Behavioural GP Diversity for Dynamic Environments: an application in hedge fund investment

Real-World Applications Track

Wei Yan
Department of Computer Science
University College London
Gower Street, London, WC1E 6BT
w.yan@cs.ucl.ac.uk

Christopher D. Clack
Financial Computing and Quantitative Finance
Department of Computer Science
University College London
Gower Street, London, WC1E 6BT
c.clack@cs.ucl.ac.uk

ABSTRACT

We present a new mechanism for preserving phenotypic behavioural diversity in a Genetic Programming application for hedge fund portfolio optimization, and provide experimental results on real-world data that indicate the importance of phenotypic behavioural diversity both in achieving higher fitness and in improving the adaptability of the GP population for continuous learning.

Categories and Subject Descriptors

I.2.M [**Artificial Intelligence**]: Miscellaneous

General Terms

Algorithms, Experimentation

Keywords

Genetic Programming, Diversity, Phenotype, Finance, Adaptation, Dynamic Environment

1. INTRODUCTION

In the context of our research into the applicability of Genetic Programming (GP) technology to the optimization of Hedge Fund portfolios, we have constructed an automated investment simulator that uses GP to derive a useful non-linear relationship between a large number of factors relating to equities. This non-linear factor model assists the system in making buy/sell choices. Our system has been trained on a basket of 24 equities from the Malaysian stock market, and displays successful investment behaviour in out-of-sample tests. Having trained the system, our intention is that it should be used in a continuous learning mode so that it can adapt to changes in the Malaysian economy by modifying (via retraining) the non-linear equation. However,

standard GP systems are characterized by the fact that the final trained population often has little diversity and this could make it difficult to adapt to sudden changes in the economy.

We have developed a new mechanism for preserving phenotypic behavioural diversity in a GP population, inspired by the risk-management practices of fund managers. In this paper we present that new system, together with the results of experiments to determine:

- the efficacy of the new system — i.e. to what extent it improves diversity, compared with a standard GP system;

- the impact of increased diversity on the ability of the GP population to adapt to a new economic environment; and

- the behaviour of an individual chromosome obtained by such a system, compared with that obtained from a standard GP system — in particular, the fitness of a newly adapted individual after retraining on different amounts of data from a new economic environment.

2. RELATED WORK

Numerous techniques have been developed to maintain population diversity. These approaches have three main research directions:

1. the preservation of genotype diversity based on formally-defined structural distance measures;

2. the preservation of phenotype diversity based on the unique individual fitness in a population;

3. the reintroduction of genetic material in various evolutionary phases.

The great majority of existing work on diversity aims to preserve genotype diversity using various measures of distance such as an edit distance [9] or a string edit distance for linear GP [1]. Monsieur and Flerackers [15] detect identical subtrees within a GP individual and delete individuals with low diversity and "DGEA" [22] uses a distance-to-average-point structure measure to alternate between decreasing and increasing diversity. Another popular technique is Fitness Sharing which smoothes the fitness landscape, initially used

for maintaining diversity in genetic algorithms [6] though later applied to GP [5, 13, 14].

In contrast to the approach of preserving genotypic diversity, there is very little work in the area of preserving phenotypic diversity. Two example techniques in this area are: (i) using entropy (the amount of disorder of the population) and free energy to measure phenotype diversity in terms of the number of unique fitness values in the population [18, 17]; and (ii) using a selection method that is uniform over the fitness values [11].

Three examples of work where genetic material is introduced are: (i) random immigrants [7] where in every generation the population is partly replaced by randomly generated individuals; (ii) hypermutation [4], which increases mutation rate drastically whenever the average best individual performance worsens, and (iii) the Restart approaches [8, 10], which advocate a complete restart of the GP to maintain diversity after a change in the environment has occurred.

2.1 Limitations of related work

None of the previously mentioned techniques are wholly satisfactory. Studies have shown [2, 3] that genotype diversity approaches "may not be useful for capturing the dynamics of a population" and phenotypic measures appear to have better run performance. However, phenotype diversity techniques do not explicitly consider behavioural diversity of individuals, although they successfully spread individuals across different fitness levels. In other words, by ignoring the behaviour of individuals, the diversity in the same fitness level is not maintained. Legg and Hutter [11] remarked that "while the total population diversity was improved, the diversity among the fit individuals was not."

The reintroduction techniques also have their limitations: "random immigrants" and "hypermutation" do not perform well in continuously and abruptly changing environments, and "Restart" techniques ignore the knowledge accumulated by the population from previous learning. Restart techniques explicitly re-introduce diversity following a change in the environment, but may reconverge on a local optimum, especially when retraining only on data from the new environment where there are few new data points. Restart retraining on a mixture of old and new data suffers from the fact that it is very difficult to know *a priori* what mix of the two would give the best results [19].

Using any one technique on its own appears not to preserve useful diversity, implying that more elaborate techniques should be explored.

3. BEHAVIOUR DIVERSITY

We present a new mechanism to preserve diversity in phenotypic behaviour, inspired by the risk management techniques used by investment managers.

Investment managers often look for equities with contra-correlated market behaviour — i.e. equities whose prices move in opposite directions as economic factors change. A diversified portfolio [20] will invest in both in order to reduce the overall risk of losing money (because if one drops in price, the other will rise) whilst in many circumstances[1] the overall return is only slightly reduced [12]. Alternatively, a Hedge Fund manager might find two highly correlated stocks

and trade long in one, short in the other — a "market neutral" strategy where returns are not dependent on market movement but only on the relative behaviours of two stocks.

We treat a population of individuals in a similar manner to a diversified portfolio of stocks. Retaining an individual in the population is equivalent to investing in a stock; we attempt to identify correlations (and contra-correlations) in the phenotypic behaviour of those individuals, and attempt to construct a non-correlated population.

Consider two individuals that differ in the fine detail of their behaviour over time-series training data. If that detailed behaviour were important for evolutionary selection, then it would be part of the calculation of fitness and individuals with different detailed behaviour would differ in fitness. However, where the detail is irrelevant for selection, fitness can reflect general behaviour and two individuals who differ in detailed behaviour might have identical fitness. In this case, we propose that differences in detailed behaviour be used to preserve diversity of phenotypic behaviour.[2]

Given an individual that has a detailed phenotypic behaviour that is different from others of similar fitness: deletion of that individual would decrease the total population diversity of phenotypic behaviour. Conversely, if a set of individuals has entirely correlated behaviour and similar fitness then we can delete all but one from this set without comprising phenotypic behavioural population diversity.

Our new GP system preserves diversity at two levels:

1. Simple phenotypic fitness diversity.

2. Phenotypic behavioural diversity within groups of the same fitness.

For the former we hold individuals in fitness segments — each segment contains individuals of similar fitness. For the latter, we establish the correlation of the phenotypic behaviour of all individuals in a segment, partition them into correlated groups and for each group we delete all but one individual — this is repeated for each segment.

In practice, determining the correlation coefficients of GP individuals is not always straightforward. However, in our financial portfolio application it is achieved by comparing the changes in Return on Investment (ROI) at particular time periods. If one individual's ROI drops and at the same time another individual's ROI increases at the same point in time, then we define the two individuals to be contra-correlated at that point. The degree of correlation can be calculated using a metric such as the Spearman Correlation test to compare behaviour at every time point — see Section 3.1. Note that the overall fitness (Sharpe Ratio over the whole time period) for two individuals might be identical yet the ROI behaviours of the individuals might be contra-correlated.

3.1 Description of the algorithm

In our system, each individual has not only a fitness value but also an individual behaviour vector storing its historical behaviour (ROI performance). Let $O(I)$ be the vector of ROIs of an individual I during the whole time period and $p_t(I)$ be the ROI of an individual at a particular time t. Then, $O(I) = \{p_1(I), ..., p_t(I)\}$. We call $O(I)$ the *individual behaviour vector.*

---

[1]In particular, where the upside gains are greater than the downside losses.

[2]Note that this is *not* a multi-objective GP system.

At the beginning of each evolutionary selection cycle, after the initial population fitness is calculated, we group individuals according to their fitness value into a number of segments. Let F_{max}, F_{min} be the maximum and minimum fitness values for a population. Segments $\{(F_{min}, F_{min} + d), (F_{min} + d + 1, F_{min} + 2 * d), ..., (F_{min} + n * d + 1, F_{max})\}$ are collections of fitness intervals of equal length with defined lower and upper bounds.

Let $averageS$ be a vector of average performance outputs per time period of all n individuals $I_{S_1}...I_{S_n}$ in a fitness segment S. We call $averageS$ the *segment behaviour vector* given by

$$averageS = \{average_1(S), ..., average_t(S)\}$$

where

$$average_t(S) = \frac{1}{n} \sum_{i=1}^{n} p_t(I_{S_i})$$

Behavioural diversity in each of the fitness segments can be determined by measuring the correlation between an individual and its segment, as the segment behaviour vector should be a good representative of the general behaviour pattern of all the individuals in the segment. In this sense, if the difference of behaviour pattern between an individual and its segment is large, then it means that this individual exhibits unique or largely different behaviour from other individuals. We measure the behaviour relationship between the individual behaviour vector $O(I)$ and the segment behaviour vector $averageS$ based on the Spearman correlation test. The Spearman correlation measure simply ranks the two variables, and makes no assumption about the distribution of the values. The Spearman correlation coefficient $\rho(O(I), averageS)$ is computed as follows:

$$1 - \frac{6 \sum_{i=1}^{N} d_i^2}{N^3 - N}$$

where N is the numbers of pairs , and d_i is the distance between (i) the rank of performance $p_i(I)$ (compared with all other $p_j(I)$, $j \neq i$) in the individual behaviour vector and (ii) the rank of $average_i(S)$ (compared with all the other $average_j(S)$, $j \neq i$) in the segment behaviour vector. The degree of correlation returned by this measure varies from -1.0 representing negative correlation, through 0.0 indicating no correlation, to 1.0 representing positive correlation.[3]

If a segment contains a set of individuals that have correlated or similar behaviour throughout the run, the behaviour diversity of the segment would be compromised without deleting individuals from this set. Conversely, if we keep an individual whose reactions to the environment are very different from other individuals, the diversity level of the segment can be maintained. Therefore, we firstly check for the correlation coefficient value ρ of each individual in each segment and then we delete individuals with high ρ (higher than a predefined correlation threshold, which we set at

---

[3]Alternatively, one could measure the correlation relationship between each pair of individuals in the same segment; in other words, create a correlation coefficient matrix for the segment. The major drawback of this approach is the computational cost ($O(n^2)$ rather than $O(n)$).

0.67). Additionally, in the case of domination of one or two particular fitness segment(s), which means that the number of individuals contained in the segment are higher than the average (defined as P/s where P is the total population size and s the number of the segments), we delete not only those individuals with high correlation coefficient but also surplus individuals in the order of decreasing correlation. In this way, the algorithm encourages the creation of individuals at all fitness levels throughout the evolutionary cycle and also preserves diversity at the global population level. After crossover and mutation, randomly generated individuals are inserted into the population in order to keep the population size constant.

4. HEDGE FUND SIMULATION

To test the efficacy of the new diversity-preservation algorithm, we simulated a long/short market-neutral hedge fund of Malaysian equities. The GP system evolved a non-linear equation that used market data to determine whether each stock should be selected to buy, or to sell.

4.1 System overview

Our test system comprised a GP subsystem utilizing the new diversity-preservation algorithm (we call this **NGP**), coupled with an investment simulator. The coupling between the two was the fitness function — the investment simulator was called each time the GP subsystem needed to determine the fitness of an individual, at which point the individual was used to control the simulation of an hedge fund of Malaysian stocks. The simulator was applied to training data giving monthly prices and other factors for a period of 41 months. Monthly returns on investment were calculated, and at the end of each year the Sharpe ratio [21] was calculated. The simulator returned to the GP system both the fitness calculated from the Sharpe Ratio and a correlation vector of ROIs over the training period.

Fitness

The fitness f for an individual is given by Equation 1.

$$f = \frac{1}{1 - |1.5 - S|} \quad where \quad S = \frac{1}{4} \sum_{i=1}^{4} \frac{\bar{x}_i - RFR_i}{\sigma_i} \quad (1)$$

In Equation 1, S is the average Sharpe Ratio over the training period (comprising 4 sub-periods), $\bar{x}_i$ is the average monthly ROI over the sub-period i, σ_i is the standard deviation of monthly ROIs over the sub-period i, and RFR_i is the average monthly Risk Free Rate for sub-period i. We set RFR_i to 0.003 for all i (equivalent to 3.6% per annum). Fund managers often set a target Sharpe Ratio, as do we — our target is 1.5, and the absolute difference between the measured Sharpe Ratio and the target is then normalized to provide a fitness value that varies between 0 and 1.

Correlation

The vector of ROIs returned by the simulator was used by the GP system to assess correlation with other individuals. This correlation data was then used to guide the preservation of phenotypic behavioural diversity as explained in Section 3.1.

4.2 The Investment Simulator

We simulated a market-neutral long/short Hedge Fund of Malaysian equities. The fund focused on a basket of 24 Malaysian stocks, which it could buy ("go long") or sell (even if it didn't own any — "go short"). Since all the stocks in this basket were quite well correlated, the market-neutral strategy simply entailed buying the profitable stocks and selling short those stocks that were performing poorly.

The training data was monthly prices (and other technical and fundamental data) over a period of 41 months. Since we had only monthly data, all trading occured at the beginning of each month and the resulting stock mix was held for the duration of the month. At the beginning of each month, we used the individual provided by the GP system as a stock selection model that quantitatively measured the attractiveness of each stock; this model was a non-linear combination of technical and fundamental factors to predict the return expectation for each stock over a 4-week forward horizon.

For each month, we applied the stock selection model to the current month data — this was a table per stock with about 20 factors and 7,680 data points. A return prediction was assigned to each stock.

Stocks were grouped into 4 sectors; within each sector all stocks were ranked according to expected return. The simulator then made the following fund management decisions:

- The long/short portfolio was both dollar neutral and sector neutral. Thus, at all times, 24 stocks were maintained in the portfolio with 12 long positions and 12 short positions equally distributed across all the sectors. According to the ranking, the top 3 stocks in each sector became the top fractile and the bottom 3 became the bottom fractile. The top fractile of each sector and the bottom fractile of each sector were chosen to hold long positions and short positions respectively in the portfolio.

- Sectors were equally weighted and each stock was given equal weight in the portfolio. Thus, each position accounted for approximately 4% of total portfolio value.

- CFDs (Contract for Differences) were used instead of conventional shares to trade on stocks. We assumed 20% notional trading requirement (margin), 0.25% trading commission, and 5% financing rate.

At the end of each month, all of the positions held in the portfolio were closed and the profit or loss of the portfolio during the month was calculated. At the beginning of the next monthly trading cycle, the simulator updated the expected return based on the new "current" data and a new desired long/short portfolio was formed.

5. METHOD

The motivation for our research is to develop a system that can be used in a continuous-learning context, where the economic environment is dynamic and unpredictable (and so the most successful non-linear factor model will vary).

Following a shift in economic context it will be necessary to continue to use the previously trained "best" individual while new data is being accumulated for retraining — we are interested in the behaviour of this previously-trained "best" solution, and how well it performs in the context of a new economic environment. When sufficient new data has been accumulated for re-training, we require that the system should re-train effectively on the new data — our expectation is that shifts in context will not normally be excessive and therefore retraining should start from the previously trained population (rather than from a random population) since the system is then less likely to converge on a local optimum, especially when there are few data points available for the new environment [16]. Initial training will use an "in-sample" data set: subsequent retraining will be on data that was "out-of-sample" in the context of the original training.

We performed experiments to answer the following three research questions:

1. Does the new technique really improve population diversity when compared with a standard GP system, and how does this affect fitness?

2. When retraining on data that comes from a different economic context, how quickly and how well does the diverse population adapt to a new environment?

3. Are trained individuals from the new system more robust when exposed to a new economic environment?

We ran three experiments, presented below. In each case, a standard GP system (SGP) was compared with our new GP system (NGP). The training data and validation data was in all cases identical for both SGP and NGP.

5.1 In-sample population dynamics

The aim of Experiment 1 was to see if NGP improved diversity, and what effect it had on fitness. Both SGP and NGP were trained (separately) using 41 months of financial time-series data for 24 Malaysian stocks, taken from the period 31/7/97 to 31/12/2000.

The following measurements were made each generation:

1. the fitness of the best (fittest) chromosome, the fitness of the worst chromosome, the average fitness across all chromosomes, and the standard deviation of all chromosome fitnesses;

2. the distribution of chromosomes in the segment vector in terms of (i) fitness and (ii) phenotypic behaviour.

After repeating the experiment 5 times, the average value (across 5 runs) of each measurement was plotted for each generation (see Fig. 1); in each case the best run and the worst run were indicated by error bars.

5.2 Retraining population dynamics

The aim of Experiment 2 was to see how well the population as a whole adapted to a new environment; in each case we started with a population that had previously been trained on the 41 months of training data (Experiment 1 above) and then retrained SGP and NGP separately on new training data that reflected a different economic climate. We used three sets of new training data and investigated the population behaviour for each:

Period a: 31/12/2000 - 31/4/2002 (16 months).

Period b: 31/4/2002 - 31/1/2002 (9 months)

Period c: 31/1/2002 - 31/12/2003 (11 months)

As with Experiment 1, this experiment was repeated five times and for each generation the best, average and worst fitnesses, and standard deviation, were measured (see Figs. 2, 3, and 4). Convergence characteristics were compared by inspection of the graph.

5.3 Retraining individual behaviour

The aim of Experiment 3 was to simulate a continuous-learning context and investigate the fitness behaviour of individuals after retraining on successively larger amounts of new-context data. Specifically, we were interested to learn how many data points of new data were needed to get a robust chromosome that behaved well on the rest of the new data. This experiment utilised Period a (the largest period), divided into four consecutive phases (i.e. the initial 25%, the initial 50%, the initial 75%, and the whole period). The experiment proceeded as follows:

1. Using the best trained individual from Experiment 1, a validation test was run on the whole of Period a. First the best trained individual from the SGP system was used, and then the best trained individual from the NGP system was used.

2. For both SGP and NGP separately, starting with the previously trained population from Experiment 1, retraining was performed using the phase 1 data (i.e. the initial 25% of the Period a data). In each case the best chromosome was obtained and a validation test was run on the remaining 75% of Period a.

3. For both SGP and NGP separately, starting with the previously trained population from Experiment 1, retraining was performed using the phase 2 data (i.e. the initial 50% of the Period a data). In each case the best chromosome was obtained and a validation test was run on the remaining 50% of Period a.

4. For both SGP and NGP separately, starting with the previously trained population from Experiment 1, retraining was performed using the phase 3 data (i.e. the initial 75% of the Period a data). In each case the best chromosome was obtained and a validation test was run on the remaining 25% of Period a.

In all cases (for both SGP and NGP) the monthly Return on Investment (ROI) was measured and recorded, the Sharpe Ratio for the entire test period was calculated, and the experiments were repeated five times.

The results for Experiment 3 are shown in Fig. 5. The figure plots the average across 5 runs for ROI for both SGP and NGP and separately shows the Sharpe Ratio comparison. Error bars are included for the Sharpe Ratio, but excluded for the ROI for clarity. At the bottom right of the figure, a Sharpe Ratio comparison is provided to indicate increasing success as the amount of retraining data from the new economic context increases. As before, error bars are included for the Sharpe Ratio.

6. DISCUSSION OF RESULTS

6.1 In-sample population dynamics

Our first research question was *"Does the new technique really improve population diversity when compared with a standard GP system, and how does this affect fitness?"*.
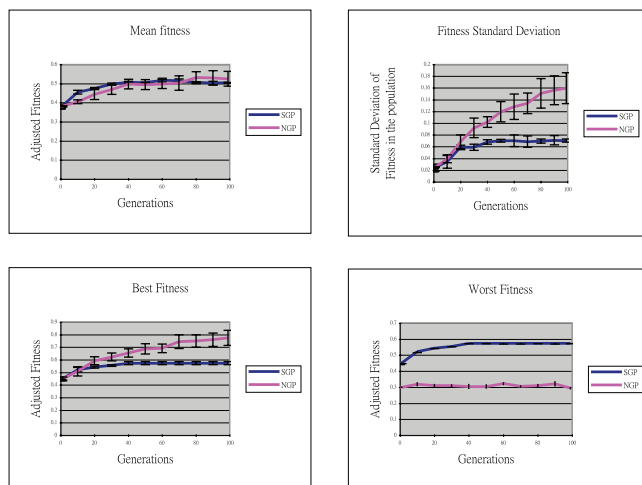

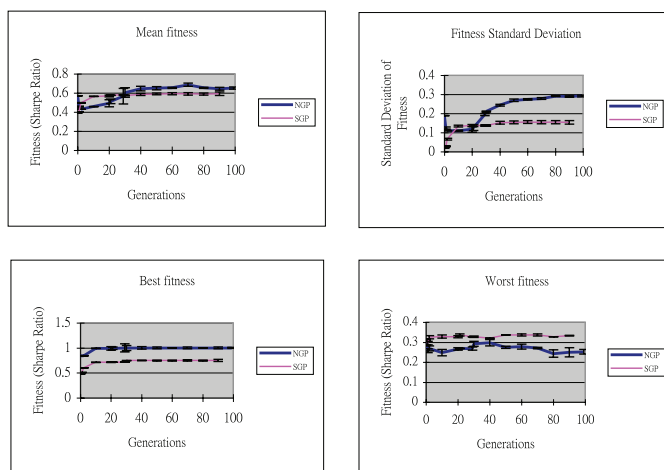**Figure 1: Experiment 1, in-sample training (population dynamics).**

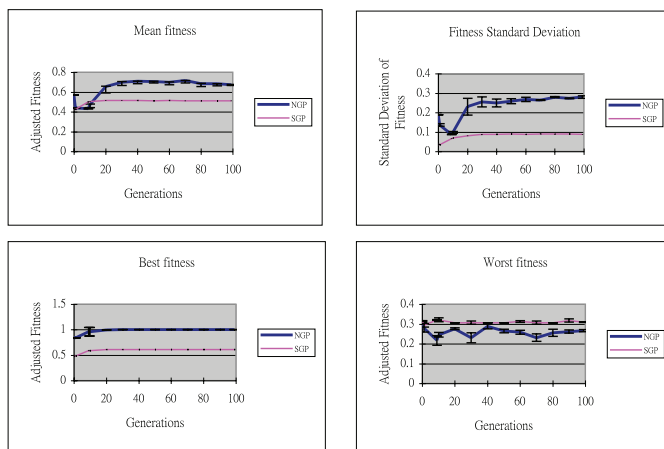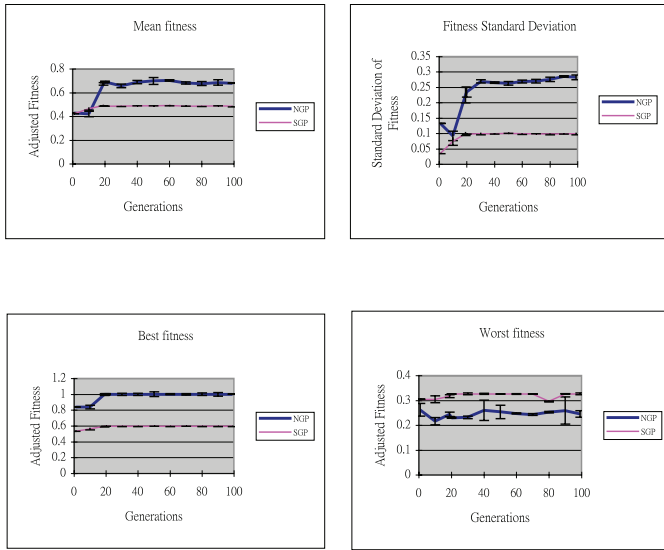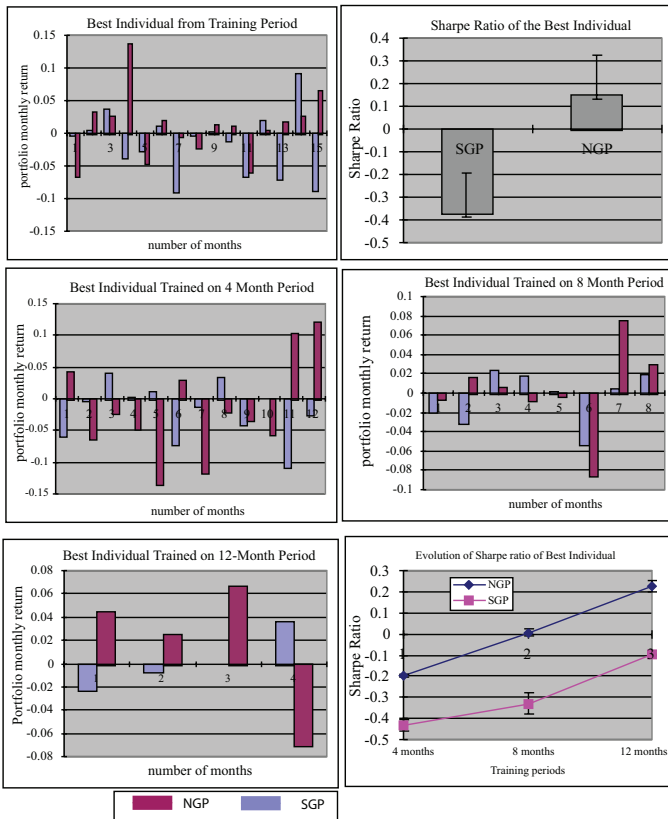**Figure 2: Experiment 2, retraining Period a (population dynamics).**

**Figure 3: Experiment 2, retraining Period b (population dynamics).**

Figure 4: Experiment 2, retraining Period c (population dynamics).

Fig.1 illustrates how the fitness of individuals in the population evolved over 100 generations while being trained on the base data of 41 months. Whilst the mean fitness difference between SGP and NGP is negligible, and remains negligible as the population evolves, the difference in standard deviation of fitnesses clearly increases as the population evolves, with NGP developing a substantially greater diversity of fitness. It is interesting to note a corollary of these two statements, which is that the fitness of the *best* individual is much greater for the NGP system — this in itself is a compelling reason to use diversity-preserving techniques, regardless of other possible benefits in terms of adaptability.

Fig. 6 gives the distribution of individuals across the fitness segments at three points during evolution (after 0, 50 and 100 generations), indicating a slight bi-modal characteristic for both SGP and NGP which may merit further investigation. Of more importance from the point of view of *phenotypic behavioural* fitness, Fig. 7 investigates the correlation coefficent of individuals in the highest fitness segment (since this is the segment from which the final solution(s) will be drawn). The Figure illustrates how phenotypic behavioural diversity increases as the population evolves — on the left is a graph of correlations from a single run of NGP showing the mean, best and worst correlations of individual behaviour vectors with the segment behaviour vector; on the right is a graph comparing the mean correlations (with error bars for 5 runs) for NGP and SGP.[4]
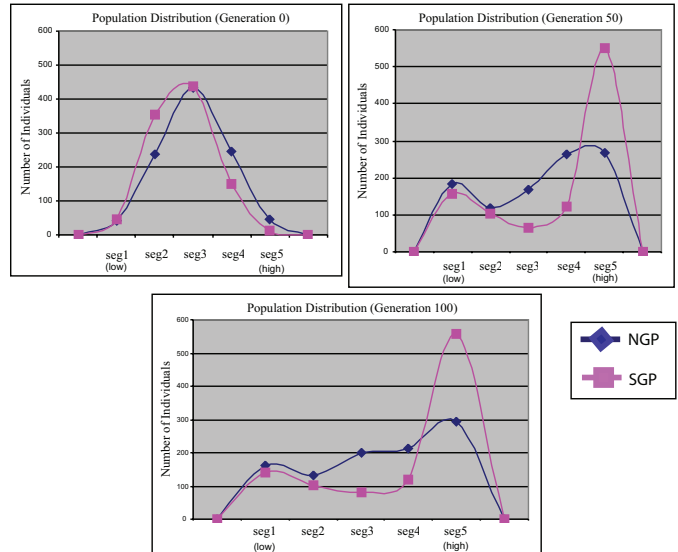
Figure 6: Phenotypic diversity for SGP and NGP

The results indicate that the new algorithm NGP provides a substantial improvement in diversity of phenotypic behaviour, when compared with SGP on this data. To analyse this further we used Student's T-test (since the number of data points is small). The T-test results (less than 0.00002

Figure 5: Validation of successive retraining of individual

---

[4]To obtain the corresponding data for SGP, it was necessary to save population snapshots for generations 0, 20, 50, 80 and 100, and determine which individuals would have been in the highest fitness segment if NGP were being run; then each such individual was used in the investment simulator and its individual behaviour vector obtained; finally, a segment behaviour vector was calulated and the correlations measured (this was repeated for each "snapshot").

Figure 7: **Evolution of behavioural phenotypic diversity for SGP and NGP**

for Generations 80 and 100) confirm that for the later generations NGP provides a highly significant improvement.

6.2 Retraining population dynamics

Our second research question was *"When retraining on data that comes from a different economic context, how quickly and how well does the diverse population adapt to a new environment?"*

Figs. 2, 3 and 4 illustrate how the populations for SGP and NGP evolve during re-training on data from different economic climates. In all three cases (Periods a, b and c) the standard deviation of fitness for NGP is significantly better for NGP than for SGP after about 20 generations. NGP always starts with a "best" fitness that is higher than SGP (indicating that the NGP individual is already better able to cope with the new economic climate than the SGP individual), and that superiority is retained as the population retrains — in fact, in all cases NGP quickly finds the "ideal" Sharpe Ratio of 1.5 (giving the best adjusted fitness of 1.0) whereas SGP never finds that ideal individal.

When retraining, both SGP and NGP appear to converge on their respective "best" individuals at about the same rate, though SGP has the edge for Period c.

An interesting result from Figs. 2, 3 and 4 is the fact that the NGP standard deviation decreases for the first ten generations of re-training, before then rising strongly. The initial generation comprises individuals trained for a different economic context, and these individuals do not have a normal distribution of fitnesses in the context of the new data. In particular, (i) the fitness peak for the highest segment that is observed at the end of the original training is unlikely to occur in the context of the new data, and it is much more likely that a peak will occur in the middle segments, and (ii) the behavioural diversity of individuals in the highest segment is likely to be low. In the early generations, the NGP algorithm will discard many individuals whose behaviours are too well correlated with others, and replacements will be drawn from random samples — the overall effect will be to increase the numbers of individuals in the middle segments (i.e. producing a more peaked distribution). The NGP standard deviation therefore dips towards that of SGP, then rises strongly after the distribution has normalised and the diversity-preserving algorithm has a stronger effect.

6.3 Retraining individual behaviour

Our third research question was *"Are trained individuals resulting from the new system more robust when exposed to a new economic environment?"*

Fig. 5 illustrates how the best trained individual performs when the economic context changes. The monthly ROI in the new economic climate is plotted for both the best NGP individual and the best SGP individual. The Sharpe Ratio over the test period is also calculated and compared. We did this for four cases:

The best individual from the 41-month training period, tested over the whole of Period a. Whilst the plot of monthly returns shows that both SGP and NGP have months with poor performance, the comparison of Sharpe Ratios indicates that NGP has significantly superior performance on this test data, giving an overall positive risk-adjusted return whereas SGP gives an overall negative risk-adusted return. This suggests that NGP might produce individuals that are more robust to changes in economic climate — though of course we need to conduct further tests on different kinds of change in the economy before making any firm claim of this nature.

The best individual after retraining on three successively larger amounts of new data, tested on the remainder of the new data. Once again we use Period a as our test case. First retraining was carried out on 4 months data (tested on 12 months data), then retraining on 8 months data (tested on 8 months data) and finally retraining on 12 months data (tested on 4 months data). This was intended to simulate a continuous-learning context where the GP system would be periodially retrained in order to accommodate changes in the economy. As with the first case (above), the monthly ROI plot indicates that both NGP-trained and SGP-trained individuals have poor performance in some months. However, the plot of Sharpe Ratio evolution for these three cases shows that both NGP and SGP performance improves as more data from the new economy becomes available for training, and that the NGP individual starts with a higher Sharpe Ratio,

Of course, a GP system would be expected to improve performance if it has more data points in the training data. However, in practice the upper bound on the size of the training data is not as important as the lower bound — i.e. how small can the training set be in order for retraining to be effective? This lower bound determines how quickly an automated investment algorithm will recover from an abrupt change in the market. The plot of Sharpe Ration evolution shows that *on this data* NGP adapts more quickly to the change: NGP improves its Sharp Ratio from -0.2 to just over +0.2 (an increase of 0.4) whereas SGP in the same time and on the same data improves its Sharpe Ratio from just under -0.4 to -0.1 (an increase of 0.3).

7. CONCLUSION

We have presented a new Genetic Programming algorithm for preserving diversity of phenotypic behaviour, and presented results of experiments based on simulation of a market-neutral long/short hedge fund of Malaysian stocks. The technique is applicable to problems that have time-series data as input and that are subject to a changing environment.

Our results confirm that, on our test data, the new algo-

rithm does what it says — it increases population diversity of phenotypic behaviour. It also increases diversity of standard fitness. Furthermore, analysis indicates that the difference from a standard GP algorithm is statistically highly significant.

Our results also provide an insight into the utility of the new algorithm (NGP) compared with a standard GP algorithm (SGP):

1. On our test data, NGP consistently produced a "best" individual with higher fitness than SGP.

2. On our test data, NGP consistently retrained (in the context of a sudden shift in the environment) faster and better than SGP.

We conclude that (i) there is good evidence to recommend the use of NGP for preserving diversity of phenotypic behaviour *in any GP context*, and (ii) there is encouraging initial evidence to recommend the use of NGP for retraining contexts where the environment is continually changing. Further work is now required: to investigate the use of different correlation measures; to undertake a parameter sensitivity analysis of our system; to establish that these results are repeatable for a much wider range of test cases and over more runs; and to obtain empirical data from a continuous-learning system.

8. ACKNOWLEDGMENTS

9. REFERENCES

[1] M. Brameier and W. Banzhaf. Explicit control of diversity and effective variation distance in linear genetic programming. In *Proceedings of the 5th European Conference on Genetic Programming*, volume 2278 of *LNCS*, pages 37–49. Springer, 2002.

[2] E. K. Burke, S. M. Gustafson, G. Kendall, and N. Krasnogor. Advanced population diversity measures in genetic programming. In *Proceedings of the 7th international Conference on Parallel Problem Solving From Nature*, volume 2439 of *LNCS*, pages 341–350. Springer-Verlag, 2002.

[3] R. Burke, S. Gustafson, and G. Kendall. A survey and analysis of diversity measures in genetic programming. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 716–723. Morgan Kaufmann, 2002.

[4] H. Cobb. An investigation into the use of hypermutation as an adaptive operator in genetic algorithms having continuous, time-dependent nonstationary environments. Technical report, Naval Research Laboratory, 1990.

[5] A. Ekart and S. Nemeth. Maintaining the diversity of genetic programs. In *Proceedings of the 5th European Conference, Foster, J. et al., editors, Genetic Programming*, volume 2278 of *LNCS*, pages 162–171. Springer-Verlag, 2002.

[6] D. Goldberg. *Genetic Algorithms in Search, Optimisation, and Machine Learning*. Addison-Wesley, 1989.

[7] J. Grefenstette. Genetic algorithms for changing environments. In *Parallel Problem Solving from Nature 2*, pages 137–144. North Holland, 1992.

[8] C. Karr. Genetic algorithms and fuzzy logic for adaptive process control. In *Intelligent Hybrid Systems, S.Goonatilake and S.Khebbal, eds.* Wiley, 1995.

[9] R. Keller and W. Banzhaf. Explicit maintenance of genetic diversity on genospaces. In *Unpublished Manuscript*, 1994.

[10] M. Kidwell and D. Cook. Genetic algorithm for dynamic task scheduling. In *In Proceeding of the IEEE 13th International Phoenix Conference on Computers and Communications*, pages 61–67, 1994.

[11] S. Legg and M. Hutter. Fitness uniform deletion: A simple way to preserve diversity. In *Proc. Genetic and Evolutionary Computation Conference, GECCO*, pages 1271–1278. ACM, 2005.

[12] B. Malkiel. *A Random Walk Down Wall Street*. W.W. Norton & Company, ISBN: 0393057828, 2003.

[13] R. McKay. Fitness sharing in genetic programming. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 435–442, 2000.

[14] R. Mckay and H. Abbass. Anticorrelation measures in genetic programming. In *Australian-Japan Workshop on Intelligent and Evolutionary System*, pages 45–51, 2001.

[15] P. Monsieur and E. Flerackers. Increasing the diversity of a population in genetic programming. In *Proceeding of Genetic and Evolutionary Computation Conference (GECCO)*, page 185. Morgan Kaufmann, 2001.

[16] C. Reeves and H. Karata. Dynamic sequencing of a multi-processor system: a genetic algorithm approach. In *Artificial Neural Nets and Genetic Algorithms*. Springer, 1993.

[17] J. Rosca. Entropy-driven adaptive representation. In *Proc. Workshop Genetic Programming: From Theory to Real-World Applications, J. Rosca, Ed.*, pages 23–32, 1995.

[18] J. Rosca. Genetic programming exploratory power and the discovery of function. In *Proc. Nat'l. Conf. on Evolutionary Programming*, 1995.

[19] C. Schoreels and J. M. Garibaldi. The effect of varying parameter on performance for adaptive agents in techical equity market trading. In *Proceedings of the 3rd IEEE Internatioal Conference on Computational Cybernetics (ICCC 2005)*, 2005.

[20] W. Sharpe. Risk, market sensitivity, and diversification. *Financial Analysts Journal (January/February 1972)*, pages 74–79, 1972.

[21] W. F. Sharpe. The sharpe ratio. *J. Portfolio Management*, 21:49–58, 1994.

[22] R. K. Ursem. Diversity-guided evolutionary algorithms. In J. Guervos, P. Adamidis, H. Beyer, J. Martin, and H. Schwefel, editors, *Proceedings of the 7th international Conference on Parallel Problem Solving From Nature*, volume 2439 of *LNCS*, pages 462–474. Springer-Verlag, 2002.