Piece Difference: Simple to Evolve?

Evan J. Hughes
Department of Aerospace, Power and Sensors,
Cranfield University,
Royal Military College of Science,
Shrivenham, Swindon, England. SN6 8LA
ejhughes@iee.org

Abstract- This paper details a study into whether the 'Piece Difference' heuristic could be evolved for the game of checkers. A co-evolutionary algorithm is used to evolve a piece-weighting system that is used as the evaluation function in a minimax checkers player. The results suggest that the 'Piece Difference' heuristic will evolve if allowed, but it is not necessarily easy. The work has also demonstrated that other common heuristics can also be evolved.

1 Introduction

The 'Piece Difference' heuristic is often found as one of the key heuristics of many game playing strategies used in the board evaluation function of many games. The piece difference heuristic will, in essence, try to take all of the opponent's pieces, before the opponent can take yours. In checkers, the piece difference strategy is often seen as an aggressive player, making moves that maximise the number of pieces taken in the smallest number of ply. The piece difference has been applied as a base heuristic in some works (Chellapilla and Fogel, 2000), with the implicit assumption that it could be evolved easily.

This paper focuses on attempting to evolve a piece difference strategy for both the men and kings independently in the game of checkers. The results demonstrate that the reverse-symmetry (white on square 32 has negative of weight on square 1 etc.) required in the white pieces with respect to the black pieces develops, although some regions of the board develop much faster than others. Even when reverse-symmetry is enforced, the weight of many of the king locations evolve only very slowly.

2 Method

2.1 Representation

The board is arranged with the top left hand black piece being square 1. The square is the first playable square and does not lie in the corner of the board, as shown in figure 1.

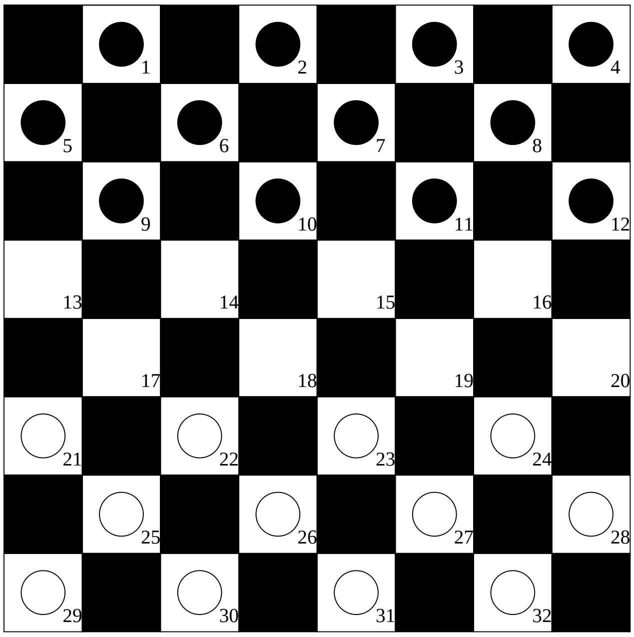

Figure 1: Layout and indexing of board

A quartet of weights is used for each board position, each quartet corresponding to black man, white man, black king and white king. Thus a total of 128 weights must be considered to represent the board evaluation. The evaluation function is designed to return +1 if black is winning and -1 if white has the upper hand. If a man crosses the board to the far row (eg. black to squares 29 to 32), it is promoted to a king immediately. Therefore, for black and white men, only 7 rows of four weights are required, leading to a total of 120 weights that must be evolved in practice. Thus a chromosome with 120 real values was used to represent the set of weights.

The evaluation function starts by generating a random value in the range [-0.05,0.05], then adds the weight values corresponding to the piece types and locations that are currently being played. The total sum is then passed through a hyperbolic tangent function, $\tanh()$, to restrict the output of the evaluation to 1 if black is considered to be winning, and -1 if white is performing better. The use of a small amount of randomness allows each game played against the same player to be slightly different. Without a small amount of randomness, the evolution very quickly stagnates into players that draw against each other, all playing near identical games. The range of randomness chosen gave consistent evolution, without destroying the behaviour of the evaluation function. Allowing the amount of randomness to evolve always resulted in the randomness being reduced to zero and stagnation of the evolutionary process.

2.2 Player Algorithm

The evaluation function was used in a minimax strategy with α–β pruning, and played to a fixed depth of 6 ply, with no extensions for forced moves. Each player played 5 games against other players selected from the population at random, and without replacement. As used by Chellapilla and Fogel (2000), for each game won, a player scored +1. A draw scored zero and a loss -2. A draw was called after 100 moves each had been played. The aggregate score over the 5 games were used as the objective value to be maximised.

2.3 Parallel Evolutionary Strategy

A parallel evolutionary strategy was used with a population of 100 chromosomes. The algorithm was a simple 'farming' model with the next member of the population (and therefore 5 games) being allocated to the next free process from a pool of 6 DEC ALPHA 667MHz EV67 processors. In every game, the choice of whether to play black or white was made at random to prevent any bias.

Each of the 120 weights in the chromosome was initialised to be a uniformly distributed random number in the range [-5,5].

The simple evolutionary strategy (Deb, 2001), in conjunction with intermediate crossover, used a probability that two chromosomes would cross of 0.7 and initial values for the standard deviation of the mutations of 2 for each gene.

3 Results

The algorithm has been run a number of times, but each time the results are very similar. Many runs would be required to generate statistics of the behaviour of the evolutionary process, but time has not permitted. At an average of 3 games per second per processor, the evolutionary strategy is capable of performing about 3000 generations per day.

3.1 Evolution to distinguish black from white

Figure 2 shows a plot of the weight values for the men and the kings after 100 generations. The solid line is the weight values for the black pieces, given the board positions as shown in figure 1. The last 4 board positions for the black men have no associated weight as the piece is promoted to a king. The dashed line shows the weight values for the white pieces.

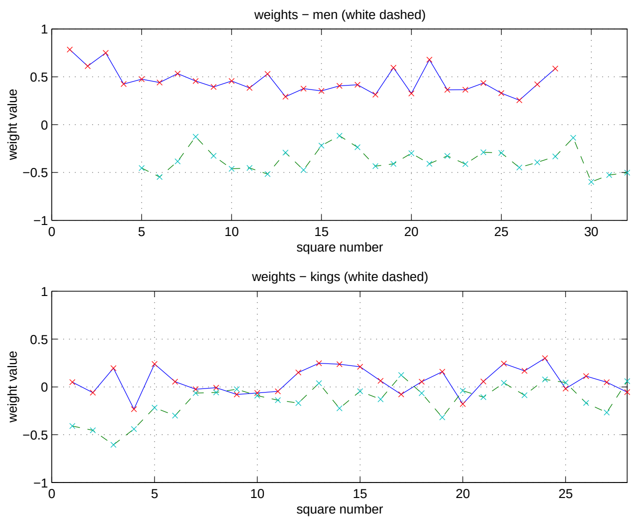

Figure 2: Weight distribution after 100 generations

It is clear from the figure that the black men have a positive value, and the white men have a negative value. It is also apparent that square 31 is almost the negated value of square 2: squares 1 and 32 corresponding to the first member of the first row of black and white respectively.

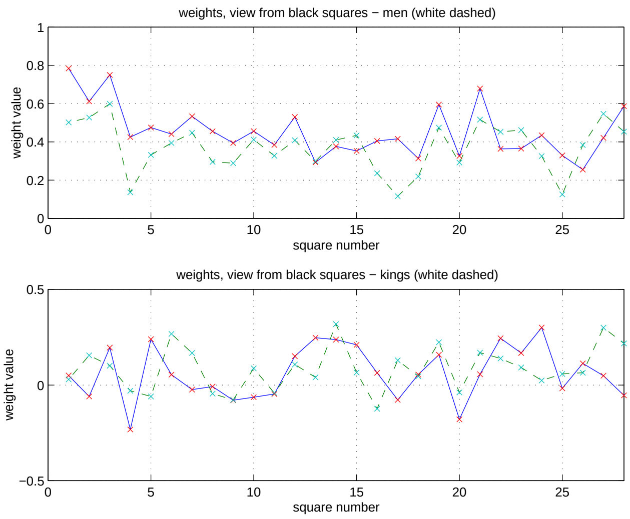

Figure 3: Symmetrical Weight distribution after 100 generations

Figure 3 shows the same data as figure 2, but the dashed line shows the negated and index-reversed weights for the white pieces, i.e. the weight shown at index 1 is the negated weight value for board position 32 etc. Thus index 1 to 4 represent the 'back rank' of both black and white etc. and the graph is referenced to the black player only.

It is very clear that the magnitude of both black and white at each square are evolving to create an 'anti-symmetry' in the board weights, with an eventual mean of zero for rotationally symmetrical piece distributions – therefore piece difference. It is clear that the central area of the board has evolved early in the optimisation, the conjecture being that a strong opening play is one key to a successful win. It is also noticeable that weights 1 to 3 are considerably higher than the other weights, indicating that the back-rank is important. The weights for the kings are inconclusive, with some still even being portrayed as negative and favouring the opponent.

Figure 4 shows a plot of the weight values for the men and the kings after 7388 generations. Figure 5 shows the difference between the lines of figure 4 (black + index reversed white in practice). It is clear that the weights for the back ranks and the centre of the board for the men are very similar and can therefore be assumed to be approaching their final values. Indeed in each run of the algorithm, very similar final values have been noted. Conversely the weights near the opponents back rank have a small bias towards white. It is also very clear that there is still a wide deviation in the weights of kings, even after over 7000 generations. As the kings are seldom in play, and also in small numbers, the selective pressure to drive the weight values is very small, and so the weights evolve very slowly.

Often kings are weighted as K times the value of a man.

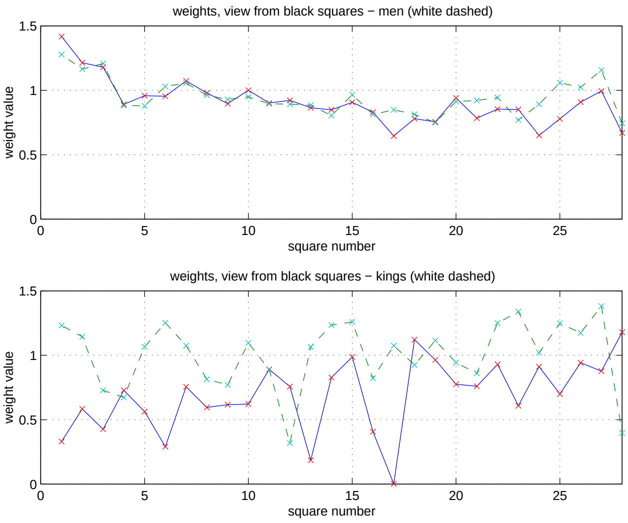

Figure 4: Symmetrical Weight distribution after 7388 generations

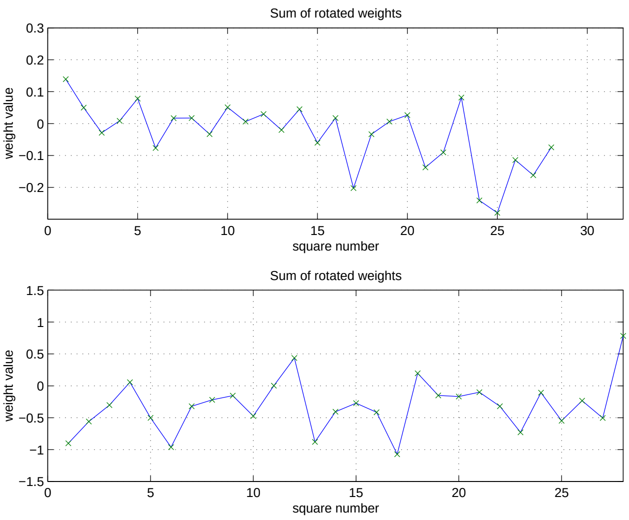

Figure 5: Difference in symmetrical Weight distribution after 7388 generations

Values of K in the range of 1.2 to 3 are common. Figure 6 shows the effective ratio of the king value with respect to the value of the man at the equivalent square index. The values range from near zero to almost 1.8, although as the weight values of the kings have not fully evolved, there is likely to be significant error. It is noticeable however that there is a general trend to prefer kings far away from the players back-rank. This feature has also been seen in all repeated runs. One conjecture is that in the short time that a king is in play, it will generally be near the area where it was created, thus biasing the evolution to concentrate on those areas.

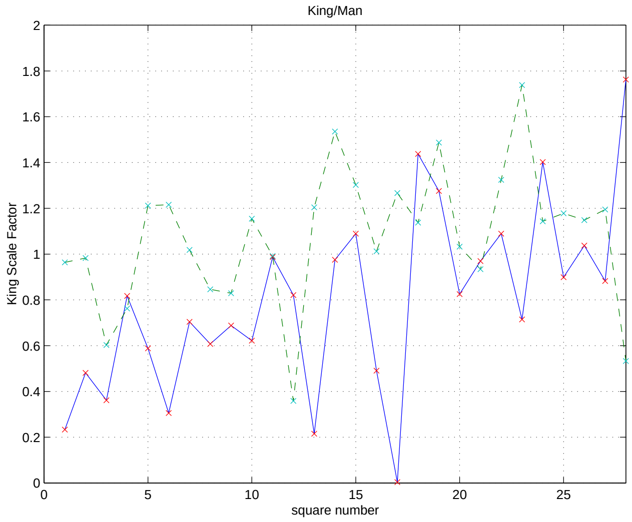

Figure 6: Ratio of King value with respect to Men

3.2 Evolution to generate weight values

Once it was determined that the evolutionary process would evolve the white weights to be rotationally assymetric to the weights for the black pieces, the rotational assymetry was enforced and the optimisation process repeated. With the anti-symmetry enforced, only 60 weights must be evolved, 28 for the men and 32 for the kings. The weights were initialised to lie within the region [-3,3]

Figure 7 shows a plot of the weight values for the men and the kings after 16594 generations. All the weight values are positive, indicating that every piece is of use to the player, no matter where it is located. It is also clear that the back-rank of the men is generally favoured more highly than the other regions, with the centre of the board being weighted the least. In the early generations, the centre of the board was often weighted relatively more highly, but as more sophisticated play evolved, the board centre became less ‘interesting’. Figure 8 shows the ratio of the king values to men. The mean of the distribution of the king value with respect to the mean value of a man is 1.23, although the kings are considered relatively more important near the opponents back rank.

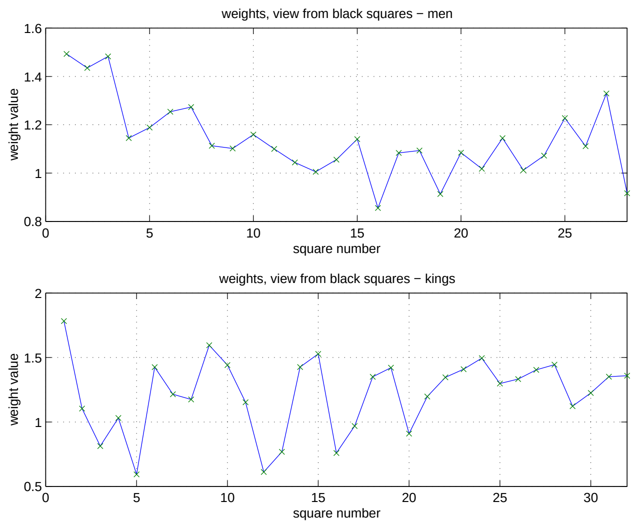

Figure 7: Symmetrical Weight distribution after 16594 generations

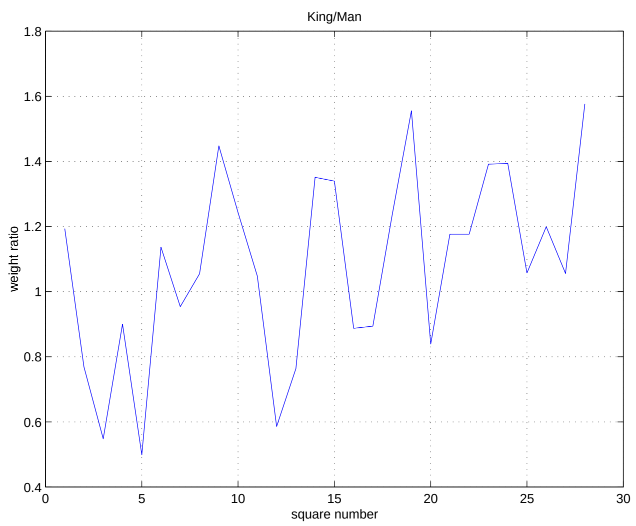

Figure 8: Ratio of King value with respect to Men

4 Conclusions

It has been shown that co-evolution can indeed discover piece difference as key heuristic in checkers. It is also clear that while key regions of the board, such as the players back-rank and the centre, evolve quite quickly, other areas, especially the king values, evolve very slowly. The weight values for the pieces across the players back-rank are favoured over the centre of the board, indicating a preference to maintain the players own back-rank, while endeavouring to disrupt the back-rank of the opponent. This is a common strategy in hand-coded evaluation functions as an intact back-rank prevents the opponent creating kings.

Thus two key heuristics have been discovered by the evolutionary process, without any additional insight, demonstrating the power of co-evolutionary methods for knowledge discovery in highly non-linear and stochastic processes.

References

Chellapilla, Kumar and David B. Fogel (2000). Anaconda defeats hoyle 6-0: A case study competing an evolved checkers program against commercially available software. In: *Proceedings of the 2000 Congress on Evolutionary Computation*. IEEE Press, Piscataway, NJ. pp. 857–863.

Deb, Kalyanmoy (2001). *Multi-objective optimization using evolutionary algorithms*. John Wiley & Sons.