  - From the information at these points, a new population is generated
  - Best points are retained
  - Method is iterated
- Useful when little information is known about the function, or if it is not easily differentiable

Simulated Annealing vs GA

- Similarities
  - Many points/iterations
  - Can be computationally expensive (parallelization)
  - Chance to select "unfavorable" points
  - Temperature decrease over time
- Differences
  - GA incorporates mutation with each generation
  - Recombination

Deceptive Problems

- Some functions are inherently more difficult to solve through GA
- Simple gradient descent may be more efficient

- Many optimizations utilize combinations

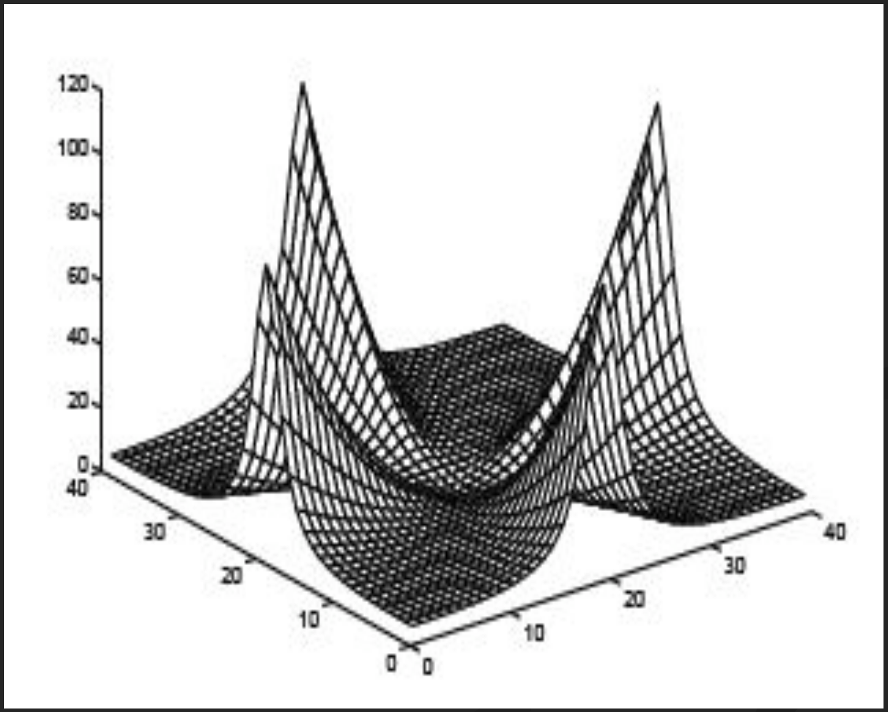

Next Generations

- Two primary variations for forming the next generation
  - $(\mu + \lambda)$
    - μ parents and λ offspring all compete for survival
  - (μ, λ)
    - λ offspring compete for survival ($\lambda > \mu$)

- Authors utilize a 1+1 strategy

Evolutionary Strategies (ES)

- Each individual is represented by a pair of real vectors, $v = (x, s)$
  - x is a point in space
  - s is a vector of standard deviations
- Mutation:

$$x^{t+1} = x^{t} + N(0, \sigma)$$

  - $N(0, \sigma)$ is a random Gaussian number with a mean of 0 and a standard deviation of s

Procedure

- Assign a value to each of the 6 cards
- Lowest two values are discarded
- Score the hand
- If the score is less than some percent, *p*, of the maximum possible score the values of the two discarded cards are updated
- Iterate, ideally until convergence

Feasibility

- 9 sample hands used
- Values constrained between 0.1 and 0.9

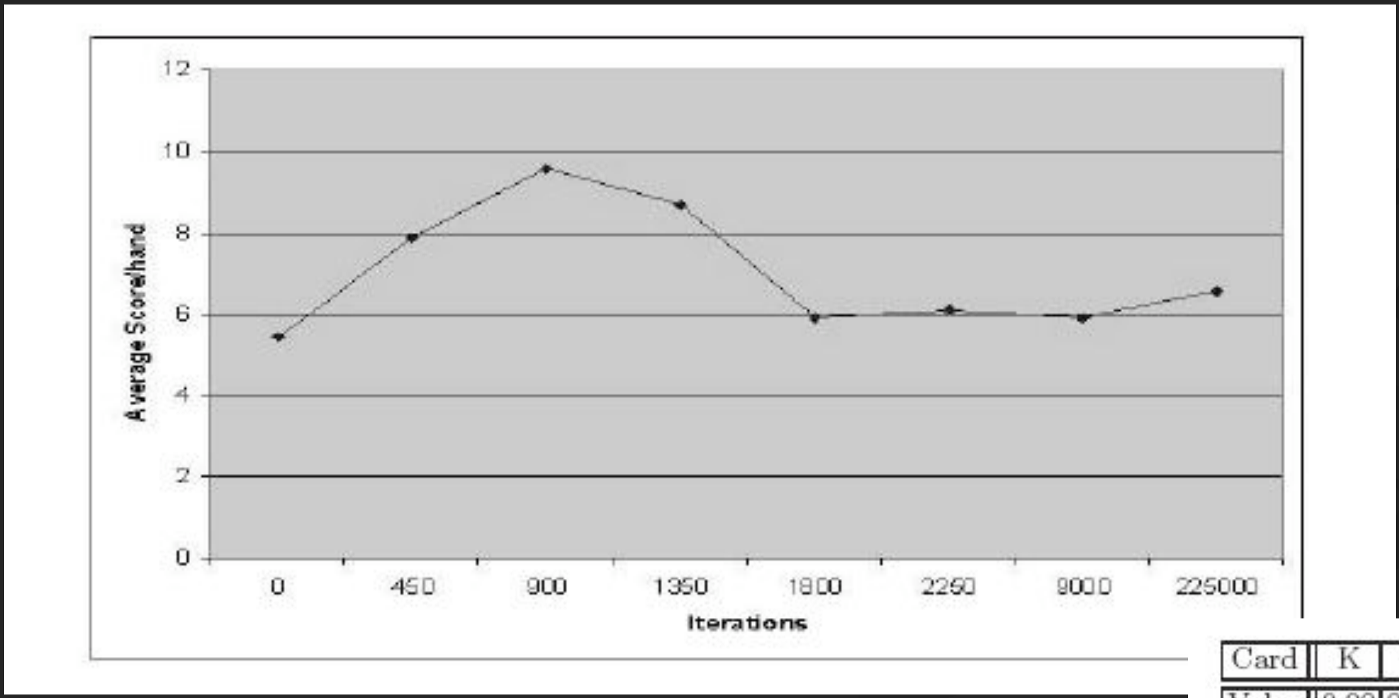

Card	K	Q	Q	10	5	4
Value	0.90	0.90	0.88	0.79	0.80	0.11

Refining *p*

- 0.65 found to be most optimal for sample hands overall
  - Tailing off at high and low values is not surprising
- Observed that for certain hands, different *p* values were better

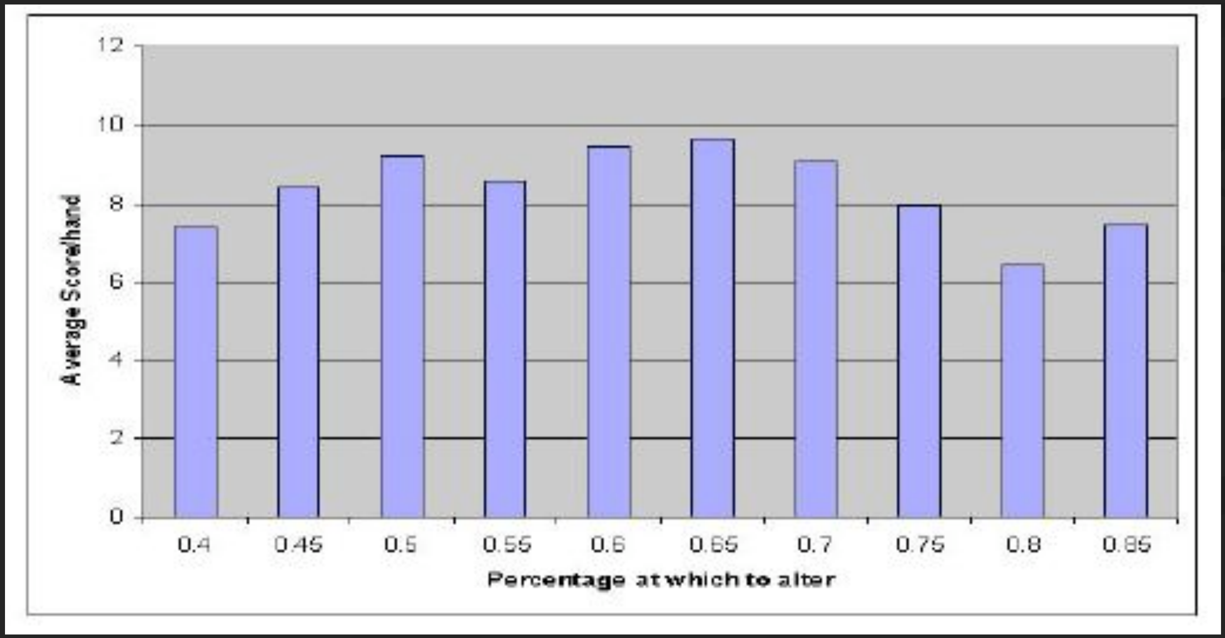

Co-Evolution

- Two players simulated
  - Each played one of the sample hands
- Scores were compared against each other
- The loser updated his statistical array

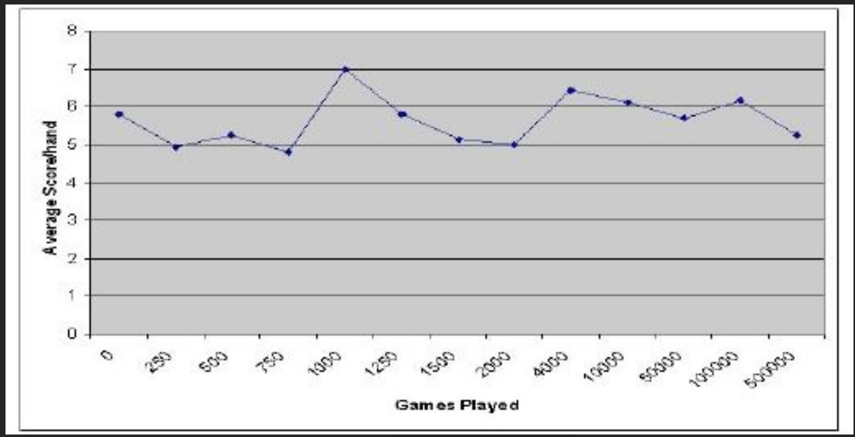

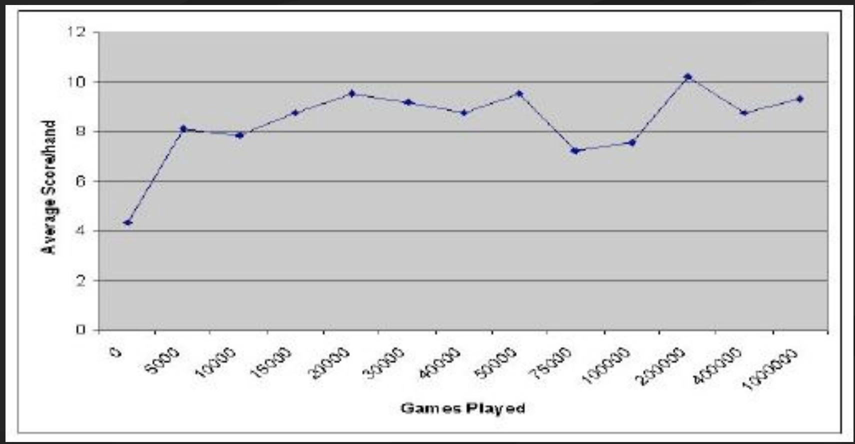

Crib

- The 2 discard cards from each player form a 4 card “hand”
  - The points for this “hand” alternate between players
- Discard strategy should consider this

Crib	10	9	8	7	6	5
Player	3.55	3.58	3.58	-1.83	-3.50	5.56
Opponent	0.73	2.54	2.60	2.50	2.52	-3.50

Final Experiment

- Tested adaptive cribbage player against a commercial product
  - Beat Easy 5-0
  - Beat Medium 3-2
    - Better discards offset by poor pegging algorithm
  - Lost to Hard 5-0
    - Comparable discards but lost due to pegging
  - Beat Hard 3-2 (without pegging)
  - Lost to Harder 5-0

Conclusion

- Evolutionary Strategies work
  - Adaptive player can learn to play competitively
  - Explicit strategy need not be pre-programmed
- Techniques show promise for extension
  - Pegging phase of the game
  - Other games with incomplete information