

Analysis of the Difficulty of Learning Goal-Scoring Behaviour for Robot Soccer

Jeff Riley

School of Computer Science & Information Technology
RMIT University
Melbourne, Australia
+61 3 8877 2967

jeff.riley@iinet.net.au

Vic Ciesielski

School of Computer Science & Information Technology
RMIT University
Melbourne, Australia
+61 3 9925 2926

vc@cs.rmit.edu.au

ABSTRACT

Learning goal-scoring behaviour from scratch for simulated robot soccer is considered to be a very difficult problem, and is often achieved by endowing players with an innate set of hand-coded skills, or by decomposing the problem into learning a set of simpler behaviours which are then aggregated into goal-scoring behaviour. When only basic skills are available to the player the fitness landscape is very flat, containing only a few thin peaks. As more human expertise is injected via hand-coded skills or a composite fitness function, more gradient information becomes apparent on the landscape and the genetic search is more successful. The work presented in this paper uses *autocorrelation* and *information content* measures to examine features of the fitness landscape to explain how the difficulty of the problem is changed by injecting human expertise.

Track: Learning Classifier Systems And Other Genetics-Based Machine Learning

Categories and Subject Descriptors

F.2.0 [Theory of Computation]: Analysis of algorithms and problem complexity.

General Terms

Algorithms, Performance, Theory.

Keywords

Genetic Algorithm, Fuzzy Systems, Fitness Landscapes.

1. INTRODUCTION

A wide variety of approaches and technologies have been used in attempts to construct good robot soccer players. These include hand-coding, genetic algorithms, genetic programming, reinforcement learning, neural networks, behaviour-based and deliberative agents, and various combinations of those. In the early years of the RoboCupSoccer [7] competition a few

researchers attempted to fine tune some low-level skills using AI machine learning techniques, but nearly all entrants used hand-coded skills and strategies [9]. Even today hand-coded players, or players with hand-coded skills, generally continue to outplay players whose skills have been entirely learned or developed automatically [8]. For example, the 2005 RoboCupSoccer 2D simulation league winner used hand-coded strategies which employed a mixture of hand-coded skills and skills developed using machine learning techniques [10]. There has been only limited success when applying standard machine learning techniques to this problem – much of the work to date has been characterised by researchers beginning their work with high expectations, then ratcheting down their expectations as the work progresses, and finally adjusting their goals (and the soccer playing behaviours and skills of the players being developed) to align with the progress being made.

The concept of fitness landscapes, and the idea that the process of evolution could be studied by visualising the distribution of fitness values across the population as a landscape, has been long-established in the field of evolutionary biology, having been first proposed by Wright in 1932 [15] and revived later in [1].

Much of the work involving fitness landscapes avoids a rigorous definition of the landscape under analysis [5], and where it is mentioned or implied at all the landscape is usually assumed to be the *single-bit mutation* landscape: the landscape generated by arranging all single-bit mutations of a chromosome represented as a string of binary digits such that chromosomes that differ by only a single bit are neighbours. On such landscapes, genetic operators such as crossover are assumed to take *hypersteps* over the fitness landscape described by mutation.

The major area of concern with fitness landscapes is that there is no generally accepted definition of what constitutes a fitness landscape. There is not much agreement in the field as to what a fitness landscape is and whether a neighbourhood relation is required to describe it, and much less agreement as to what the neighbourhood relation should be. Jones' "one operator one landscape" approach [5] does have a core following, however work continues to try to present a coherent, consistent view of fitness landscapes and the neighbourhood relations that define them, as well as the methods that are used to measure them.

The goal of this work is to describe a method for analysis of the fitness landscape described by the problem of learning goal-scoring behaviour, and to use the analysis to better understand the difficulty of the problem and how progress might be made.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '06, July 8–12, 2006, Seattle, Washington, USA.

Copyright 2006 ACM 1-59593-186-4/06/0007...\$5.00.

2. EVOLVING THE PLAYER

For this work a messy-coded genetic algorithm [2] is used to evolve a single robot soccer player in a simulated soccer environment [11]. In this work a trial begins with a single player and the ball being placed in random locations on the field, and continues for a fixed period of time. Any time the player scores a goal the player and the ball are set to new random locations and the trial continues until the time expires. The behaviour of the player is governed by a fuzzy inferencing system [4] with the ruleset for the fuzzy inferencing system being evolved by the genetic algorithm.

The player being evolved is endowed with a configurable subset of soccer-playing skills taken from the full set of skills shown in Table 1. In addition, if the player is unable to determine an action to be taken based on the information known to it, the player will perform a hand-coded default hunt action (on the basis that the most likely cause for a player not being able to determine an action is that the ball is not visible). The hand-coded hunt actions available to be configured as default actions are listed in Table 2.

The player will perform one of the available actions, or the configured default action, in response to external stimuli; the specific response being determined by the fuzzy ruleset. The external stimulus used as input to the fuzzy inference system is the visual information supplied by the soccer simulator, after undergoing fuzzification. The output of the fuzzy inference system is an (*action, value*) pair which defines the action to be taken by the player, and the degree to which the action is to be taken. For example:

(KickTowardGoal, power)
 (RunTowardBall, power)
 (Turn, direction)

where *power* and *direction* are crisp values representing the defuzzified fuzzy set membership of the action to be taken. An example rule developed by the genetic algorithm is:

if Ball is Left and Goal is Left then Turn Slightly Left

Chromosomes for the genetic algorithm are not fixed length: the length of each chromosome in the population varies with the length of individual rules and the number of rules on the chromosome. The number of clauses in a rule and the number of rules in a ruleset is only limited by the maximum size of a chromosome which for this work was 68 genes.

Table 2. Available Hunt Actions

Action	Description
Hunt Action 1 <i>Goto Ball</i>	if the ball is not visible then dash in a randomly chosen direction else if ball is not in kickable distance then dash toward the ball else do nothing
Hunt Action 2 <i>Locate Ball</i>	if the ball is not visible then dash in a randomly chosen direction else do nothing
Hunt Action 3 <i>Random Turn</i>	turn 90° in a randomly chosen direction

Table 1. Available Player Skills

Skill	Description
Turn	The player turns through the angle specified. Argument: angle.
Dash	The player dashes in the direction specified with the power specified. Arguments: direction, power.
Kick	If the ball is within a kickable distance from the player, the player kicks the ball in the direction specified with the power specified. Arguments: direction, power.
RunTowardGoal	If the direction to the player's goal is known, the player dashes once in that direction, otherwise no action is taken. Argument: power.
RunTowardBall	If the direction to the ball is known, the player dashes once in that direction, otherwise no action is taken. Argument: power.
GoToBall	If the direction to the ball is known, the player dashes towards the ball and continues to dash in that direction until the ball is within the kickable distance, otherwise no action is taken. Argument: power.
KickTowardGoal	If the direction to the player's goal is known, and the ball is within the kickable distance, the player kicks the ball once in the direction of its goal, otherwise no action is taken. Argument: power.
DribbleTowardGoal	If the direction to the player's goal is known, and the ball is within the kickable distance, the player kicks the ball once in the direction of its goal, then dashes once in the same direction. If the direction to the player's goal is not known, or the ball is not within the kickable distance, no action is taken. Argument: power.
Dribble	If the ball is within the kickable distance, the player kicks the ball once in the direction it is facing, then dashes once in the same direction. If the ball is not within the kickable distance, no action is taken. Argument: power.
DoNothing	The player takes no action.

3. THE FITNESS LANDSCAPE

3.1 Fitness Landscape Definition

A review of relevant literature [e.g. 3, 5, 6, 12, 13, 14] indicates there are several possible definitions of, and representations for, fitness landscapes, and choosing the definition and representation which best describes the combination of the problem being studied and the algorithm being used to study it is extremely important. A fitness landscape is most often defined by three basic attributes:

- a search space
- a relation that defines which points are neighbours in the search space
- a fitness function that assigns a fitness value to each point in the search space

The neighbourhood relation and its specification is extremely important because any discussion of landscapes invariably involves the terms “peaks” and “valleys”, and no peak or valley can exist without the notion of neighbourhood – a peak is only a peak because it is higher than its neighbours.

For this work the fitness landscape is considered to be defined by the overall operation of the genetic algorithm. Consider an observer watching a genetic algorithm searcher perform a random walk on a fitness landscape and assume that although the observer is able to discern the granularity of the search (the genetic algorithm’s single steps), the means by which the GA determines where each step takes it is hidden from the observer. The random walk is conducted as follows:

- An individual i_0 is randomly selected from the search space
- For each step $s, s = 1 .. \text{maxsteps}$
 - i_{s-1} undergoes mutation with probability P_{mutation}
 - Another individual $i, i \neq i_{s-1}$, is randomly selected from the search space
 - Crossover is performed between i and i_{s-1} with probability $P_{\text{crossover}}$, resulting in two new individuals i'_1 and i'_2 , both of which are neighbours of (a single step from) i_{s-1}
 - Set $i_s = i'_1$ and step to i_s

The observer sees the searcher walking randomly over the landscape and considers points on the landscape one step apart to be neighbours. The definition of the neighbourhood relation is of no consequence to, and is not required by, the observer since the searcher is defining neighbouring points by performing the walk. If the random walk performed by the genetic algorithm searcher was sufficiently long, and the “altitude” (fitness) at each step recorded for the observer, the entire fitness landscape would be determined by observation. The landscape so determined would be the precise fitness landscape defined by the search algorithm.

This “black box” view of the genetic algorithm operation and consequential determination of the neighbourhood relation and fitness landscape satisfies the requirement that the landscape neighbourhood relation be defined by the search algorithm, and is the definition used for the robot soccer problem addressed by this work.

3.2 Fitness Landscape Measures

The methods used to measure and analyse the structure of fitness landscapes in this work are the autocorrelation method suggested by Weinberger [14], and the information content approach suggested by Vassilev et al. [12, 13].

3.2.1 Autocorrelation and Correlation Length

Weinberger’s autocorrelation definition [3, 14]:

Given measurements, Y_1, Y_2, \dots, Y_N at time X_1, X_2, \dots, X_N , where N is the number of measurements, and

$$\bar{Y} = \frac{1}{N} \sum_{i=1}^N Y_i, N > 0$$

the time lag k autocorrelation function r_k is defined as

$$r_k = \frac{\sum_{i=1}^{N-k} (Y_i - \bar{Y})(Y_{i+k} - \bar{Y})}{\sum_{i=1}^N (Y_i - \bar{Y})^2}, N > k$$

If $|r_k| \approx 1.0$ there is much correlation between the points k steps apart in the series, whereas if $|r_k| \approx 0.0$ there is little correlation.

Weinberger proposed that a random walk be generated on the fitness landscape, where each step on the walk is taken between neighbouring points, with the neighbour to which the step is taken selected randomly. The fitness values for the points visited during the random walk form a time series of numbers. The autocorrelation function can then be used as a measure of the ruggedness of the landscape described by the random walk.

The correlation length of a series of numbers is the largest distance, or time lag, between points for which some correlation exists. Hordijk [3] defines the correlation length of a time series as one less than the first time lag for which the autocorrelation falls inside the region bounded by the two-standard-error bound (i.e. one less than the first time lag at which the autocorrelation becomes statistically equal to zero, making the correlation length the largest time lag for which the correlation between two points is still statistically significant). This is the method used for calculating the correlation length in this work. The two-standard-error bound e is defined as

$$e = \pm \frac{2}{\sqrt{N}}$$

so the correlation length ℓ is defined in this work as the first lag k for which

$$|r_k| < \left| \frac{2}{\sqrt{N}} \right|$$

3.2.2 Information Content

Vassilev et al. propose three information measures that characterise the structure of a fitness landscape from a series of points generated by a random walk over the landscape [12, 13]:

- Information Content – characterises the ruggedness of the landscape.
- Partial Information Content – measures the modality of the landscape.
- Information Stability – the sensitivity of the information content measures.

These measures are calculated by generating a random walk of length n on the fitness landscape, with the aim being to extract information by characterising the series of points as an ensemble of objects.

To calculate the information content, a string

$$S(\varepsilon) = s_1 s_2 \dots s_n, s_i \in \{\bar{1}, 0, 1\}$$

representing a group of objects generated from a random walk over the fitness landscape. $S(\varepsilon)$ is enumerated according to the function:

$$s_i = \psi_{f_i}(i, \varepsilon), \text{ for } t = 1..n$$

and

$$\psi_{f_i}(i, \varepsilon) = \begin{cases} \bar{1} & \text{if } f_i - f_{i-1} < -\varepsilon \\ 0 & \text{if } |f_i - f_{i-1}| \leq \varepsilon \\ 1 & \text{if } f_i - f_{i-1} > \varepsilon \end{cases}$$

Thus the string $S(\mathcal{E})$ defines a sequence of objects where each object is represented by a substring $S_i S_{i+1}$ being a sub-block of length two of the string $S(\mathcal{E})$.

The parameter \mathcal{E} is a real number taken from the interval $[0.0, 1.0]$ (in this case) which defines neutral fitness and determines the accuracy with which the string $S(\mathcal{E})$ is defined. If the absolute fitness difference between neighbouring points is less than \mathcal{E} the points are considered to be of equal fitness. This means that as \mathcal{E} increases from 0.0 to the maximum possible fitness difference between points along the walk (1.0), the amount of fitness change (entropy) and the sensitivity of Ψ_{f_i} decrease to zero.

The information content is defined as the entropic measure of the group of sub-blocks of length two of string $S(\mathcal{E})$, and is given by

$$H(\mathcal{E}) = -\sum_{p \neq q} P_{[pq]} \log_6 P_{[pq]}$$

$P_{[pq]}$ are frequencies of the possible blocks pq of elements from the set $\{\bar{1}, 0, 1\}$ given by

$$P_{[pq]} = \frac{n_{[pq]}}{n}$$

where $n_{[pq]}$ is the number of occurrences of pq in $S(\mathcal{E})$ and n the length of $S(\mathcal{E})$.

The partial information content is a measure of the modality (the number or frequency of local optima) of the landscape, and is calculated by filtering out elements of the string $S(\mathcal{E})$ which are not essential for measuring modality to create a new string $S'(\mathcal{E})$, then measuring the length μ of the new string $S'(\mathcal{E})$.

The string is defined as

$$S'(\mathcal{E}) = s_{i_1} s_{i_2} \dots s_{i_j}, s_{i_j} \neq 0, s_{i_j} \neq s_{i_{j-1}}, j > 1$$

Thus the string $S'(\mathcal{E})$ has the form " $\bar{1} \bar{1} \bar{1} \dots$ ", representing the slopes of the path taken by the random walk over the landscape, and so is empty if $S(\mathcal{E})$ contains only 0s. The partial information content is given by

$$M(\mathcal{E}) = \frac{\mu}{n}$$

Note that when $M(\mathcal{E})$ is 1, the path taken by the random walk over the landscape is considered to be maximally multimodal, and when $M(\mathcal{E})$ is 0, the path is flat.

The information stability \mathcal{E}^* is defined as the smallest value of \mathcal{E} for which the landscape becomes flat (i.e. for which $S'(\mathcal{E})$ is empty). Since \mathcal{E} governs the sensitivity of the information content and partial information content measures, \mathcal{E}^* is a measure of the difference in fitness between neighbouring points on the random walk.

4. EXPERIMENTS & ANALYSIS

A number of experiments were performed to compare the effect on the performance of the evolutionary search of varying the player skillset, the default action and the fitness function. In addition to the evolutionary search, five random walks (as described earlier) were conducted for each experiment, each walk starting at a randomly selected point on the fitness landscape and continuing for a duration of 100,000 steps. The statistics gathered during the walks are also analysed.

For each experiment the following information is presented for analysis and comparison:

- Data from the evolutionary search:
 - Line graphs showing the population average fitness and individual best fitness for generations 1 to 500.
 - A bar chart showing the cumulative fitness distribution for all individuals evaluated during the 500 generations, showing the percentage of all individuals evaluated which failed to kick the ball (denoted by "nm" on the graph x-axis), moved the ball closer to the goal (graduated), kicked one goal, two goals, three goals etc.
- Data from the random walks over the fitness landscape:
 - A correlogram showing the autocorrelation data for time lags 1 to 100 for the five random walks and the two-standard error bounds.
 - A graph showing the information content data for $0 \leq \mathcal{E} \leq 1.0$ for the five walks.

Two different fitness functions were compared – a *composite* fitness function and a *simple, goals-only* fitness function. For both fitness functions implemented the fitness values range from 0.0 to 1.0, with 1.0 being the worst fitness possible, and optimal fitness values approaching 0.0.

The simple goals-only fitness function rewards a player for goals scored only – the greater the number of goals scored the greater the reward. The goals-only fitness function was implemented as:

$$f = \begin{cases} 1.0 & , goals = 0 \\ \frac{1.0}{2.0 \times goals} & , goals > 0 \end{cases}$$

where *goals* is the number of goals scored by the player.

The composite fitness function rewards, in order of importance:

- the number of goals scored in a trial
 - minimising the distance of the ball from the goal
- This combination was chosen to reward players primarily for goals scored, while players that do not score goals are rewarded on the basis of how close they are able to move the ball to the goal on the assumption that a player which kicks the ball close to the goal is more likely to produce offspring capable of scoring goals. This decomposes the original problem of evolving goal-scoring behaviour into the two less difficult problems:
- evolve ball-kicking behaviour that minimises the distance between the ball and goal, and
 - evolve goal-scoring behaviour from the now increased base level of skill and knowledge

The composite fitness function was implemented as:

$$f = \begin{cases} 1.0 & , kicks = 0 & , goals = 0 \\ 0.5 + \frac{dist}{2.0 \times fieldLen} & , kicks > 0 & , goals = 0 \\ \frac{1.0}{2.0 \times goals} & & , goals > 0 \end{cases}$$

where *goals* = the number of goals scored by the player
kicks = the number of times the player kicked the ball
dist = the minimum distance of the ball to the goal
fieldLen = the length of the field

4.1 Exp.1: Comp. Fitness, All Skills, Hunt Action 1

In the first experiment the search algorithm has been given the most help in the form of initial player skills and the default action: the player has been given all available skills and hunt action 1 as the default action. The fitness function is the composite fitness function, rewarding ball movement as well as goal-scoring. The data shown on the graph of population average fitness (Fig 1a) tends to indicate that the population as a whole ceases to improve after 30 to 40 generations though, as evidenced by the graph of best fitness values, individuals of good fitness continue to be found beyond that point. The percentage of the population exhibiting ball-kicking or goal-scoring behaviour is reasonably high, as shown by the frequency distribution (Fig 1b).

The autocorrelation graph and associated correlation length (Fig 1c) indicate that the fitness landscape for this problem (as described by the random walk) offers a reasonable amount of gradient information that the search algorithm can use to guide the search. With an autocorrelation of ~ 0.32 for points on the random walk a single step apart and a fairly steep descent for points further apart, the correlation between next and near neighbours on this fitness landscape is not so high that a search algorithm is led unerringly to a solution, but with a good correlation and a long correlation length the problem, in this form, should be readily solved by a search algorithm able to take advantage of the landscape features.

The information content graph (Fig 1d) supports the autocorrelation data for this experiment. Information stability is quite high at 0.885, indicating a high difference in fitness among neighbouring points, so pointing to some good gradient information being present in the landscape. $H(0.0)$ is not particularly large, indicating that the diversity of shapes on the landscape is not high. Similarly $M(0.0)$ is relatively small, indicating that the degree of modality of the landscape is low.

4.2 Exp.2: Comp. Fitness, Base Skills, Hunt Action 1

The difference between this experiment and experiment 1 is that instead of the player being endowed with all available skills, the player in this experiment has only the base skills of turn, kick and dash. The player has hunt action 1 configured as the default action. The fitness function is the composite fitness function.

The data shown on the graph of population average fitness (Fig 2a) indicate that improvement of the population stops at about generation 150, and although the graph of best fitness values indicates that individuals exhibiting goal-scoring behaviour continue to be found, terminating the search after generation 150 would not have adversely affected the result. Figure 2b shows that the percentage of the population exhibiting goal-scoring behaviour is extremely small, with a very large proportion of the population not kicking the ball at all.

These results show clearly the effect of removing from the players a range of mid-level hand-coded skills, and raise the question of what effect removing those skills has on the structure of the fitness landscape and how that affects the search.

The autocorrelation graph and associated correlation length (Fig 2c) indicate that the fitness landscape for this problem offers only a limited amount of useful gradient information that the search algorithm can use to guide the search. With an autocorrelation of ~ 0.1 for points on the random walk a single step apart and falling to zero for points just a few steps further

apart, the correlation between next and near neighbours on this fitness landscape indicates that the structure of the fitness landscape is close to random and not as conducive to search as was the fitness landscape of experiment 1, thus increasing the difficulty of the problem. The information content graph for this experiment (Fig 2d) supports the autocorrelation data. Information stability is relatively low at 0.371, indicating a low difference in fitness among neighbouring points. With the autocorrelation data indicating a near random landscape, and information stability indicating a low fitness variation among neighbouring points, there is almost no useful gradient information in the landscape to guide the search. $H(0.0)$ is very small, indicating that the diversity of shapes on the landscape is very low. Similarly $M(0.0)$ is extremely small, indicating that the landscape lacks any real degree of modality. Both values further indicate the lack of useful landscape data to guide the search.

The data presented all indicate that the removal of a set of mid-level, hand-coded skills has changed the relative difficulty of the problem, and that this is a result of the structure and features of the fitness landscape being altered by the problem representation – what was a landscape reasonably rich in features that helped guide the search has become a relatively barren landscape lacking in information useful for search.

4.3 Exp.3: Comp. Fitness, All Skills, Hunt Action 3

For experiment 3 the player is again given all available skills, but the default action is limited – in this case the default action is just to turn 90° in a randomly chosen direction. The fitness function is the composite fitness function.

The results show that this limiting of the default hunt action negatively affects the search. The population average fitness remains high (Fig 3a), and the percentage of the population exhibiting ball-kicking or goal-scoring behaviour is low, as shown by the frequency distribution (Fig 3b) – significantly lower than the previous experiments. The autocorrelation and information content data shown in Figures 3c and 3d paint a similar picture – from this data it is clear that the fitness landscape lacks much of the gradient information seen in the earlier experiments. This is entirely due to the removal of any sort of intelligence from the default hunt action.

4.4 Exp.4: Goals Fitness, All Skills, Hunt Action 1

In experiment 4 the player is given all available skills, and the default action is configured to be hunt action 1. The fitness function for this experiment is the goals-only fitness function, so the player is rewarded only for scoring goals.

The results for experiment 4 are very similar to those for experiment 1, with the major difference being the autocorrelation and information content data (Fig 4c & 4d). The population average fitness is, as expected for a goals-only fitness function, higher for experiment 4, but the best fitness and fitness frequency graphs are almost identical (for goal-scoring behaviour in the case of the fitness frequency graph) (Fig 4a & 4b). The autocorrelation and information content graphs (Fig 4c & 4d) indicate that the fitness landscape has somewhat less gradient information useful for search, but still sufficient to facilitate a successful search. This is another indication that when the players are given the full complement of hand-coded skills the difficulty of the problem is reduced significantly.

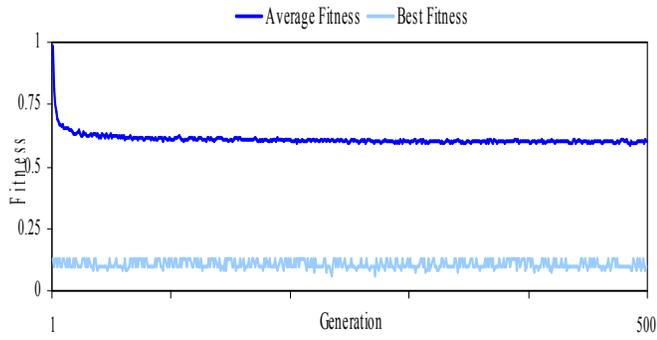


Figure 1a

Note that a lower fitness indicates more goals scored.
The best players are scoring between 4 and 6 goals

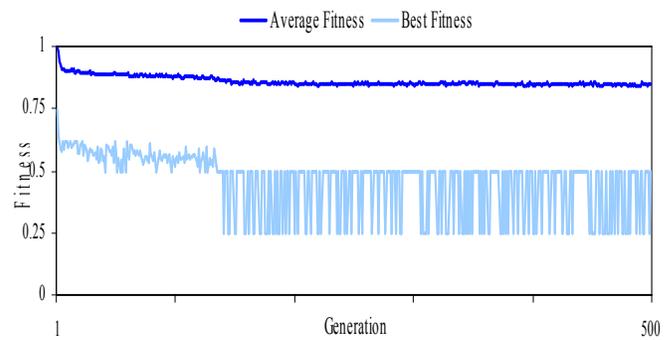


Figure 2a

After generation 130 the best player's score alternates between 1 and 2 goals

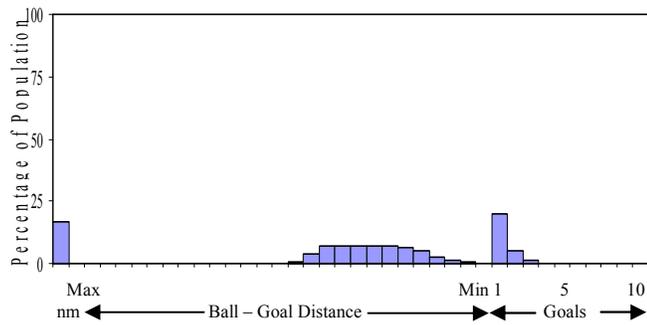


Figure 1b

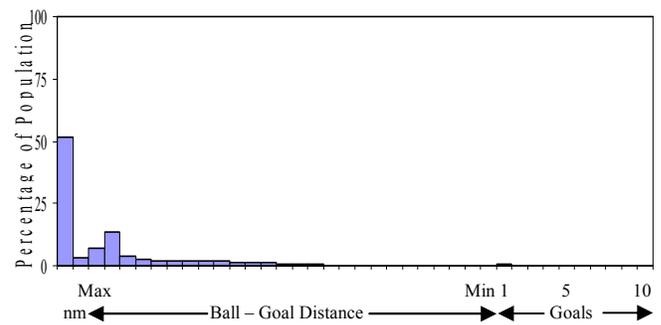


Figure 2b

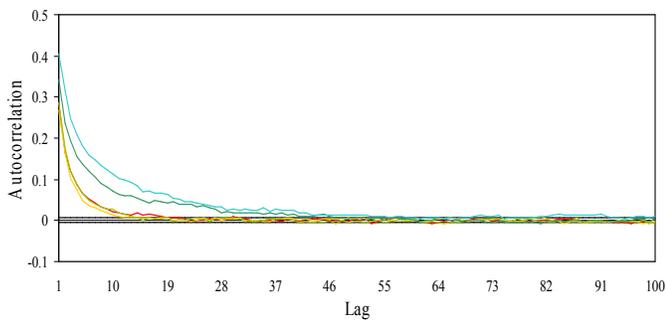


Figure 1c

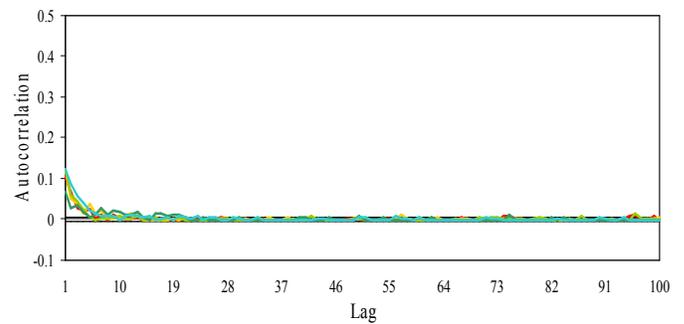


Figure 2c

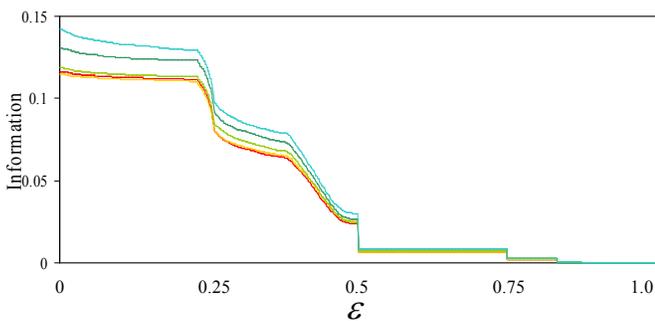


Figure 1d

Figure 1 Composite Fitness, All Skills, Hunt Action 1

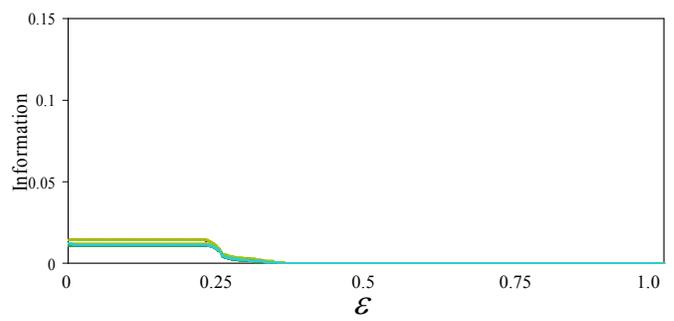


Figure 2d

Figure 2 Composite Fitness, Base Skills, Hunt Action 1

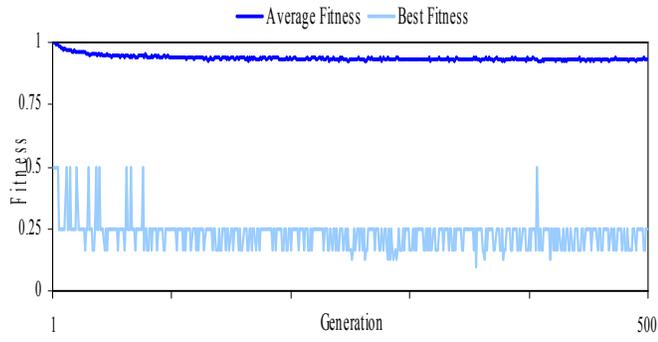


Figure 3a

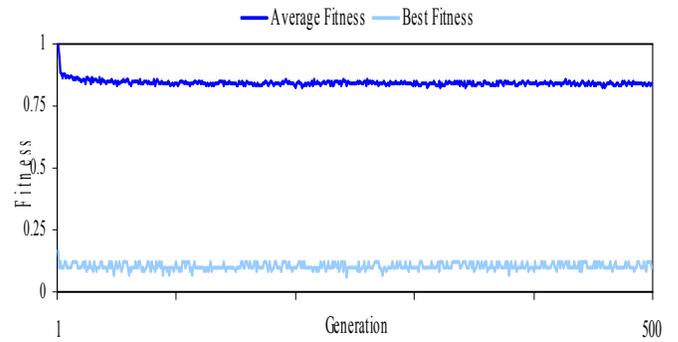


Figure 4a

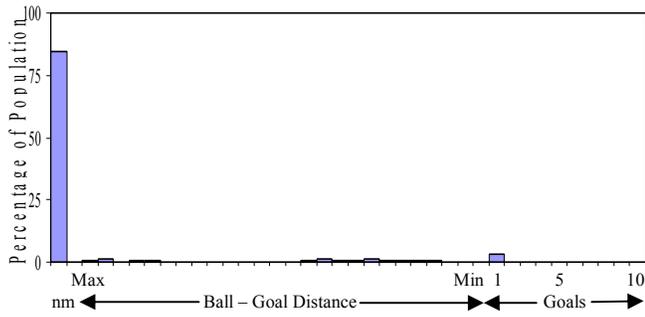


Figure 3b

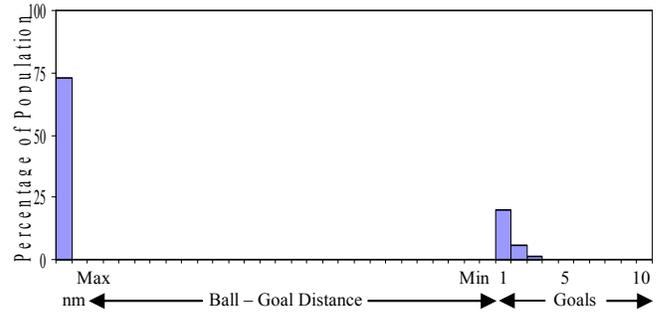


Figure 4b

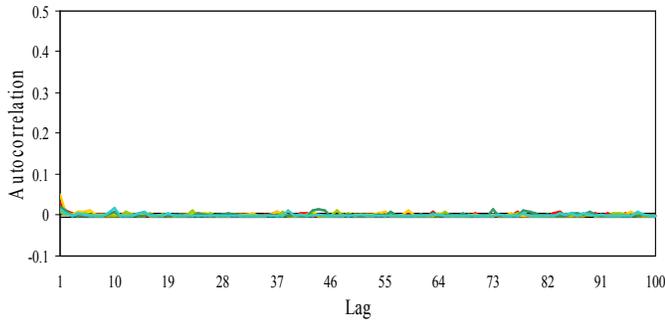


Figure 3c

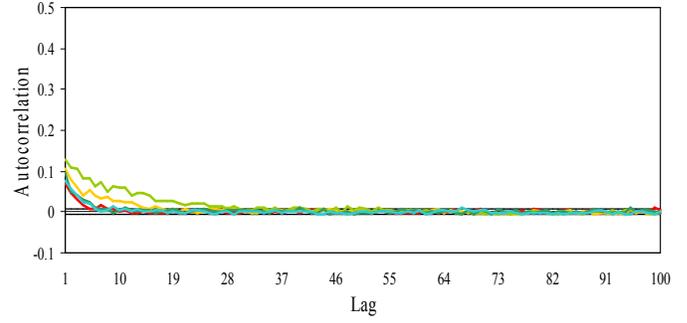


Figure 4c

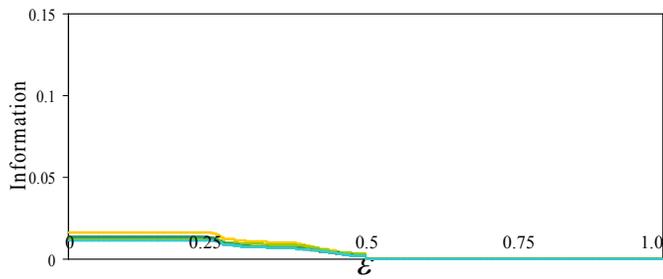


Figure 3d

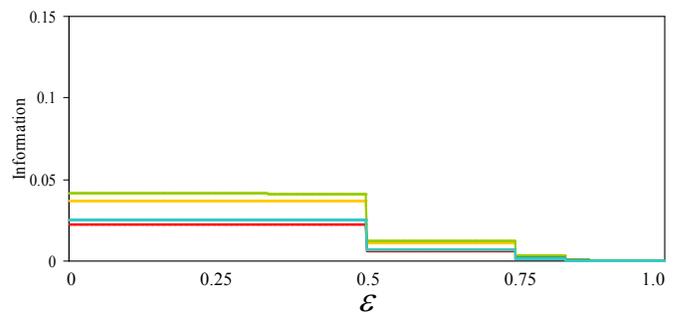


Figure 4d

Figure 3 Composite Fitness, All Skills, Hunt Action 3

Figure 4 Goals-Only Fitness, All Skills, Hunt Action 1

5. CONCLUSIONS & FURTHER WORK

This work has presented a method for the analysis of the fitness landscape described by the problem of learning goal-scoring behaviour using a genetic algorithm. The analysis indicates that when human expertise and expert knowledge is injected into an evolutionary search algorithm via hand-coded innate skills, smart default actions and a composite fitness function to guide the search, the problem of learning goal-scoring behaviour for robot soccer is solvable. That is, the genetic algorithm is able to evolve individuals which display goal-scoring behaviour to the extent that they are able to consistently score multiple goals in the tests conducted.

The fitness landscape analysis further indicates that as human expertise and expert knowledge is removed from the algorithm by restricting the hand-coded innate skills and smart default actions available to the players, or by using a simple goals-only fitness function, at some threshold the problem becomes intractable for the evolutionary algorithm. This suggests that while there may be gradient information in the fitness landscape, as the human expertise is reduced the density of solutions in the search space becomes very low the “mountain ranges” in the landscape begin to become isolated from each other, and the landscape begins to appear as a flat plain, sparsely populated by individual peaks – so the problem begins to resemble a *needle-in-a-haystack* problem. In this case since the genetic algorithm is not able to locate the sparsely distributed gradient information any way other than by randomly sampling the search space it performs little or no better than random search - confirmed by a small number of tests performed comparing genetic and random search. This is an indication that the injection of human expertise and expert knowledge acts like a magnifying glass to the searcher - as more expertise and knowledge is injected the fitness landscape features conducive to search are magnified, and as the expertise and knowledge is removed those landscape features become less discernable. As the granularity of the injected knowledge is increased (e.g. a richer set of skills) the modality of the landscape decreases and the gradients between peaks becomes smoother.

This is one of the underlying causes of the difficulty of the robot soccer problem for evolutionary algorithms, and the analysis presented in this work suggests that with a difficult problem such as robot soccer an evolutionary algorithm will only find a reasonable solution if one of:

- a rich skill set (placing the initial population closer to the desired solution)
- a composite fitness function (providing a solution recipe)

is present - if both of those components are absent the problem becomes very difficult for evolutionary algorithms.

Further work to ascertain the best balance between the two components identified as being necessary for successful evolutionary search (a rich skill set or a composite fitness function) would be useful, as would work to determine if there is a limit in the level of initial skills at which a difficult problem becomes intractable.

This work proposed a definition of the fitness landscape described by the combination of the genetic algorithm and the problem being investigated, but further work needs to be done to develop a

consistent definition of the fitness landscape, or landscapes, described by the operation of an evolutionary algorithm. With a consistent landscape definition, more work can be done to develop measures that will aid researchers in tuning algorithms and search methods based on landscape analysis.

6. REFERENCES

- [1] Frauenfelder, H., et al., eds. Landscape Concepts in Physics and Biology. Special Issue of Physica D. Vol. 107, p. 2-4. 1997. Elsevier Science: Amsterdam, Holland.
- [2] Goldberg, D., B. Korb, and K. Deb, Messy Genetic Algorithms : Motivation, Analysis, and First Results. Complex Systems 3, 1989.
- [3] Hordijk, W., A Measure of Landscapes. Evolutionary Computation, 4(4): p.335-360. 1996.
- [4] Jang, J.-S., C.-T. Sun, and E. Mizutani, Neuro-Fuzzy and Soft Computing. 1997. Prentice Hall.
- [5] Jones, T., One Operator, One Landscape, in Technical Report 95-02-025. 1995. Santa Fe Institute: Santa Fe, NM.
- [6] Kauffman, S.A., Adaptation on Rugged Fitness Landscapes, in Lectures in the Sciences of Complexity, D. Stein, Ed. p.527-618. 1989. Addison Wesley.
- [7] Kitano, H., et al. The Robot World Cup Initiative in Working Notes of the 1995 IJCAI Workshop on Entertainment and AI/Alife, p.19-24. 1995. Montreal, Canada.
- [8] Lima, P., et al., RoboCup 2004 Competitions and Symposium: A Small Kick for Robots, a Giant Score for Science, in AI Magazine 6(2). 2005.
- [9] Luke, S. Genetic Programming Produced Competitive Soccer Softbot Teams for RoboCup97 in J. Koza et al., eds, Proceedings of the Third Annual Genetic Programming Conference, p.204-222. 1998. Morgan Kaufmann, San Francisco.
- [10] Riedmiller, M., et al. Brainstormers 2D - Team Description 2005 in Team Description Papers, Proceedings of RoboCup 2005 (CD) (to appear). 2005.
- [11] Riley, J. The SimpleSoccer Machine Learning Environment in S.-B. Cho et al., eds, Proceedings of the First Asia-Pacific Workshop on Genetic Programming, p.24-30. 2003. Canberra, Australia.
- [12] Vassilev, V., T. Fogarty, and J. Miller, Information Characteristics and the Structure of Landscapes. Evolutionary Computation, 8(1): p.31-60. 2000.
- [13] Vassilev, V.K. Information Analysis of Fitness Landscapes in P. Husbands and I. Harvey, eds, Proceedings of the Fourth European Conference on Artificial Life. 1997. Brighton, England: The MIT Press.
- [14] Weinberger, E., Correlated and Uncorrelated Fitness Landscapes and How to Tell the Difference. Biological Cybernetics, 63: p. 325-336. 1990.
- [15] Wright, S. The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution in D.F. Jones, ed., Proceedings of the Sixth International Congress on Genetics, p.356-366. 1932. Brooklyn NY: Brooklyn Botanic Garden.