

FOOTSTEPS


- A psychological game .
- Pencil and paper
- Five minutes
- A little game of surprising subtlety

Background

This simple-to-play psychological game combines mind-reading with bluff, mental arithmetic and just a little luck. A proprietary version is marketed as Quo Vadis (Invicta).

Play

The 'board' consists of seven circles or marks in a straight line, spaced so that each circle is large enough to accommodate a small token. A line is drawn through the central circle. The three circles on each side of it belong to the respective players (see Figure 39).



Each player starts with 50 points. On each turn, both players write down the number of points that they wish to spend. The minimum allowed is 1. Numbers are then compared and the player with the larger number advances the token, which is placed at the start of each game in the central circle, one space into his opponent's territory. If the numbers are equal, the token is not moved. Both players now deduct the points they have used from their allocation of 50 to give their remaining total of points, and the next turn starts. The object of the game is to get the token on to your opponent's end circle; or, put another way, to win three more turns than your opponent. If one player uses up all his points, the other can have successive plays. If both use up their allocations, the game can either be agreed drawn or a half-victory awarded to the player in whose territory the token does not lie. And that is all there is to the game - but don't underestimate it.

Strategy

The subtleties of Footsteps will not become apparent until you have played a few games. The ideal to be aimed at, of course, is to win each turn by a single point. To allow your opponent a big lead on points is to face certain loss. In the early stages a margin of 5 or 6 points is acceptable, but later this could be severe. An extreme case is where you have got the token to within one space of victory but have used up all your points. Your opponent, however, has 5 points left. Clearly he has only to play 1 point each turn to win.

Be frugal with your points until danger threatens. Consider the situation where the token is on the centre and you have each used 30 points. You decide to allocate 9 to your next turn (you wouldn't do this if you'd played the game twice!) and you find that your opponent has only allocated the minimum permissible of 1. Now you are two spaces from victory but your opponent has 19 points left to your 11. If you now play 6, for example, your opponent may again play 1, and although this puts you within a space of home, you have only 5 points left to your opponent's 18 - an unacceptable margin. Your opponent next plays 5 (in case you gambled all your remaining points) and thereafter will force a win by a series of Jow plays: remember that you are obliged to use at least 1 point on each turn. It will now be dear where the mind-reading comes in. Some people seem to have second sight at this game, particularly when playing with close friends or relatives.

Psychic players could claim that it is an advantage to let the opponent write his number down first. The way round this is to keep the action concealed with, say, a one-minute time limit per move.

Good advice for Poker is good advice for Footsteps: never play routinely. If you always play low your opponent can advance quickly by pitching a few points higher each time. If he reaches the penultimate space you are obliged to equal his point count, so it is better to stop him one space further back. But he may read your thinking on this and - well, as I said, the subtleties of the game will soon become apparent.

Variants

The seven-space board and the fifty-point allocation give a quick and balanced game, but you may prefer other parameters.

Incidentally, a couple of pocket calculators are a useful aid for Footsteps. Start with 50 on the display and at each turn compare negative inputs.

Evolving Strategies for the Game Footsteps

Robert Morris

Faculty of Computing Sciences
and Engineering
De Montfort University
Leicester, LE1 9BH
UK
robert.morris@email.dmu.ac.uk

Tim Watson

Faculty of Computing Sciences
and Engineering
De Montfort University
Leicester, LE1 9BH
UK
tw@dmu.ac.uk

Abstract

Footsteps is a two-player board game in which the objective is to move a counter, one step at a time, to one's own end of the board, by spending points. Though the game is simple in conception, it is unusual and strategically quite complex, and it captures the essence of the important tactical dilemma of whether to attack or to sit back and then counter attack. In this paper, deterministic and probabilistic strategies are evolved for the game of Footsteps. The noteworthy aspects of the genetic algorithm used here are discussed, and it is concluded that unpredictability is the key to success.

1 Introduction

For decades, game theory [8] has been of great use in many different academic disciplines, particularly biology and economics. Some examples of topics that have been enlightened by game theoretic analysis are: the Battle of Bismark Sea in the Second World War [14], aggressive behaviour of animals [17], and whether or not to invest venture capital [18]. Perhaps the best known, and certainly the most studied [5], is the *Prisoner's Dilemma* (PD), often in an iterated form (IPD). The value of this game lies primarily in the fact that it captures, in as simple a way as possible, the essence of the dilemma of whether to co-operate or defect.

This paper reports an experimental investigation into the board game 'Footsteps' [1, 2]¹, a two-player game of complete information. The rules of the game are as follows.

The board (figure 1) has seven sections, separated by vertical lines; the left end is player one's end, and the right is player two's. A counter is placed in the middle section, and the two players are allocated 50 points each. In each turn, the players simultaneously and independently sub-

(P1) ...|...|...|○|...|...|... (P2)

Figure 1: The layout of the board for Footsteps

mit some of their points (between zero and their current total): if one player submits more points than the other, he wins the round, and moves the counter one step towards himself; if both players submit the same number of points, the round is tied, and the counter does not move. The object of the game is to win the counter. If the counter reaches one of the outer sections, the player whose end it is wins the game, and his opponent loses; if, however, the players both run out of points before the counter has reached an extremity, the game is tied.

The size and simplicity of this game are in stark contrast with the seemingly intractable complexity of its game-play. In the absence of knowledge of one's opponent's intentions, the question of how many points to play next is very difficult to answer convincingly. Crudely speaking, in any given turn, a player must decide whether to submit a 'large' number of points (to attract the counter), zero points (to absorb a potentially large number from his opponent), or a 'small' number, as a gamble. And when the counter is one or two steps away from someone's end, the pressure on these decisions rises.

Whereas the Prisoner's Dilemma models co-operation, Footsteps models attack-versus-counterattack. This topic has not yet been addressed directly by game theorists, with the closest work being that pertaining to predatory pricing and the chain-store paradox [16]. There are many manifestations of attack-or-counterattack in the real world, from sporting and military strategies to political election campaigning, so a game like Footsteps merits investigation as strongly as any other abstract game.

In section 2, the genetic algorithm that was

1. Also known as 'Citadel' and 'Quo Vadis?'

used to evolve the strategies is presented. In section 3, the results of the experiments are discussed, and section 4 is the conclusion.

2 The Representation and Evaluation of the Strategies

There have been numerous instances of the computational evolution of strategies for games [12, 6, 15], e.g. the IPD [11, 9], Chequers [7], Rock-Paper-Scissors [4], and Awari [10]. The strategies themselves have varied from look-up charts to multivariate functions to tree searches, and many of them have proved very successful, against both human and machine opponent.

2.1 The Representations

In Footsteps there is a space of every possible state the game can occupy. Each state consists in three variables: the number of points player one has (0–50), the number of points player two has (0–50), and the position of the counter (1–7). There are thus $51 \times 51 \times 7 = 18,207$ points in the state space, the majority of which are viable game states (some points cannot be game states, e.g. (50,50,3) because the counter cannot have reached board-section 3 without player one having spent a single point). A game of Footsteps can be viewed as a journey through this state space, starting always in the initial state (50,50,4), and whose length and progression are determined by the decisions of the players.

The state space for regular Footsteps is very large, so for this investigation, a scaled-down version was used, where the players started with 20 points and there were only 5 sections on the board. This smaller version was computationally less taxing, easier to analyse, and receptive to more selective pressure than full Footsteps would have been, whilst being big enough to preserve the key aspects of game-play (the important first decision, and the challenge of being one turn away from losing).

2.1.1 The Deterministic Strategies

The representations of the strategies are based on this conception. For the deterministic kind, every gene on the chromosome corresponds to a unique game state, with the value of the gene being the instruction of how many points to play there. The lowest number of points allowed to be played was raised from zero to one, to prevent the possibility of games running indefinitely, so the chromosome length was $20 \times 20 \times 3 = 1200$

genes. (The 3 signifies the middle 3 sections – the outer 2, when reached, represent the end of a game, so they do not require next-turn instructions.) At a given locus, the gene value must lie between one and the number of points the player currently has; for example, the gene for state (15,12,3) must be an integer in [1..15].

When every gene in a given individual's chromosome has been set to a viable value, it is ready to play against any opponent. For every scenario that can arise, there is a legal instruction of what to do next.

2.1.2 The Probabilistic Strategies

The probabilistic strategies are an extension of the deterministic ones. Whereas those have *one* gene for state (p, q, c) , its value in [1.. c], these have c genes, where each gene's value – in [1..99] – represents the relative probability of that number of points being submitted. This raises the chromosome length to 13,167 genes², an order of magnitude increase. As an example of how they work, if the genes corresponding to state (4,9,3) were {30,60,70,40}, then the individual – when in that state in a game – would have a $30/(30+60+70+40) = 3/20$ likelihood of playing one point, a $6/20$ likelihood of playing two points, a $7/20$ likelihood of three points, and a $4/20$ likelihood of four.

It is again the case that when every gene in a given individual's chromosome has been set to a viable value, it is ready to play against any opponent. For every scenario that can arise, there is a selection of weighted legal next-moves, from which a probabilistic choice can be made.

Both strategy types were evolved separately in distinct experiments. They were not evolved together for this paper because trial runs failed to produce notable new results, at the cost of complicating and extending the parameter set.

2.2 The Fitness Function

Populations of 100 individuals were used, and in every generation there was a round-robin league. For the deterministic strategies, each fixture comprised one game, but for the probabilistic strategies, because very many games could arise between any two players, each fixture comprised a sequence of 10 games. In games where there was a winner, the winner was awarded 3 points

2. Some of these genes relate to game states that cannot occur. To have weeded out every such state would have been to complicate the chromosome without gaining anything, so they were left in.

and the loser none; in games that were drawn, both players were awarded 1 point. At the end of the generational leagues, the point totals became the fitnesses.

3 The Results

The parameters were essentially the same in both sets of experiments. Standard proportionate selection was used, there was uniform crossover across 60% (on average) of each newly-selected population, and the mutation rate was set to a value to cause every three individuals to experience two mutations between them (so if the chromosome was N genes long, the gene-wise mutation probability was $\frac{2}{3N}$). The question of mutation will be discussed later in more detail.

3.1 The Deterministic Results

Figures 2 and 3 (at the rear) contain the key results pertaining to the deterministic strategies. Figures 2(a) and 2(b) are snapshots of the extent of the convergence of a population. In particular, they show the standard deviations of the gene values at the loci corresponding to the game states where the counter is in the middle section (so they represent exactly $\frac{1}{3}$ of the chromosome). For example, the bottom-right slots correspond to the initial state, (20,20,3). Square brackets indicate that that state was visited – in the course of a game – on at least one occasion during the generational league.

In figure 2(a), taken after 20 generations, the local standard deviations vary greatly, showing that the population is very genetically diverse. The large quantity of square brackets shows that many different game states were visited in that generation. This was because different individuals had different instructions for what to do in the states that arose.

In 2(b), taken after 200 generations, around $\frac{2}{3}$ of the local standard deviations are zeroes, suggesting that the population is *close* to convergence. Only four game states were visited in this generation, and in those loci, the population is converged. This communicates the fact that though the population was unconverged *genotypically*, it had converged *phenotypically*. This is because the phenotypes here are the games themselves; every game in that generation was played out in the exact same way³, because in the loci corresponding to its states, the individuals all had the exact same genes.

Figure 3 gives further information about how

evolution progressed in this GA. 3(a) shows how the total number of game-states visited varied between the 50th and 250th generations of a run. The plateaux indicate situations where one – or a few – phenotypes are prevalent; where only one – or a few – games are played out in the league. The rises that terminate the plateaux show those occasions where a mutation occurred which changed the value of a gene that was contributing to a phenotype. Sometimes, the mutant was a loser and so died off, and the old order was restored, but when it was not a loser, copies of it persisted in the population.

This is the key result of the foregoing experiments. In strategy GA's like this, the environment is rooted in the population, making it dynamic ([9], §1.1), so phenotypic convergence is synonymous with environmental stationarity. Mutations (and random immigrations [13], and any other introductions of diversity) therefore change the environment, in a way that can be trivial or profound. In other words, an individual's fitness is wholly dependent on which other individuals are in the population alongside it.

In the setting of Footsteps, this means that a given deterministic strategy is 'good' or 'bad' according to its opponent(s). Any given sensible⁴ strategy, whilst being able to beat many opponents, can also be defeated by many opponents. Because of this, rather than the populations evolving in the traditional manner ('upwards, towards a good solution') they enter a state that is analogous to *neutral equilibrium* [3]. They remain in a phenotypically converged state until some mutation displaces them into a new state, and so on.

3.2 The Probabilistic Results

Figures 4–6 (at the rear) contain some selected results pertaining to the probabilistic strategies. Figure 4 shows the mean values (for the whole population) of the genes that determine the number of points to play at the start of a game. There are 20 genes at that part of the chromosome, each one representing the relative probability (from 1 to 99) of choosing that particular number of points. The probability distributions are shown from 500, 1000, and 1500 generation milestones, and the similarities (by seed) reveal that the populations were nearly

3. The exact progression was: (20,20,3), (11,11,3), (5,5,3), (2,2,3), (1,1,3), (0,0,3) – game tied.

4. "Sensible" here means: not needlessly spending so many points in a given round that one cannot realistically go on to win the game.

converged there (the standard deviations – not shown in the figures – were mostly zeroes, corroborating this observation).

There are two important things to notice in these plots. Firstly, that the rightmost genes acquired small values (in the non-dotted curves, you can see this physically happening). In the game, this translates as: “Do not play 16–20 points on your first turn.” This was to be expected, because such moves make it practically impossible for the player to go on to win.

Secondly, across the range of possibilities that do *not* put the player in big trouble, there is no single dominant value. In fact, there is a *wide range of choices* of what to do, making the player unpredictable at this point in the game. Additionally, a wide range of likely choices (in general) can contain a small number of less-likely ones without reducing the overall width – this supports the notion that where some values are small for one seed, it was because they were not essential for the range, and not because they were inherently bad.

Figure 5, whose plots are in the same format as those in figure 4, shows the probability distributions for two situations where the game could end on the next turn. In (a), the top plot shows the range of possibilities for a player in state (15,9,4), and the bottom shows the range for state (9,15,2), the exact opposite situation. In the top plot, as before, it can be seen that the moves which make the game un-winnable have low probabilities of being chosen, and across the remaining options there is an (approximately) unpredictable choice. The situation is similar in the bottom plot of (a), but where the lower probabilities at the high end represent the fear of failing to win, more than the fear of losing in that turn.

In 5(b) – which represents a similar pair of situations to (a) – the curves are somewhat noisier. The top plot, with its relative uniformity at the left and its pair of very small values at the right, has the appearance of being only ‘half evolved’ compared to its counterpart in (a). This was probably on account of the low mutation rate, which is now discussed.

A characteristic of this GA is that not all (in fact, very few) genes in the chromosome contribute to the phenotype. This means that instead of selective pressure being continually applied everywhere, it is only applied across a different subset of loci each time. The absence of selective pressure for variable periods of time on a given gene means that it can be mutated negatively with no consequences. A non-low

mutation rate therefore becomes a dangerous thing, because good genetic building-blocks will most probably be corrupted while they are non-expressing. And if the chromosome is long (as is the case here), a low mutation rate equates to little exploration, so a difficult trade-off arises: safe building blocks & insufficient exploration versus sufficient exploration & unsafe blocks.

Figure 6 shows what happens when there is too much mutation/exploration. In 6(a), it can be seen that though a desirable *average* genotype emerged after 1000 generations, the standard deviations of the population-wide alleles were generally high, indicating variation. Had this GA been seeking a precise genotype, as opposed to the sort it has been reported to have been, then this variation would have represented corruption.

6(b) shows the result of an attempt to bypass the corruption problem of non-low mutation whilst still having good exploration. Random immigrants were added in every generation, with the hope being that good genes from good immigrants would be able to enter the population, and bad genes would not. The lower plot in (b) shows that in fact, the immigration scheme represented an effective *raising* of the standard mutation rate. This was because enough immigrants were surviving to pass their genes into the population via crossover, injecting more than a moderate mutation operator would. And the dotted curve in the upper plot of (b) communicates the fact that having a mere 5% of immigrants per generation is comparable to randomising the individuals.

4 Conclusions

Deterministic and probabilistic strategies were evolved for the game Footsteps. It was found that the deterministic kind are flawed because of their exploitability, and that the probabilistic kind – when they are unpredictable – are good. The first conclusion of this paper is therefore that in situations of attack or counterattack, the best general strategy is unpredictability (which includes deceptiveness as well as randomness). If an opponent knows how you will behave (even if you are using biased probabilities), then he has the opportunity to exploit you. An example of this from the real world is the football match between Arsenal FC and Aston Villa FC in the 2002/03 season of the English Premier League. Arsène Wenger’s Arsenal had acquired a reputation for fast counterattacking football, so Graham Taylor had his Villa side play defensively

for most of the match. The game was drawn (giving Villa one point instead of zero), and afterwards, Wenger criticised Taylor's negative attitude. Taylor's reply was that the alternative would have been to hand them the game.

The second conclusion is that in GA's where only *some* of the genes are used in any given phenotypic development, convergence should be measured phenotypically, not genotypically. This is because the population could converge phenotypically on an acceptable solution, without being converged genotypically, giving the impression – from the genetic point of view – that no solution had been found. The consequences of using a misleading convergence metric include wasting time and energy on unnecessary searching, and overlooking solutions altogether.

The third and final conclusion concerns GA's with long chromosomes where the selective pressure is not universally applied, for whatever reason. In these GA's, what was already a difficult search is further compounded by the problem of mutation. The higher the mutation rate, the more the chromosomes will be damaged by unrejected mutations, but the lower the rate, the weaker the search and the less the capacity to evolve. If the representation of a problem puts its GA into this category, then a more sophisticated technique than artificial evolution ought to be used instead. This is not a criticism of evolutionary computing, but a delimitation in its range of viable application.

References

- [1] http://everything2.com/?node_id=888736.
- [2] <http://vying.org/games/footsteps>.
- [3] <http://www.diracdelta.co.uk/science/source/e/q/equilibrium/source.html>.
- [4] F. F. Ali, Z. Nakao, and Y. W. Chen. Playing the rock-paper-scissors game with a genetic algorithm. *Proceedings of the 2000 Congress on Evolutionary Computation*, pages 741–745, 2000.
- [5] R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. *Genetic Algorithms and Simulated Annealing*, pages 32–41, 1987.
- [6] K. Chellipilla and D. B. Fogel. Evolution, neural networks, games, and intelligence. *Proceedings of the IEEE, Vol. 87, No. 9*, pages 1471–1498, 1999.
- [7] K. Chellipilla and D. B. Fogel. Anaconda defeats hoyle 6-0: A case study competing an evolved checkers program against commercially available software. *Proceedings of the Congress on Evolutionary Computation*, pages 857–863, 2000.
- [8] A. M. Colman. *Game Theory and its Applications in the Social and Biological Sciences*. International Series in Social Psychology. Routledge.
- [9] P. J. Darwen and X. Yao. *On Evolving Robust Strategies for Iterated Prisoner's Dilemma*, pages 276–292. Progress in Evolutionary Computation, Vol. 956. Springer, 1995.
- [10] J. E. Davis and G. Kendall. An investigation, using co-evolution, to evolve an awari player. *Proceedings of the Congress on Evolutionary Computation*, pages 1408–1413, 2002.
- [11] D. Fogel. Evolving behaviours in the iterated prisoner's dilemma. *Evolutionary Computation, Vol. 1*, pages 77–97, 1993.
- [12] J. J. Grefenstette. *Handbook of Genetic Algorithms*, chapter 14: Strategy Acquisition with Genetic Algorithms. Von Nostrand Reinhold, 1991.
- [13] J. J. Grefenstette. Genetic algorithms for changing environments. *Parallel Problem Solving from Nature 2*, 1992.
- [14] O. G. Haywood. Military decision and game theory. *2nd Journal of the Operations Research Society of America*, pages 365–385, 1954.
- [15] D. Koller and A. Pfeffer. Representations and solutions for game-theoretic problems. *Artificial Intelligence, Vol. 94*, pages 167–315, 1997.
- [16] R. W. Rosenthal. Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, pages 92–100, 1981.
- [17] J. M. Smith. The evolution of behavior. *Scientific American*, pages 136–145, 1978.
- [18] A. Wald. Statistical decision functions which minimize maximum risk. *Annals of Mathematics*, pages 265–280, 1945.




Figure 2: (Deterministic strategies.) The population-wide standard deviations of the gene values at the loci that correspond to the game states where the counter is in the middle of the board. The left half shows them after 20 generations, and the right half after 200, from the same run. Square brackets indicate that that state was visited at least once by at least one individual in that generation's league. It can be seen in the right half of the figure that though the population is unconverged genotypically, it is converged phenotypically.




Figure 3: (Deterministic strategies.) The number of different game-states visited in total in each generation of a run: (a) between generations 50 and 250, and (b) between generations 0 and 250 of the same run.

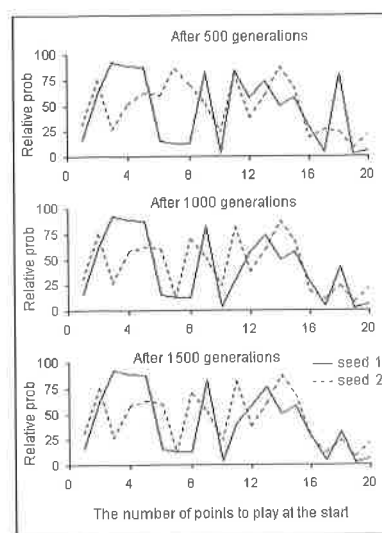


Figure 4: (Probabilistic strategies.) The population-wide averages of the gene values that refer to the starting state; in other words, the probability distribution for the first turn. Shown for two random seeds and at three generation marks. The selective advantage in low values at the right and mixed values elsewhere is visible. That the populations converged in that part of the chromosome is also visible.

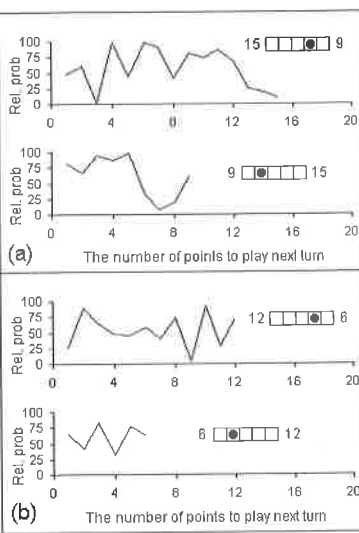


Figure 5: (Probabilistic strategies.) The population-wide averages of the gene values referring to both sides of two situations where the game could be won in the next turn, taken after 1000 generations. The inset diagrams show the exact situations, and the plot belongs to the left player. A reluctance to render the game unwinnable (i.e. small values at the right) is visible in the top plot in (a), but not in the same place in (b).

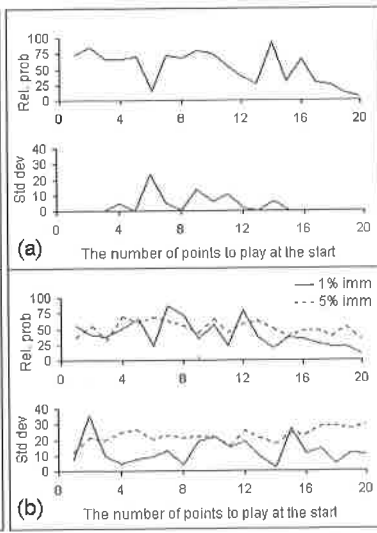


Figure 6: (Probabilistic strategies.) The population-wide means and standard deviations of the gene values that refer to the starting state, taken after 1000 generations, (a) for a mutation rate 10 times bigger than was used for the other results, and (b) with 1% and 5% random immigrants added in every generation. The standard deviations – which were mostly zeroes for the other results – show the level of persistent diversity, an indicator of corruption.