Evaluating Coevolution on a Multimodal Problem *

Tirtha Ranjeet, Philip Hingston, Chiou Peng Lam and Martin Masek
School of Computer and Security Science
Edith Cowan University
2 Bradford St, Mt Lawley, Western Australia 6050
t.ranjeet, p.hingston, c.lam, m.masek@ecu.edu.au

ABSTRACT

Competitive coevolutionary algorithms can be used to study problems in which two sides compete against each other and must choose a suitable strategy. Often these problems are multimodal. In this paper, we introduce a scalable multimodal test problem for competitive coevolution, and use it to investigate the effectiveness of some common coevolutionary algorithm enhancement techniques.

Categories and Subject Descriptors

I.2.8 [**Computing Methodologies**]: Artificial Intelligence—*search*

General Terms

Algorithms, Experimentation

Keywords

coevolution, multimodal

1. INTRODUCTION

In this paper, we focus on how to adapt a competitive coevolutionary algorithm to improve its performance when dealing with a multimodal problem.

Coevolutionary algorithms are especially suited for determining good strategies in an adversarial situation, such as games, negotiations and tactical planning. Often these problems are multimodal — there is more than one strong strategy for each side. In this paper, we introduce a scalable multimodal test problem for competitive coevolution, and use it to investigate the effectiveness of some common coevolutionary algorithm enhancement techniques.

There have been many studies testing **evolutionary** algorithms on multimodal problems (e.g. [2, 3, 7]). However, we have been unable to locate any previous work on multimodal test problems for competitive **coevolution**.

2. A MULTIMODAL TEST PROBLEM

We introduce an *n-peaks* problem - there are n equally good strategies for each side. The challenge for a coevolutionary algorithm is to locate these peaks.

*

GECCO'12 Companion, July 7–11, 2012, Philadelphia, PA, USA.
ACM 978-1-4503-1178-6/12/07.

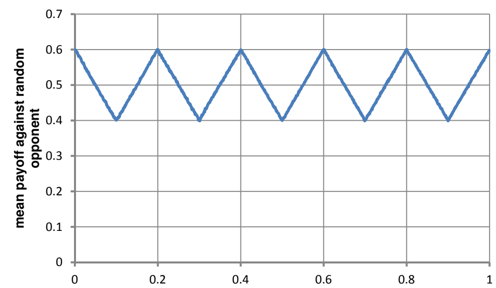

Figure 1: Mean payoffs against random opponents for solutions to the 5-peaks problem with H=L=1.

We picture the domain as divided into n equal intervals. Parameters H and L control the heights of the peaks and troughs. Figure 1 illustrates the case $n = 5$, $H = L = 1$. When x and y compete, the outcome depends on which intervals they belong to, and on the distances from the centres of their intervals. If x and y are in the same interval, then x gets a payoff of H if it is *further from* the centre of the interval than y is (otherwise 0). If y is in the next interval to the right of x, then x gets a payoff of L if it is *closer to* the centre of its interval than y is to the centre of its interval. For this purpose, the "next interval to the right" of the rightmost interval is considered to be the leftmost interval - i.e. the domain wraps around. If y is two intervals to the right of x, then x gets a payoff if it is furthest from the centre of its interval. This pattern continues, with wrapping if necessary, so that the domain is actually circular, rather than linear. If y is an even number of intervals to the right of x, then it is good for x to be nearer its boundary (payoff $= H$), while if y is an odd number of intervals to the right, then it is good for x to be nearer the centre of its interval (payoff $= L$).

3. ALGORITHM AND MEASURES

In order to illustrate the difficulties posed by multimodality, we carried out experiments to test the performance of a simple competitive coevolutionary algorithm, along with some popular variations, on an *n-peaks* problem. In this section we describe the algorithm and variations that we used.

As a base case, we use a simple, naïve, competitive coevolutionary algorithm which we call *CEAN*. We then define variations on *CEAN* which include a fitness sharing mechanism, or a Hall of Fame, or both, and we also vary the mutation rate. We used competitive fitness sharing [4]. This works by penalising population members that are similar to others in the population. Modifying CEAN to use this selection procedure gives an algorithm variant we call CEAFS. As an archive mechanism, we implemented a Hall of Fame (HOF) [4]. For each population, we maintain an archive, known as a Hall of Fame, consisting of fittest individuals from each earlier generation. In this CEAHOF variant of CEAN, the fitness calculation is modified to include payoffs of the individual in question against members of the opposing population *as well as* the members of the archive. Finally, we also created a fourth variant which uses both fitness sharing and a Hall of Fame, CEACFH. In this variant, fitness values are calculated using the Hall of Fame as for CEAHOF, and then these values are adjusted to obtain shared fitness values as for CEAFS.

One aspect of performance is *generalisation*: how well do solutions found for one side in a contest, learned via a coevolutionary algorithm, generalise to compete well against arbitrary strategies for the other side? Chong et al. [1] proposed a suitable set of related measures for generalisation.

In the case of a multimodal problem, another relevant aspect of performance is how well an algorithm does at locating *as many peaks as possible* – that is, can the algorithm locate many different representative solutions with high generalisation performance, rather than simply *any of them*. We used the two measures *peak ratio* and *success ratio* [5].

4. EXPERIMENTS AND RESULTS

To investigate the effects of diversity maintenance via fitness sharing and/or mutation, and of an archive (Hall of Fame), on the 5-peaks problem, we ran each algorithm 60 times, with mutation rates varying from 2.5% to 100% in steps of 2.5%. In each generation, we recorded generalisation performance, and peak finding ability (peak ratio and success ratio).

Figure 2 presents the results in a series of profile plots. Each data point is an average over 60 executions of the mean value over the final 60 generations. There is a data point for each algorithm variant and mutation rate.

5. CONCLUSION

In this paper, we have examined the performance of a competitive coevolutionary algorithm on a multimodal problem. We create the *n peaks* problem, a scalable multimodal test problem in which the number and amplitude of the peaks in the fitness landscape can be manipulated.

We then used an instance of the problem to test a naïve competitive coevolutionary algorithm, as well as variants with an archive (Hall of Fame) and a diversity maintenance mechanism (competitive fitness sharing), in terms of generalisation ability, and peak finding ability. We found that, for this problem, best results were obtained with the combination of an archive and diversity maintenance, with a moderately low level of mutation.

In future work, it remains to investigate other instances of the problem with different fitness landscapes, and in higher dimensions. In addition, other methods for handling multimodality can be tested.

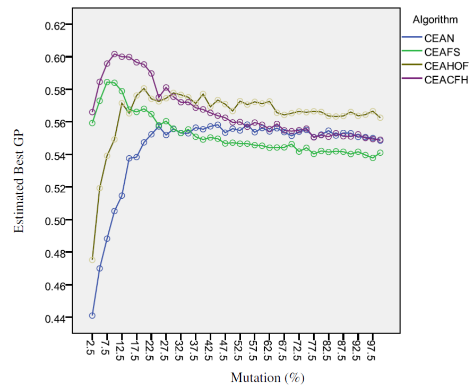

(a) generalisation

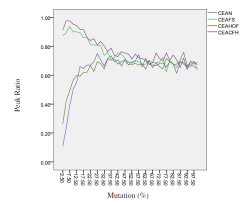

(b) peak ratio

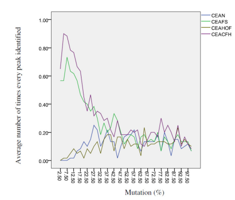

(c) success ratio

Figure 2: Generalisation and peak finding performance (mean over the final 60 generations) versus mutation rate for each algorithm variant.

6. REFERENCES

[1] S. Chong, P. Tino, and X. Yao. Measuring generalisation performance in coevolutionary learning. *IEEE Transactions on Evolutionary Computation*, 12(4):479–505.

[2] K. Deb and D. Goldberg. An investigation of niche and species formation in genetic function optimization. In *Proceedings of the Third International Conference on Genetic Algorithms*, 1989.

[3] B. Miller and M. Shaw. Genetic algorithms with dynamic niche sharing for multimodal function optimization. In *Proceedings of IEEE International Conference on Evolutionary Computation*, 1995.

[4] C. Rosin and R. Belew. New methods for competitive coevolution. *Evolutionary Computation*, 5(1):1–29.

[5] R. Thomsen. Multimodal optimization using crowding-based differential evolution. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 1382–1389, 2004.

[6] R. A.Watson and J. B. Pollack. Coevolutionary dynamics in a minimal substrate. In *Proceedings of the 2001 Genetic and Evolutionary Computation Conference(GECCO'01)*, pages 702–709, 2001.

[7] E. Yu and P. Suganthan. Ensemble of niching algorithms. *Information Sciences*, 180:2815–2833.