Investigating a New Paradigm for Designing Evolutionary Optimisation Algorithms Using Social Behaviour Evolution

Mikdam Turkey* and Riccardo Poli†

School of Computer Science and Electronic Engineering
University of Essex, Colchester CO4 3SQ, United Kingdom

*mturkey@essex.ac.uk, †rpoli@essex.ac.uk

ABSTRACT

This paper describes a new approach for building evolutionary optimisation algorithms inspired by concepts borrowed from evolution of social behaviour. The proposed approach utilises a set of behaviours used as operators that work on a population of individuals. These behaviours are used and evolved by groups of individuals to enhance a group adaptation to the environment and to other groups. Each group has two sets of behaviours: one for intra-group interactions and one for inter-group interactions. These behaviours are evolved using mathematical models from the field of evolutionary game theory. This paper describes the proposed paradigm and starts studying its characteristics by building a new evolutionary algorithm and studying its behaviour. The algorithm has been tested using a benchmark problem generator with promising initial results, which are also reported. We conclude the paper by identifying promising directions for the continuation of this research.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods and Search—*Heuristic methods*

General Terms

Algorithms, Theory.

Keywords

Evolutionary Optimisation, Social Behaviour Evolution, Evolutionary Game Theory, Social Adaptive Groups.

1. INTRODUCTION

Nature has provided computer science with many sources of inspiration to develop a variety of optimisation approaches, of which natural selection or the Darwinian principle of "the survival of the fittest" has a lion's share [4]. While many types of Evolutionary Algorithms (EAs) have been developed based on Darwin's theory and our modern knowledge of genetics, rarely if ever EAs, in their original form, have naturally shown the full range of properties exhibited by natural evolution. In particular, a varieties of extensions and modifications have been necessary in order to obtain EAs that could deal with multi-modal optimisation, multi-objective optimisation and dynamic optimisation problems [13]. Under the pressure of selection, individuals with higher fitness survive for longer and/or reproduce more often. It stands to reason that, with most genetic operators and representations, this leads the population to converge into an area in the vicinity of an optimum in the fitness landscape, thereby losing diversity, and with it the ability to, for example, identify more than one optimum or to track a moving optimum. It is clear that when the natural selection process is based merely on an individual's fitness, losing diversity is an anticipated result and countermeasures have to be used. Consequently, to enable EAs to tackle these and other sorts of problems various techniques have been proposed (e.g. [2, 13]).

In this paper we propose a different approach to building EAs which can potentially deal with the problems mentioned above where populations show a natural tendency to maintain diversity and form groups. We take inspiration from the evolution of social behaviour. The approach uses a notion of *fitness of groups* which takes different measures related to a group's survival and performance into account. Each group has a set of social behaviours (operators) that individuals use in interacting with other individuals from the same or different groups. The exact nature of such behaviours is determined by a probability distribution which is tuned by an evolutionary process so as to maximise group fitness. Each behaviour serves a specific purpose and contributes to a group's survival or to the group's interaction with other groups. The behaviour probability distributions of each group are updated dynamically during the optimisation process using a dynamic mathematical model from evolutionary game theory [6].

Game theory was first introduced into evolutionary theory by Maynard Smith and Price who used it to model natural selection [9]. Subsequently many researchers have proposed models to deal with social behaviour evolution and population dynamics. In this paper, we use a simple dynamic mathematical model presented in [11] to evolve the social behaviours of groups. A distinguishing feature of our proposed approach is that the whole system is built based on notion of social behaviour evolution and evolutionary game theory. However there is some previous relevant work which was inspired by similar ideas. For example, an approach that incorporates ideas from game theory and social interaction into standard genetic algorithm to modify fitness values of individuals to slow convergence and avoid local optima was proposed in [7]. This approach uses models from game theory to represent social interaction and which improved the capability of problem solving of the standard genetic cycle. Their approach is somehow related to co-evolutionary approaches (e.g., [10]) in the dependency of an individual's fitness on its relationship with other individuals. It is worth mentioning that evolutionary game theory is different in many respects from classic game theory [12], especially in evolving the strategy (behaviour) distributions which represents the corner stone in our approach. In the general population structure and organisation, our approach has also some similarity to multi-population approaches and niching techniques (e.g. [1, 3]).

GECCO'11, July 12–16, 2011, Dublin, Ireland.

The rest of this paper is organised as follows. Section 2 introduces the proposed paradigm. Section 3 reports experimental results of initial implementation. In Section 4 we indicate some possible avenues for future work.

2. PROPOSED APPROACH

In the real world, individuals affect each other's fitness value by social behaviours. These behaviours are used and evolved according to their impact on the collective performance, so the good behaviours will survive and be adopted, while bad behaviours will die away and disappear. Social interaction behaviours can be classified into four categories according to the change (increase or decrease) they cause to the fitness values of the initiator and the recipient. These four categories are: *altruism*, *spite*, *selfishness* and *cooperation* [5]. The pay-off of some behaviours is not immediate or direct to an individual's fitness. Instead, it may increase the *relative fitness* of the group in general, which in turn enhances the individual's fitness indirectly. Based on this concept, the environment that the individual needs to be adapted to includes not only the actual environment (the fitness landscape) but also the other individuals from the same or different groups that interact with, and have influence on the individual's fitness.

In the proposed paradigm, social behaviours are used as operators to move individuals in the search space, where individuals move and form groups as they socially interact. Social behaviours are dealt with as a trait of a group of individuals that describe the way individuals behave toward other individuals from the same or different group. An appropriate representation for group behaviours is needed as well as an evolving mechanism.

The proposed evolutionary system can be described as a tuple $E = \langle X, G, V, B_{intra}, B_{inter}\rangle$ where $X = \{x_1, \dots, x_n \,|x_i \in R^{dim}\}$ represents a population of n real-valued individuals of length *dim*; G is the set of all possible groups, where $G \supset G_t = \{\, g_1, \dots, g_{N_t}\}$ represents the set of groups formed by individuals at time t; V is the group *behaviour probability distribution update function*; and, finally, B_{intra} and B_{inter} are two sets of transformations (operators) which represent the *intra-group and inter-group behaviours* used in pairwise interactions between individuals, respectively. The transformations are defined as follows:

$(x_i', x_j') = b(x_i, x_j)$, where $b \in B_{intra}$

$(x_i', x_j') = b'(x_i, x_j)$, where $b' \in B_{inter}$

where b and b' are functions that transform two individuals into two new individuals, as x_i interact with x_j. The behaviours cause to change the position of individuals in the search space. So, individuals move as they interact.

A group $g \in G_t$ is defined as $g = \langle M_t, \alpha_t, \beta_t\rangle$ where

$M_t = \{x_i, x_j \in X \,|S(x_i, x_j) \leq \tau\}$, and

$\alpha_t \in R_+^{|B_{intra}|}$ and $\sum_{b \in B_{intra}} \alpha_t(b) = 1$

$\beta_t \in R_+^{|B_{inter}|}$ and $\sum_{b' \in B_{inter}} \beta_t(b') = 1$

The function S measures the similarity between a pair of individuals and τ is a threshold. This definition means that individuals can form groups based on their similarity. In our initial implementation, we use the Euclidean distance as a similarity function (i.e., $S(x, y) = \sqrt{\sum_i (x_i - y_i)^2}$). The formation of groups is a dynamic process as individuals move freely around the search space as a result of interactions. α_t and β_t are probability distributions over B_{intra} and B_{inter} at time t, respectively, and $\alpha_t(b)$ denotes the probability of using behaviour b by the group g at time t. A function V *is* used to evolve behaviours. This is done by updating their probability distributions.

The process of evolving behaviours is based on assessing the effect of the behaviour on the group, and uses it to calculate the behaviour *pay-off*, $u(b) = F(g)e_{\theta_t}(b)$, where θ can be either α or β, $F(g)$ denotes a fitness function for groups, for $b \in B, e_{\theta_t}(b)$ represents the effect rate of behaviour b, where B= {*Cooperative, Selfish, Spiteful, Altruistic*} could be either B_{intra} or B_{inter}. The factors that should be included in calculating the fitness of a group must reflect different aspects of the group well-being and must not be based merely on the individual direct fitness values. The effect value of a behaviour measures the rate at which that specific behaviour contributes to the group fitness.

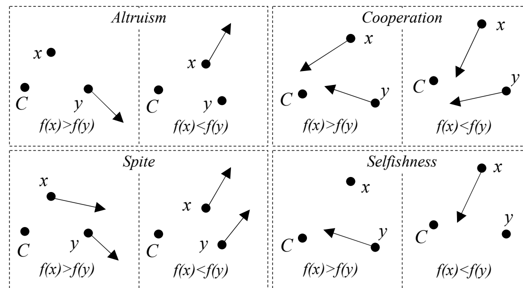

a. Intra-group interaction behaviours

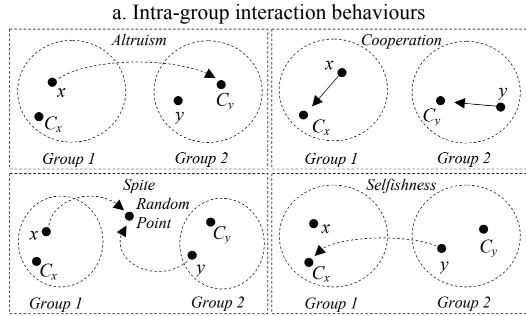

b. Inter-group interaction behaviours

Figure 1: The proposed interaction behaviours

Intra-group behaviours deal with moving individuals within the area where the group resides, whereas inter-group behaviours move individuals across group areas. The intra-group behaviour directs individuals to the promising locations in the area occupied by a group, while, at the same time, exploring the surrounding areas and maintaining a good spread in the distribution of individuals. The inter-group behaviour, instead, moves individual between groups and also move individuals randomly to new spots in the fitness landscape to investigate the possibility of forming new groups there, in case the new area has enough resources to sustain a group. In order for a group to decide how to move individuals according to the sets of behaviours, a group uses information on the local area of the fitness landscape perceived by group members. This information is synthesised in a quantity we call *group centre*. For group g, the centre is defined as $= Centre(g) = \frac{\left(\sum_{i=1}^{N_{top}} top(g,i)\right)}{N_{top}}$, where *top(g,i)* is a function that returns the i^{th} ranked member of group g according to individual fitness and N_{top} = $0.4*$ $|g|$ represents forty percent of the group size. The motion of individuals caused by social behaviour interactions uses the group centre as a reference. The change in an individual's position takes the form $x' = x + \Delta x$ where Δx is a displacement vector. Δx has to be computed in such a way to bring an individual closer to some target point and/or to push it away from some other point. Figure 1 depicts the two sets of social behaviours. In intra-group interaction behaviours (Figure 1(a)), the direction of movement of an individual is decided on the basis of the fitness value of the individual with which the individual interacts, and also the position of that individual and the centre of the group. For example, if we want to move x closer to both y and the centre of the group C, then we need to compute $d_1 = x\text{-}y$ and $d_2 = x\text{-}C$, where $\Delta x = -r_1 * bias_y * d_1 - r_2 * bias_C * d_2$, where $bias_y$ and $bias_C \in [0,1]$ are suitable constants and r_1 and $r_2 \in [0,1]$ are two random numbers. If, instead, we want to move x away from the centre and closer to y then the change in its position can be computed as $\Delta_x = -r_1 * bias_y * d_1 + r_2 * bias_C * d_2$. And so on. Computing Δx in inter-group interaction behaviours (Figure 1 (b)) requires something different. If we are moving x closer to Z, where Z can be the centre of another group, a random point in the search space or the centre of x's group itself, then the displacement is computed as $\Delta x = -d * r$, where $d = x - Z$ and $r \in [0.95, 1.05]$ is a random number.

The group fitness function is a linear combination of three values which represent three different aspects of group quality: the ranking, the size, and the volume of the space occupied by the group. Formally the group fitness is defined as follows:

$$F(g) = \frac{|G_t| - Rank(g)}{|G_t|} + SizeFitness\left(Size_g, \frac{|X|}{|G_t|}\right) + \frac{Volume_t}{\frac{Volume_{pop}}{|G_t|}}$$

where g is a group and *Rank(g)* is a function that gives the ranking of g among other groups. For the purpose of ranking, groups are sorted in descending order. The sorting is based on the value of the expression *0.75*BestFitness+0.25*AverageFitness*. The top group's rank will be 0. The *SizeFitness* is as follows:

$$SizeFitness(S, Max_S) = \begin{cases} \frac{S}{Max_S} & if\ S < Max_S \\ 1 - \frac{S - Max_S}{Max_S} & otherwise \end{cases}$$

The output of this function increases as the value of S increases until it is greater than the value Max_S, beyond which the output starts to decrease. This function rewards groups with the "right size", bigger or smaller sizes leading to less group fitness. The volume of the group, $Volume_t$ is the volume of a *dim*-dimensional sphere, the radius of which is computed as one half the diameter of the group (i.e., the distance between the two individuals further apart in the group). $Volume_{pop}$ is the volume of the search space (typically a multi-dimensional box).

The process of evolving behaviours tries to find the right combination of intra- and inter-group behaviours to put the groups in some state of dynamic equilibrium. The evolution process updates the behaviours to provide a group with the required operators to cope with different environmental changes, including changes that are caused by other groups as they compete or cooperate. After a round of interactions, the procedure that evolves behaviours works out how each behaviour has influenced the relative fitness (group fitness), so we can apportion blame and credit. For intra-group behaviours of group g, the effect rate is computed as follows:

$$e_{\alpha_t}(b) = \frac{\omega_t(b)}{\Omega_t(g)}\left(w_1(b)\frac{Size_t - Size_{t-1}}{Size_{t-1}} + w_2(b)\frac{A_t - A_{t-1}}{A_{t-1}} + w_3(b)\frac{Volume_t - Volume_{t-1}}{Volume_{t-1}}\right)$$

where $b \in B_{intra}$, $\omega_t(b)$ is the number of occurrences of behaviour b and $\Omega_t(g)$ is the total number of behaviours that caused changes to the group, by interaction behaviours initiated by group members or by members of other groups. A_t is the average fitness of group members at time t. The values of weighting parameters $w_i(b)$ are shown in Table 1

Table 1: The values of weighting parameters $w_i(b)$

Behaviours (b)	$w_1(b)$	$w_2(b)$	$w_3(b)$
Cooperative	0	0.5	-0.5
Selfish	0	0.5	-0.5
Spiteful	-0.33	-0.33	0.33
Altruistic	-0.33	-0.33	0.33

The effect rates of inter-group behaviours B_{inter} are given by:

$$e_{\beta_t}(Cooperative) = 1 - \left|\frac{|G_t| - |G_{t-1}|}{|G_{t-1}|}\right| + 1 - \left|\frac{Size_t - Size_{t-1}}{Size_{t-1}}\right|$$

$$e_{\beta_t}(Selfish) = 1 - \left|\frac{|G_t| - |G_{t-1}|}{|G_{t-1}|}\right| + \frac{Size_t - Size_{t-1}}{Size_{t-1}}$$

$$e_{\beta_t}(Spiteful) = \frac{|G_t| - |G_{t-1}|}{|G_{t-1}|} - \frac{Size_t - Size_{t-1}}{Size_{t-1}}$$

$$e_{\beta_t}(Altruistic) = 1 - \left|\frac{|G_t| - |G_{t-1}|}{|G_{t-1}|}\right| - \frac{Size_t - Size_{t-1}}{Size_{t-1}}$$

After computing the effect rates and the group's fitness, we can update the behaviour distributions of the group and prepare for next round of interactions. First we find the behaviours *pay-off* $u(b) = F(g)e_{\theta_t}(b)$ where $b \in B$, and B could be either B_{intra} or B_{inter}. The average of the pay-off of the two (intra- and inter-group) mixed behaviours is as follows:

$$U(\alpha_t) = \sum_{b \in B_{intra}} u(b)\alpha_t(b) \text{ and } U(\beta_t) = \sum_{b \in B_{inter}} u(b)\beta_t(b)$$

Then we use the *replicator equation* [11] to find the new distributions of group behaviours. Namely,

$$\dot{\alpha}_t(b) = \alpha_t(b)\big(u(b) - U(\alpha_t)\big)$$

$$\alpha_{t+1}(b) = \dot{\alpha}_t(b) + \alpha_t(b) \text{ for } b \in B_{intra}$$

$$\dot{\beta}_t(b) = \beta_t(b)\big(u(b) - U(\beta_t)\big)$$

$$\beta_{t+1}(b) = \dot{\beta}_t(b) + \beta_t(b) \text{ for } b \in B_{inter}$$

Figure 2 shows the pseudo-code of our social adaptive groups evolutionary system. It is worth mentioning that all random numbers are generated uniformly. The relative frequency of inter- and intra-group interactions is an important parameter of the algorithm that needs to be correctly set.

```
t=0
Generate an initial random population X
Evaluate population individual fitnesses
Form groups set G_t
Initialise behaviour distributions α_t and β_t for all g ∈ G_t
Repeat
  Repeat //interaction round
    Randomly select between Inter- or Intra- group interaction
    If intra-group interaction then
      Randomly select x and y from a group g ∈ G_t
      Randomly select b from B_intra according to α_t of g
    Else //inter-group interaction
      Randomly select x ∈ g_1 and y ∈ g_2 where g_1 ≠ g_2 and g_1, g_2 ∈ G_t
      Randomly select b from B_inter according to β_t of g_1
    End if
    Compute (x', y') = b(x, y)
    Replace x and y with x' and y', respectively
  Until maximum number of interactions per iteration
  t=t+1
  Evaluate population individual fitnesses
  Form groups set G_t
  Update behaviour distributions α_t and β_t for all g ∈ G_t
Until t reaches maximum number of iterations
```

Figure 2: Pseudo-code for proposed evolutionary system

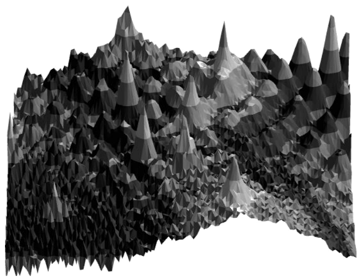

a. Fitness landscape

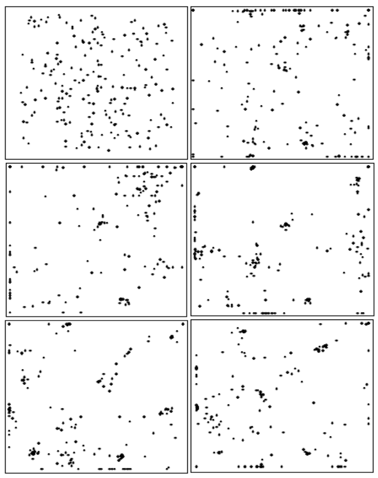

b. Individual distribution in the search space

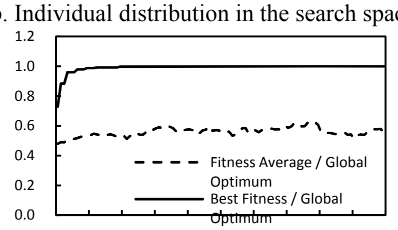

c. Performance of the proposed algorithm

Figure 3: Performance in a 2-dimensional environment

3. EXPERIMENTAL RESULT

In order to better study the performance of the proposed algorithm and the general behaviour and progress of groups as they move in (and, thus, explore) the fitness landscape, as mentioned above we conducted experiments in two-dimensional search spaces. The fitness landscapes were created using benchmark problem generator described in [8]. Figure 3 illustrates one such landscape and the result of a typical sample run. More specifically, Figure 3(a) shows the fitness landscape. In Figure 3(b), the distribution of individuals and the process of group formation are illustrated taking snapshots of the population at 20-iteration intervals (with the top-left panel showing the initial random population). Figure 3(c) describes the general behaviour of the algorithm from the point of view of the average of population fitness, the best-fitness-so-far in the run.

4. FUTURE WORK

We plan to use the proposed approach on higher dimension problems and use various measures of performance such as peak coverage and population diversity to analyse and make a better understanding of the proposed approach. Dynamic optimisation problems will be used and the response of the approach to dynamism will be studied. We plan also to investigate using our approach on a wider range of computational problems, such as combinatorial and discrete problems. We also plan to investigate using different population dynamics and replicators models for behaviours evolving mechanism. New features related to group will be studied. Such features can be used to make a better assessment of the behaviour effects and also can be used to improve the way we evaluate group fitnesses. Further features of groups related to increasing the perception of the group to the fitness landscape and enhancing the collective performance of the population, will also be studied and incorporated in the proposed model.

REFERENCES

[1] Blackwell, T. and Branke, J. Multi-swarm optimization in dynamic environments. *Applications of Evolutionary Computing*, *3005*. 489-500.

[2] Branke, J., Kaussler, T., Schmidt, C. and Schmeck, H. A multi-population approach to dynamic optimization problems. *Evolutionary Design and Manufacture* (2000). 299-307.

[3] Cedeno, W. and Vemuri, V.R. On the use of niching for dynamic landscapes. *Proceedings of 1997 Ieee International Conference on Evolutionary Computation (Icec '97)*. 361-366.

[4] De Jong, K.A. *Evolutionary computation : a unified approach*. MIT Press, Cambridge, 2006.

[5] Hamilton, W.D. The genetical evolution of social behaviour I & II. *Journal of Theoretical Biology*, *7* (1964). 1-52.

[6] Hofbauer, J. and Sigmund, K. *Evolutionary games and population dynamics*. Cambridge University Press, Cambridge 1998.

[7] Lahoz-Beltra, R., Ochoa, G. and Aickelin, U., Cheating for problem solving: a genetic algorithm with social interactions. in *Genetic and Evolutionary Computation Conference (GECCO 2009)*, (Montreal, Canada, 2009), ACM, 811-818.

[8] Li, C., Yang, S., Nguyen, T., Yu, E., Yao, X., Jin, Y., Beyer, H. and Suganthan, P. Benchmark generator for CEC'2009 competition on dynamic optimization, University of Leicester, University of Birmingham, Nanyang Technological University, Tech. Rep, 2008.

[9] Maynard Smith, J. and Price, G. The logic of animal conflict. *Nature*, *246* (1973). 15-18.

[10] Potter, M.A. and De Jong, K.A. Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, *8* (1). 1-29.

[11] Taylor, P. and Jonker, L. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, *40* (1978). 145-156.

[12] Thomas L. Vincent and Brown, J.S. *Evolutionary game theory, natural selection, and Darwinian dynamics* Cambridge University Press, Cambridge :, 2005.

[13] Yang, S., Ong, Y.-S. and Jin, Y. (eds.). *Evolutionary computation in dynamic and uncertain environments* Springer, Berlin 2007.