

A Comparison of Evaluation Methods in Coevolution

Ting-Shuo Yo
Utrecht University
PO Box 80.089
3508 TB Utrecht
The Netherlands
tyo@cs.uu.nl

Edwin D. de Jong
Utrecht University
PO Box 80.089
3508 TB Utrecht
The Netherlands
dejong@cs.uu.nl

ABSTRACT

In this research, we compare four different evaluation methods in coevolution on the Majority Function problem. The size of the problem is selected such that an evaluation against all possible test cases is feasible. Two measures are used for the comparisons, i.e., the objective fitness derived from evaluating solutions against all test cases, and the objective fitness correlation (OFC), which is defined as the correlation coefficient between subjective and objective fitness. The results of our experiments suggest that a combination of average score and weighted informativeness may provide a more accurate evaluation in coevolution. In order to confirm this difference, a series of t-tests on the preference between each pair of the evaluation methods is performed. The resulting significance is affirmative, and the tests for two quality measures show similar preference on four evaluation methods.

Categories and Subject Descriptors

F.0 [General]

General Terms

Algorithms, Experimentation, Performance

Keywords

Coevolution, evaluation, performance comparison, objective fitness correlation, OFC

1. INTRODUCTION

Coevolution offers an approach to adaptively select tests for the evaluation of learners [10, 20, 12, 19, 18, 5]. Using coevolution, the evaluation function is adapted as part of the evolutionary process. This approach can be useful if the quality of individuals can be assessed using some form of *tests*. For such *test-based problems*, the identification of an informative set of tests can reduce the amount of required

computation, while potentially providing more useful information than any static selection of tests. Since an adaptive test set can render evaluation unstable, an important question is how coevolution can be set up to be sufficiently reliable.

A recent insight in coevolution research is that the design of a coevolutionary setup should begin with a consideration of the desired *solution concept* [6]. A solution concept specifies which elements of the search space qualify as solutions and which do not. Examples of solution concepts include: Maximum Expected Utility (maximizing the expected outcome against a randomly selected opponent, which for uniform selection is equivalent to maximizing the average outcome against all opponents), the Pareto-optimal set resulting from viewing each test as a separate objective, and Nash-equilibria, in which no candidate solution or test can unilaterally deviate given the other candidate solutions and tests without decreasing its payoff.

For several of the main solution concepts used in coevolution, archive methods exist guaranteeing that when sufficiently diverse sets of new individuals are submitted to the archive, the archive will produce monotonically improving approximations of the solution concept. Recent examples of such archive methods are the Nash Memory [6, 9], which guarantees monotonicity for the Nash equilibrium solution concept; the IPCA algorithm, which guarantees monotonicity for the Pareto-optimal equivalence set [3]; and the Max-Solve algorithm [2], which guarantees monotonicity for the Maximum Expected Utility solution concept.

While theoretical guarantees of monotonic progress are important, so far no bounds or guarantees regarding the improvement of the approximation to the solution concept over time are available. Thus, approximating the solution concept to a desired degree of accuracy may take an infeasible amount of time. An important current practical question is therefore: how can coevolutionary algorithms be set up such that their dynamics lead to quick improvement over time? By using such efficient algorithms as generators of new individuals and coupling them to monotonic archives, thereby combining the guarantee of monotonic progress with efficiency, a principled approach to designing robust and efficient coevolution algorithms is obtained.

Efficiency in coevolutionary algorithms depends on selection, see e.g. [7], and evaluation. In this research we focus on evaluation. Our aim is to compare the efficiency, reflected in the improvement over time of an objective quality measure, that can be achieved using different coevolutionary evaluation methods. Since a main question is how sufficiently

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'07, July 7–11, 2007, London, England, United Kingdom.

Copyright 2007 ACM 978-1-59593-697-4/07/0007 ...\$5.00.

accurate evaluation may be achieved, the testing environment is chosen such that evaluating individuals on all tests is feasible; while this is not the case in practical applications of coevolution, this provides a possibility to compare evaluation methods with the maximally informative situation in which information about all possible tests is available. This setup permits investigating two important questions:

1. Given all information that may be relevant to evaluation, how can this information be used optimally?
2. Compared to evaluation based on all relevant information, how do different coevolutionary evaluation methods perform?

In this paper, we focus on the second question. Four different coevolutionary evaluation methods are compared to each other and to the baseline of testing against all tests. The test problem is a small variant of the Majority Function test problem [17] [16] [12] chosen such that evaluation against all test cases (initial conditions) is feasible. A new tool named the Objective Fitness Correlation (OFC) [4], the correlation between the subjective and the objective fitness measures, is used to assess the evaluation accuracy of the different methods.

The paper is structured as follows. In section 2 we discuss the evaluation methods and algorithms used in this research. The design of experiments, parameters and performance measures are described in section 3. The results are presented in section 4, and the discussions and concluding remarks are shown in section 5.

2. EVALUATION METHODS

In this section, we describe the evaluation methods conducted in this study. Since we put our main focus on the test-based problems, we start with defining some terminologies and fitness measures. Evaluation methods based on those fitness measures are introduced in the following sections.

2.1 Interaction matrix

For a test problem, there are two sample spaces: one is the test cases, T , and the other is possible solutions, S . The term *interaction* is defined as the result of letting one solution interact with one test case. For the ease of analysis, an *interaction function*, $G(T, S)$, is designed to return one scalar outcome of the interaction between a pair of test case and a solution. To simplify the following discussion, we assume that the interaction function returns a binary result which represents whether the solution succeeds in solving the test case or not. An affirmative interaction is preferred by the solution, but is unfavorable to the test case.

By evaluating all solutions, $\mathbf{S} = \{S_j | j = 1, 2, \dots, m\}$, against all test points, $\mathbf{T} = \{T_i | i = 1, 2, \dots, n\}$, an interaction matrix, $I = \{I_{ij} | i = 1, \dots, n; j = 1, \dots, m\}$, is obtained. The sum of column j , I_j , of this interaction matrix represents the number of test cases that S_j has solved, which is a common performance measure for a solution. Similarly, the sum of row i , I_i , is the number of times this test case T_i being solved, which represents the difficulty of the test case.

In the simplest case, only the interaction outcomes are considered, so that $-I_i$ can be used as the fitness of T_i (i.e., the more difficult the better), and I_j for S_j (i.e., the more powerful the better).

2.2 Distinctions

Distinctions are defined as the ability to distinguish between good and bad solutions and may provide important information for selecting a proper set of test cases. This concept was first proposed in [8], and here we follow the notation in [5] for our experiments. Accordingly, a three-dimensional matrix $dist(T, S, S)$ is defined to represent "if test case T_i distinguishes S_k from S_l ", or mathematically:

$$dist(T_i, S_k, S_l) \Leftrightarrow G(T_i, S_k) > G(T_i, S_l) \quad (1)$$

where $T_i \in \mathbf{T}$ and $S_k, S_l \in \mathbf{S}$.

The number of distinctions that test case T_i has made may be represented by summing up all m^2 solution pairs in $dist$, i.e., $dist_T(T_i) = \sum_{kl} dist(i, k, l)$. This value represents the ability of a test case to maintain diversity of the solutions, and it may also be considered as a fitness measure for test cases.

Although we explain the concept of distinction in terms of "selecting a proper set of test cases", the same procedure can also be applied to solutions. This yields the distinction of solutions, and is also a reasonable fitness measure for evaluating solutions. With those measures, four different evaluation methods are defined as follows.

2.3 Four evaluation methods

For each evaluation method, we define the fitness for test cases $F(T)$ and solutions $F(S)$ as follows.

2.3.1 Average score

For a solution, the average-score is defined as "the proportion of test points it has succeeded in solving." And for a test point, this value is represented by "the proportion of solutions that it has failed". With an interaction matrix $I(i, j)$ as described earlier, we define:

$$\begin{aligned} F_{AS}(T_i) &= 1 - \sum_{j=1}^m I(i, j)/m, \\ F_{AS}(S_j) &= \sum_{i=1}^n I(i, j)/n \end{aligned} \quad (2)$$

The application results of F_{AS} on test cases and solutions are both between $[0, 1]$, and a higher value represents a better performance.

2.3.2 Weighted score

Instead of giving each interaction an equal weight, the average scores described above may be used as weights for each test case and solution. That is to say, a test case earns more credits by failing a more powerful solution, and a solution gets a higher score when it succeeds in a tough test case. The mathematical expression of this evaluation can be

represented as follows:

$$\begin{aligned}
W_T(T_i) &= \left(\sum_{j=1}^m I(i, j) \right)^{-1}, \\
W_S(S_j) &= \left(\sum_{i=1}^n I(i, j) \right)^{-1}, \\
W'_T(T_i) &= \frac{W_T(T_i)}{\sum_{i=1}^n W_T(T_i)}, \\
W'_S(S_j) &= \frac{W_S(S_j)}{\sum_{j=1}^m W_S(S_j)}, \\
F_{WS}(T_i) &= 1 - \sum_{j=1}^m W'_S(S_j) \cdot I(i, j), \\
F_{WS}(S_j) &= \sum_{i=1}^n W'_T(T_i) \cdot I(i, j) \quad (3)
\end{aligned}$$

Weights are checked to avoid being divided by zero and normalized to ensure the weighted scores are ranged from 0 to 1. Hence F_{WS} also returns a value between $[0, 1]$, and the higher values are more preferable. This definition is similar in spirit to the niching methods of Rosin [20] and Juille [11], though their methods are not considered while we developed the weight function.

2.3.3 Average informativeness

The informativeness of a test measures the amount of information it provides about a given set of candidate solutions. In [1], a definition for informativeness based on the incomparable and equal elements of the order induced by a test is provided. Here, we measure the informativeness of a test based on Ficici's notion of distinctions [8]. Since each distinction a test makes contributes to its informativeness, and since the set of all possible distinctions is sufficient to provide ideal evaluation [5], we measure the informativeness of a test as the normalized number of distinctions it makes. In this study, we define the *distinction score*, D , as "the number of distinctions one test case (solution) makes." This value can be derived from the distinction matrix described in (1), and represents the informativeness. For the convenience of further computation, we normalize the distinction score with its maximum and minimum values, which are returned by the function $Max(D)$ and $Min(D)$. Because the informativeness represents the individual's ability to maintain diversity of the other population, we want to integrate it with the average scores to form the fitness. Here we simply use a linear combination of two scores, with weights of (0.3, 0.7) for distinction and average score, respectively. These weights are based on the observation that the average scores are generally lower than the distinction scores due to the computational scheme we used, as well as experiences from the pilot experiments. The average informativeness is therefore defined as the linear combination of the normalized distinction score and the average score.

$$\begin{aligned}
D_T(T_i) &= \sum_{k,l \in [1,m]} dist(T_i, S_k, S_l), \\
D_S(S_j) &= \sum_{p,q \in [1,n]} dist(S_j, T_p, T_q), \\
D'_T(T_i) &= \frac{D_T(T_i) - Min(D_T(T_i))}{Max(D_T(T_i)) - Min(D_T(T_i))}, \\
D'_S(S_j) &= \frac{D_S(S_j) - Min(D_S(S_j))}{Max(D_S(S_j)) - Min(D_S(S_j))}, \\
F_{AI}(T_i) &= 0.3 \cdot D'_T(T_i) + 0.7 \cdot F_{AS}(T_i), \\
F_{AI}(S_j) &= 0.3 \cdot D'_S(S_j) + 0.7 \cdot F_{AS}(S_j) \quad (4)
\end{aligned}$$

2.3.4 Weighted informativeness

Similar to the weighted score, each distinction that has been made can be weighted differently when the distinction score is derived. For example, the distinction made by all test cases, $\sum_{i=1}^n dist(T_i, S_k, S_l)$, provides the information of "how many test cases have made the distinction on S_k and S_l , such that $G(T_i, S_k) > G(T_i, S_l)$." The inverse of this value can be used as the weight of this particular distinction, so that when more test cases can make this distinction, the less worthy this distinction is. This operation can also be applied to solutions, and the weighted informativeness can be defined mathematically as:

$$\begin{aligned}
W_d(S_k, S_l) &= \left(\sum_{i=1}^n dist(T_i, S_k, S_l) \right)^{-1}, \\
W_d(T_p, T_q) &= \left(\sum_{j=1}^m dist(S_j, T_p, T_q) \right)^{-1}, \\
D_T(T_i) &= \sum_{k,l \in [1,m]} W_d(S_k, S_l) \cdot dist(T_i, S_k, S_l), \\
D_S(S_j) &= \sum_{p,q \in [1,n]} W_d(T_p, T_q) \cdot dist(S_j, T_p, T_q), \\
D'_T(T_i) &= \frac{D_T(T_i) - Min(D_T(T_i))}{Max(D_T(T_i)) - Min(D_T(T_i))}, \\
D'_S(S_j) &= \frac{D_S(S_j) - Min(D_S(S_j))}{Max(D_S(S_j)) - Min(D_S(S_j))}, \\
F_{WI}(T_i) &= 0.3 \cdot D'_T(T_i) + 0.7 \cdot F_{AS}(T_i), \\
F_{WI}(S_j) &= 0.3 \cdot D'_S(S_j) + 0.7 \cdot F_{AS}(S_j) \quad (5)
\end{aligned}$$

In the following discussion, the four evaluation methods described above are referred as **AS**, **WS**, **AI**, and **WI**, respectively.

2.4 Algorithms for experiments

There are two main algorithms used in this study, i.e., a single population genetic algorithm (GA) and coevolution (CO). The former uses only one single population for solutions and evaluates the population on all possible test cases to obtain the interaction defined earlier. Afterward, the average scores given by F_{AS} are used as the fitness in the GA. The results of this algorithm are used as the baseline for comparisons. The second algorithm uses coevolution between test cases and solutions, and all four evaluation methods are tested.

Algorithm 1 and 2 describe GA and CO, respectively.

Functions used in the algorithms, e.g., INTERACTION, EVALUATE, SELECT and BREED, are specific to problems and experiments. The general concepts of INTERACTION and four implementations of EVALUATE are already discussed in this section. The SELECT and BREED together define the selection, reproduction and replacement cycle in a generation. The choice of these two functions is problem specific and they are specified in the experimental design section.

Algorithm 1 Single Population GA

```

TC ← all  $n$  test cases
Sol(0) ← randomly generate  $m$  solutions
while  $t < \text{MAX\_GENERATION}$  do
  I ← INTERACTION(TC, Sol( $t$ ))
  Sol( $t$ ) ← EVALUATE Sol( $t$ ) with  $F_{AS}$ 
  Sol'( $t$ ) ← SELECT Sol( $t$ )
  Sol( $t + 1$ ) ← BREED Sol'( $t$ )
   $t = t + 1$ 
end while
return Sol( $t$ )

```

Algorithm 2 Coevolution

```

TC(0) ← randomly generate  $n$  test cases
Sol(0) ← randomly generate  $m$  solutions
while  $t < \text{MAX\_GENERATION}$  do
  I ← INTERACTION(TC( $t$ ), Sol( $t$ ))
  TC( $t$ ) ← EVALUATE TP( $t$ )
  Sol( $t$ ) ← EVALUATE Sol( $t$ )
  TC'( $t$ ) ← SELECT TC( $t$ )
  Sol'( $t$ ) ← SELECT Sol( $t$ )
  TC( $t + 1$ ) ← BREED TC'( $t$ )
  Sol( $t + 1$ ) ← BREED Sol'( $t$ )
   $t = t + 1$ 
end while
return (TP( $t$ ), Sol( $t$ ))

```

3. EXPERIMENTAL SETUP

In this section, the settings of our experiments are described. The software used in this research is implemented as an extension of ECJ [15], which is developed by George Mason University’s Evolutionary Computation Laboratory. All simulations use the basic evolutionary loop provided by ECJ, plus the evaluation methods and problem-specific functions in the extension. The design of the experiments, the parameters used in the test problem, and the performance measures are introduced as follows.

3.1 Design of experiments

As mentioned in the previous section, two algorithms and four evaluation methods are used in this study. Table 1 shows the design of our experiments. Experiment 1 is a combination of the single population genetic algorithm and the average score evaluation method. This GAAS evaluates the solutions against all possible test cases and serves as the baseline experiment. Experiment 2, 3, 4, and 5 are four different evaluation methods combined with the coevolution algorithm. Each experiment is run for ten times with ten different random seeds.

Exp No	1	2	3	4	5
Algorithm	GA	CO	CO	CO	CO
Evaluation Method	AS	AS	WS	AI	WI
Number of Runs	10	10	10	10	10

Table 1: Design of experiments.

Parameter	Value
MAX_GENERATION	200
Size of population	64/100
Elitism	12/20 (20%)
Selection	linear rank selection
Crossover	one point crossover
Mutation rate	0.01

Table 2: Parameters used for the majority function problem.

3.2 Parameters for the test problem

In this study, we perform the experiments on the majority function problem. This problem is also known as the one dimensional cellular automata problem. Mitchell and his colleagues have detailed discussion on this problem in [17] and [16]. Two major parameters are used for this problem, i.e., the radius of the neighbourhood (r) and the size of the one dimensional lattice (N). Boolean vectors are used to represent both the initial conditions (test cases) and the rules (solutions). In order to evaluate against all test points, we chose $r = 2$ and $N = 9$. Other parameters, e.g., the size of populations, elitism, crossover type, and mutation probability are selected based on Mitchell’s work [16], and are summarized in Table 2.

In the coevolution experiments (2, 3, 4, and 5), a symmetric setup is used, i.e., the same set of EVALUATE, SELECT and BREED functions are used for both population of test cases and the population of solutions. However, there is one exception: the population size. Because there are only 512 possible test cases in total, the population size for test cases is set as 64 in the convenience of comparison, while the population size for solutions is set as 100. Accordingly, the number of elites for two populations are also different, while their proportions to the size of populations are the same (20%). The initial populations are both randomly created and are both checked to ensure there are no redundant individuals in the beginning of the experiments.

3.3 Performance measures

In order to compare all experiments fairly, an objective performance measure is required. Hence we define the *subjective fitness* and *objective fitness* separately. The *subjective fitness* is the fitness value returned by the EVALUATE function used in the coevolution algorithm. In this study, the *objective fitness* is defined as the fitness used in experiment 1, i.e., to evaluate the average scores against all possible test cases. For each generation in experiment 2 ~ 5, both the subjective and the objective fitness are recorded. All experiments are compared with the their objective fitness, based on the number of interactions. In coevolution experiments, the number of test cases is 1/8 of that in the GAAS experiment, therefore during the analysis we compare the




Figure 1: The best objective fitness of solutions as a function of the number of interactions

objective fitness of the coevolution experiments every 8 generations to the fitness of the controlled experiments every generation. Since these values are all derived from exhaustive testing, the progress over interactions may be seen as the performance of each method.

In addition to the progress of the fitness, the correlation coefficients between subjective and objective fitness are also computed. This correlation is defined as the **Objective Fitness Correlation**, OFC, in [4] as a new objective measure for evaluating coevolutionary algorithms. In our experiments, this measure is collected for each generation of every run. Since OFC is always 1 in the baseline experiment, we only compare OFC among the four evaluation methods with coevolution. Comparisons are made with these two measures averaged upon ten independent runs.

4. RESULTS

4.1 Best objective fitness

Figure 1 shows the best objective fitness of the solutions averaged over 10 runs. As illustrated in the figure, the baseline experiment is outperformed by all coevolution experiments, especially when the number of total interactions is small. Among four evaluation methods with coevolution, the weighted informativeness shows a constantly better performance than others.

A series of paired, one-tailed t-tests are performed to examine the significance of the differences between each evaluation methods. The best fitness of each generation is averaged over 10 runs, and is paired with the same quantity of another evaluation method for the t-test. The results of t-tests shows significance on $WI > AI$, $WI > WS$, $WI > AS$, $AI > AS$, and $WS > AS$ (with p-values $< 1 \times 10^{-10}$), but not on $AI \neq WS$ with $\alpha = 0.05$. These results are summarized in table 3.

4.2 Objective fitness correlation

The OFCs for different evaluation methods with coevolution are shown in figure 2. The values shown in the figure are averages of 10 independent runs, and therefore the fluctuations of OFCs are already smoothed. In a single run of one experiment, the OFC can be negative for some generations.




Figure 2: The OFC for experiment COAS, COWS, COAI and COWI.

P-value for test on $i < j$	WS	AI	WI
	AS	1.4×10^{-21}	9.9×10^{-15}
WS		9.8×10^{-1}	7.2×10^{-81}
AI			1.5×10^{-64}

Table 3: Significant level in paired t-tests for objective fitness among experiment COAS, COWS, COAI and COWI. A p-value smaller than 0.05 is usually considered as significant.

In all experiments, the OFC always starts from a high value. This is because the initial populations are randomly generated, and this can be seen as they are sampled from the set of all possible cases with a uniform distribution. Ideally, this random sampling may create a set well represents the original search space, i.e., the all possible cases, and results in a higher correlation between the subjective and the objective fitness. As the coevolution proceeds, the populations move toward certain direction rather than a uniformly random sampling, and hence the OFC decreases over generations. However, as argued in [4], the OFC may still remain as a performance measure for comparing different evaluation methods in a coevolutionary algorithm.

As demonstrated in figure 2, the weighted informativeness shows a constantly higher OFC through generations than other methods, while other evaluation methods do not show clear differences between one another due to the fluctuations.

A set of t-tests is employed to verify the difference in OFC between each pair of evaluation methods. The results show significance on all of following relations: $WI > AI$, $WI > WS$, $WI > AS$, $AI > WS$, $AI > AS$, and $WS > AS$. The preferences for WI against other methods are very significant, with p-values $< 1 \times 10^{-15}$, which is the limit of precision in the computing software. These results are summarized in table 4.

P-value for test				
on $i < j$	WS	AI	WI	
AS	4.2×10^{-5}	6.5×10^{-21}	2.9×10^{-69}	
WS		4.6×10^{-9}	3.4×10^{-59}	
AI			1.1×10^{-39}	

Table 4: Significant level in paired t-tests for OFC among experiment COAS, COWS, COAI and COWI.

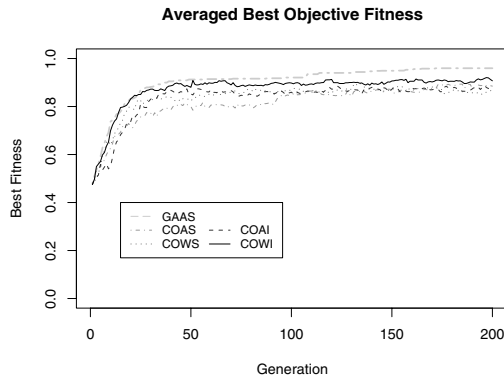


Figure 3: The objective fitness over generations for experiment GAAS, COAS, COWS, COAI and COWI.

5. DISCUSSION

5.1 Evaluating with the weighted informativeness

The results of our experiments suggest that in coevolution, a combination of performance (average score) and diversity (weighted distinction) can achieve the accuracy of full evaluation with less computational cost. This statement holds for both conditions when considering the improvement of an objective quality measure over time or the correlation between the subjective fitness and objective measure. This advantage is even clearer if we compare the objective fitness according to the generation instead of the number of interactions. As shown in figure 3, the WI evaluation method with coevolution progresses as fast as evaluating against all possible test cases, while other evaluation methods are apparently slower than the baseline.

In our study, the OFC is used as a quality measure in addition to the objective fitness. If the preference orderings of evaluation methods are considered, this measure provides information similar to the objective fitness. However, the difference between WS and AI shows significance in OFC but not in the objective fitness. This implies that two measures may contain different information, and a detailed discussion on OFC can be found in [4].

5.2 The majority function problem with different sizes

This research is the first time the OFC is calculated on a real problem, but in exchange we are only able to experiment on problems with a small number of possible test cases in

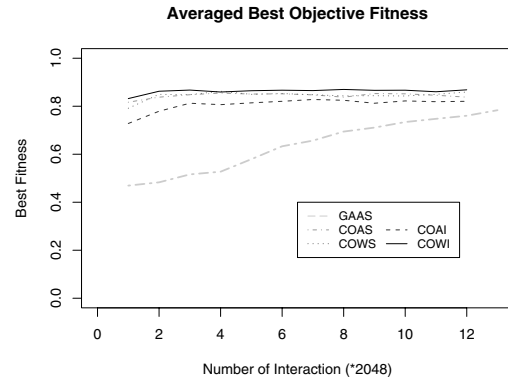


Figure 4: The averaged best objective fitness for the $N = 11$ majority function problem.




Figure 5: The averaged OFC for the $N = 11$ majority function problem. The upper lines are AI and WI, and the lower lines are AS and WS.

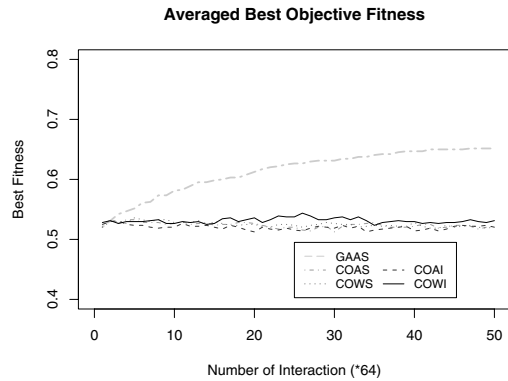


Figure 6: The averaged objective fitness for the 6-odd-parity problem.

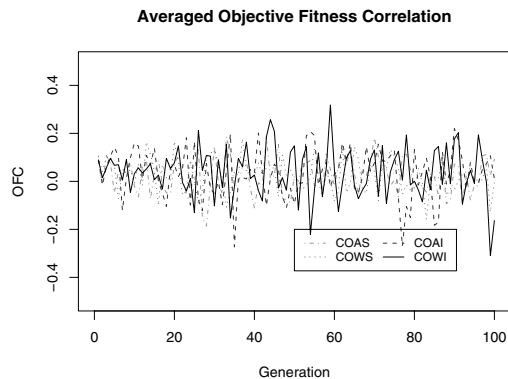


Figure 7: The averaged OFC for the 6-odd-parity problem.

total. The majority function problem with different sizes ($N = 7$ and $N = 11$) are also tested. The results are similar to the finding presented in the previous section, and the experiments of $N = 11$ are summarized in figure 4 and 5. From these figures it is shown that OFC can distinguish AI and WI from AS and WS while the objective fitness can not. This argument is consistent with the results of t-tests (not shown): t-tests on the objective fitness show no significant difference between pairs among AS, WS, and AI, while those on OFC show the same preference as in the $N = 9$ experiments.

5.3 The parity problem

In addition to the majority function problem, the same comparison is also performed on the n -odd-parity problems. Although Lee, Xu, and Chau [14] have proposed that the boolean parity problem may be transformed into a cellular automata problem, here we use Koza’s Genetic Programming approach in [13].

Figure 6 and 7 show the results for the 6-odd-parity problem. As shown in figure 6, the WI still performs the best among four evaluation methods in coevolution, and this is also confirmed by the t-tests (not shown). However, we also

find some disagreement between results from the majority function problem and the parity problem. First, in the parity problem, the single population GA outperforms all CO experiments in this genetic programming approach, while figure 1 shows the opposite in the majority function problem. Second, OFC in the parity problem shows no significant pattern (as shown in figure 7), while the OFC distinguishes different evaluation methods better than the objective fitness in the majority function problem. Finally, the t-tests hardly show any significant difference among AS, WS, and AI in the parity problem.

A detailed analysis on the output has been done and suggests a few possible reasons for the disagreement. First, solving the parity problem with the genetic programming approach is not a symmetric test-based problem in nature. That is to say, while the test cases are represented with boolean vectors as they are in the majority function problem, the solutions for the parity problem are represented with tree-like structures. As a result of this asymmetry, the same evolutionary operator (e.g., selection methods, mutation rate, crossover methods, etc.) may have different effects on two populations, and hence a symmetric setting for coevolution is not suitable. Second, the initial population of the tree-like solutions is not randomly selected from “all possible solutions”. In Koza’s approach, the sizes of the solutions start from smaller values and grow over generations. This “biased sampling” may explain the disfavor of OFC in the parity problem. Finally, since Koza’s approach has been studied for several years, the evolutionary operators and parameters are already well tuned. We believe that the disagreement may be reduced by developing proper asymmetric coevolution schemes.

Despite the disagreement, the results of the parity problem still shows a favor to WI in coevolution.

6. CONCLUSION

In this research, we compare four different evaluation methods in coevolution on a test-based problem. Two measures are used for the comparisons among average score (AS), weighted score (WS), average informativeness (AI), and weighted informativeness (WI). In addition to an objective quality measure, the objective fitness correlation (OFC) is also computed.

The experimental results show a strong preference on WI, which suggest that a combination of the performance and the ability to create distinctions may provide more accurate evaluation in coevolution. The resulting significance from t-tests show a similar preference when two quality measures are used, separately. This study also uses the recently proposed OFC to evaluate the accuracy of coevolutionary evaluation methods on a concrete test problem.

Although we have shown the advantages of using WI in coevolution, the way we combine these two measures is simply using a weighted summation. It may worth exploring more sophisticated methods to fuse this information together. Currently a multi-objective approach is in progress, and this may lead to a more detailed investigation on how to use both measures in coevolutionary algorithms.

7. REFERENCES

- [1] A. Bucci and J. B. Pollack. A mathematical framework for the study of coevolution. In

- Foundations of Genetic Algorithms (FOGA-2002)*, San Francisco, CA, 2003. Morgan Kaufmann.
- [2] E. D. De Jong. The MaxSolve algorithm for coevolution. In H.-G. Beyer, editor, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-05*, pages 483–489. ACM Press, 2005.
- [3] E. D. De Jong. A monotonic archive for pareto-coevolution. *Evolutionary Computation*, 15(1), 2007. to appear.
- [4] E. D. De Jong. Objective fitness correlation. In Dirk Thierens et al., editor, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-07*. ACM Press, 2007.
- [5] E. D. De Jong and J. B. Pollack. Ideal evaluation from coevolution. *Evol. Comput.*, 12(2):159–192, 2004.
- [6] S. G. Ficici. *Solution Concepts in Coevolutionary Algorithms*. PhD thesis, Brandeis University, 2004.
- [7] S. G. Ficici, O. Melnik, and J. B. Pollack. A game-theoretic and dynamical-systems analysis of selection methods in coevolution. *IEEE Transactions on Evolutionary Computation*, 9(6):580–602, 2005.
- [8] S. G. Ficici and J. B. Pollack. Pareto optimality in coevolutionary learning. In *ECAL '01: Proceedings of the 6th European Conference on Advances in Artificial Life*, pages 316–325, London, UK, 2001. Springer-Verlag.
- [9] S. G. Ficici and J. B. Pollack. A game-theoretic memory mechanism for coevolution. In Cantú-Paz, E., et al., editor, *Genetic and Evolutionary Computation – GECCO-2003*, volume 2723 of *LNCS*, pages 286–297, Chicago, 12-16 July 2003. Springer-Verlag.
- [10] D. W. Hillis. Co-evolving parasites improve simulated evolution in an optimization procedure. *Physica D*, 42:228–234, 1990.
- [11] H. Juillé and J. Pollack. Dynamics of co-evolutionary learning. In *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 526–534. MIT Press, 1996.
- [12] H. Juillé and J. B. Pollack. Coevolving the ”ideal” trainer: Application to the discovery of cellular automata rules. In *Proceedings of the Third Annual Genetic Programming Conference*, 1998.
- [13] J. R. Koza. *Genetic programming: on the programming of computers by means of natural selection*. MIT Press, Cambridge, MA, USA, 1992.
- [14] K. M. Lee, H. Xu, and H. F. Chau. Parity problem with a cellular automaton solution. *Phys. Rev. E*, 64(2):026702, Jul 2001.
- [15] S. Luke. ECJ 15: A Java evolutionary computation library. <http://cs.gmu.edu/~eclab/projects/ecj/>, 2006.
- [16] M. Mitchell, J. P. Crutchfield, and P. T. Hraber. Evolving cellular automata to perform computations: mechanisms and impediments. *Phys. D*, 75(1-3):361–391, 1994.
- [17] M. Mitchell, P. T. Hraber, and J. P. Crutchfield. Revisiting the edge of chaos: Evolving cellular automata to perform computations. *Complex Systems*, 7:89–130, 1993.
- [18] L. Pagie and P. Hogeweg. Evolutionary consequences of coevolving targets. *Evolutionary Computation*, 5(4):401–418, 1998.
- [19] J. Paredis. Coevolutionary computation. *Artificial Life*, 2(4), 1996.
- [20] C. D. Rosin and R. K. Belew. New methods for competitive coevolution. *Evolutionary Computation*, 5(1):1–29, 1997.