

Using Fitness Comparison Disagreements as a Metric for Promoting Diversity in Dynamic Optimisation Problems

Edgar Galván-López*

School of Computer Science and Statistics,
Trinity College Dublin, IE
edgar.galvan@scss.tcd.ie

Ouassim Ait ElHara

TAO, INRIA & LRI, CNRS & U. Paris-Sud,
Université Paris-Saclay, FR
ouassim.ait_elhara@inria.fr

Abstract—In Evolutionary Algorithms (EAs), it is well-known that the adoption of diversity is highly beneficial for evolutionary search. This has also been explored and confirmed in Dynamic Optimisation Problems (DOPs) using EAs. Multiple works have been proposed to encourage diversity in EAs in the face of a change, where the most common form to promote diversity is to replace a number of individuals by new genetic material. A common element when adopting this form of diversity is the fact that, frequently, the number of individuals to be replaced is picked rather arbitrarily. In this work, we propose the adoption of the Kendall tau distance that quantifies pairwise dissimilarities among two lists (of fitness values) with the hope to make a better informed decision in terms of the number of individuals that need to be replaced in a population by new individuals. Results on continuous fitness-valued cases indicate that the adopted distance is beneficial in DOPs.

I. INTRODUCTION

Multiple definitions have been proposed to describe optimisation in dynamic environments. In this work, we adopt the most common definition of dynamic optimisation problems (DOPs) which are defined as problems that are solved online by an optimisation algorithm as time progresses [12].

Multiple evolutionary algorithm (EA) elements have been reported to be beneficial in DOPs (see e.g., diversity [12], immigrants [9], memory schemes [10]).

One key element is diversity. Diversity is a key element of the biological theory of natural selection and it is used in EAs to describe, for instance, structural [4] or behavioral variety [2], [5]. A large amount of research has proved that diversity is necessary to escape from local optima. Intuitively, this means that if the vast majority of the population has converged to a local or global optimum (e.g., this majority being structurally similar among themselves) and the environment has changed (consequently, the global optimum changing too), it is unlikely that the population would be able to escape from it. Thus, it is expected that an EA with a mechanism to promote diversity will greatly improve its performance.

Multiple approaches have been proposed to keep diversity in EAs e.g., changing a genetic operator rate, changing the

size of a population. One commonly adopted approach is the replacement of individuals in a population by newly generated genetic material. A common element observed when adopting the latter is that, frequently, researchers used an arbitrary approach to decide the number of individuals that need to be replaced. However, this process is purely intuitive and often expensive due to its trial-and-error nature.

To the best of our knowledge, however, there are no studies that have focused their attention on a more informed way to determine the number/proportion of individuals that need to be replaced in DOPs and this work intends to shed some light on this. To this end, we use the fitness values of individuals in the population as indicators to determine how big/small a change is, and consequently, use this to determine the number of individuals to be generated to promote structural diversity. Any population-based EA can adopt our proposed approach, described in Section III, and in this work, we use a Genetic Programming (GP) system [7].

This paper is organised as follows. In Section II, we discuss related works. Section III shows the proposed approach. Section IV provides details on the experimental setup used. The results are presented in Section V, and finally, conclusions and future work are drawn in Section VI.

II. PREVIOUS WORK ON PROMOTING AND MAINTAINING DIVERSITY

As indicated previously, multiple works have been proposed to promote and maintain diversity in EAs (see [12] for a summary of works). In this section, we focus only on DOPs tackled by GP (see [12], [18] for a more general treat on the subject). Among those approaches proposed to promote diversity in the face of DOPs using GP are: (a) injection of new genetic structural material, (b) adaptable genetic operators, and (c) behavioural diversity.

One of the easiest forms of promoting diversity is adopting the injection of new genetic material into the GP population. The generation of GP individuals is done by using common techniques, like the adoption of the ramped half-and-half method [7]. This can take place when, for instance, detecting a change [8] or when bloat (dramatic increase of tree sizes

*Research conducted during Galván's stay at TAO, INRIA and LRI, CNRS & U. Paris-Sud, Université Paris-Saclay, France.

as evolution proceeds) reaches a limit and there is a need to substitute individuals contained in the population by new GP programs [16]. Injecting new GP individuals into the population has also been promoted via culling [13]. That is, removing the worst (less fit) individuals and replacing them by randomly generated GP programs. Variable population size [15] also promotes diversity by adding new GP individuals into the population. A common element in all these works is that, often, the number of individuals to be replaced by the same number of newly created individuals is chosen rather arbitrarily.

Another well-known approach to promote/increase diversity in GP facing DOPs is to use adaptable genetic operators. The idea is straightforward and easy to implement in a GP system: to change (increase/decrease) the frequency of applying a genetic operator to individuals in a GP population. Examples of works inspired by this technique include [8], where the authors increased the subtree mutation rate when a change in the problem is detected. The adoption of this technique has also been adopted when there is a change e.g., in the fitness of the best individual for a number of generations [13], where the authors decreased the crossover rate if the fitness value increased (for a maximisation problem) over generations and increased the crossover rate, otherwise. An opposite approach has been conducted when using elitism (keeping the best solution found so far), that is, the elitism rate was increased when the fitness increased, for a maximisation problem, and decreased up to a certain threshold, when the fitness deteriorated [13].

Finally, the study of the behavioural diversity of GP has increased dramatically over the last years as a result of multiple scientific studies indicating that the presence of behaviour (semantic) diversity can dramatically improve the performance of a GP system. Multiple semantic methods have been proposed in GP, and so, it is not possible to give a unique definition of this. However, the common element present in these approaches is to capture how a GP system behaves on a given problem. The most popular form of semantics GP, as originally proposed in [11], is defined as the vector of output values computed by a (sub)tree for each set of input values in turn (a.k.a. fitness cases in most cases). From this, one can compute the semantic (behaviour) diversity among two individuals. That is, one can say that two individuals are behaviourally different if their output vectors are different (or some of the values contained in these vectors are different). Surprisingly, only a few scientific studies have focused their attention on promoting behavioural diversity in DOPs. Yan and Clack are among those few researchers promoting this type of diversity in a dynamic setting [17] by using the common semantic diversity method explained before.

A. Final Comments on Diversity

By far, the most used technique to promote diversity in GP on DOPs is the adoption of structural diversity. As discussed earlier, this has been adopted differently in various research works. However, one predominant element in all these works is

that the decision on how many individuals should be replaced is made rather arbitrarily. Next, we present an approach that aims to overcome this limitation.

III. PROPOSED APPROACH

From the previous section, we know that there is strong evidence indicating that the adoption and/or encouragement of structural diversity in GP search on DOPs is highly beneficial. Normally, when adopting this type of diversity, researchers have focused their attention on setting arbitrarily a number of individuals to be generated and then used them to e.g., replace the worst GP individuals in a population. The major drawback with this approach is that, very often, this process is based on trial and error and can be computationally expensive.

We believe that it is possible to adopt a more informed way of determining the number of individuals that should be replaced from a population by using fitness values. The use of these values as indicators to perform a specific task (e.g., prediction of problem hardness) is common in EAs. The most well-known example of this is the fitness-distance correlation [6], where these fitness values are used in conjunction with a metric that informs us how distant two individuals are in the search space to determine the difficulty of a problem.

In this work, we use a distance that accounts for pairwise disagreements between two lists of ranked fitness values. We hope that these disagreements can inform us on whether an evolved population is useful/relevant in the face of a change. Our proposed approach works in three phases:

- 1) Firstly, it is necessary to account for a method that can indicate when a change is about to take place. We do this in a non-expensive matter: before a new generation is about to take place, we use one individual (the elitist individual), whose fitness (f_e^g) is assessed again in the next generation ($g + 1$).
- 2) Secondly, if f_e^g and f_e^{g+1} are different, then we regard this as a change in the environment and we then proceed to compute the Kendall tau rank distance (defined in Eq. 1) between the ranking of the fitness values of all individuals at generation g and the next generation ($g + 1$). This distance counts the number of pairwise disagreements between two ranked lists and it is normalised by the maximum number of possible disagreements. This distance gives a discrete value $k = [0, 1]$ and this is used to generate a percentage of C_k new individuals with respect to the size of the population.
- 3) Thirdly, the worst individuals (less fit individuals) at $g + 1$ are replaced by the newly generated individuals keeping the size of the population constant (using ramped half-and-half initialisation method, details are discussed in Section IV).

The Kendall tau distance between two ranked lists is defined as,

$$k(\tau_1, \tau_2) = \sum_{(i,j) \in P} \bar{k}_{i,j}(\tau_1, \tau_2) \quad (1)$$

TABLE I
SUMMARY OF PARAMETERS.

Parameter	Value
Population Size	800
Generations	200
Type of Crossover	Any node
Crossover Rate	0.80
Type of Mutation	Subtree
Mutation Rate	0.20
Selection	Tournament (size = 7)
Initialisation Method	Ramped half-and-half
Initialisation Depths:	
Initial Depth	2
Final Depth	5
Maximum Length	1200 nodes
Maximum Final Depth	8
Independent Runs	50
Changes	Every 50 generations Random [1,50)

where, P is the set of pairs of elements in τ_1 and τ_2 , $\bar{k}_{i,j}(\tau_1, \tau_2) = 0$ if i and j are in the same order in both τ_1 and τ_2 ; and 1 if i and j are in opposite order in τ_1 and τ_2 .

It is worth mentioning that when the change to the objective function is monotonically increasing (order preserving), the computed Kendall tau distance will be 0. This is a good property of the distance because in this case the evolved individuals are expected to behave well in the changed objective function, so there is no need to replace individuals. A mirror image is seen in the presence of a monotonically decreasing change of the objective function, which will yield the maximal normalised distance of 1, meaning that the order of both fitness lists are completely different. The latter will indicate that our proposed approach based on the Kendall tau distance will replace the entire population by newly generated genetic material.

IV. EXPERIMENTAL SETUP

To test the proposed approach, we have adopted four symbolic regression functions, shown in Table II. The goal on these problems is to find a program whose output matches the outputs of these functions. Thus, the fitness of an individual in the population reflects how close the output of this individual are to the target. It is common to define the fitness as the sum of absolute errors measured at different values of the independent variable x , in this case in the range $[-1.0, 1.0]$. In this study we have measured the errors for $x \in \{-1.0, -0.9, -0.8 \dots 0.8, 0.9, 1.0\}$. We have defined an arbitrary threshold of 0.01 to indicate that an individual with a fitness less than the threshold is regarded as a correct solution, i.e. a “hit”. The function set is $F = \{+, -, *, /\}$, where $/$ is protected division.

To simulate a dynamic environment, we changed the objective function (changed the last ‘+’ sign symbol in each function (f_i) by a ‘-’ symbol (g_i), as shown in Table II) either every 50 number of generations or a random number of generations in the range of $[0, 50)$. We then repeated the

TABLE II
SYMBOLIC REGRESSION BENCHMARKS PROBLEMS USED IN OUR WORK.

Function	Objective function
f_1	$x^3 + x^2 + x$
g_1	$x^3 + x^2 - x$
f_2	$x^4 + x^3 + x^2 + x$
g_2	$x^4 + x^3 + x^2 - x$
f_3	$x^5 + x^4 + x^3 + x^2 + x$
g_3	$x^5 + x^4 + x^3 + x^2 - x$
f_4	$x^6 + x^5 + x^4 + x^3 + x^2 + x$
g_4	$x^6 + x^5 + x^4 + x^3 + x^2 - x$

process until the number of generations has been reached (e.g., generations $[0, 49]$, $[50, 99]$, $[100, 149]$, $[150, 199]$ using f_1 , g_1 , f_1 , g_1 , respectively).

Moreover, for comparison purposes, we also used an arbitrary approach in terms of the number of individuals that are replaced once a change has occurred as commonly adopted in EAs DOPs. In both methods, including our Kendall tau-based distance, we generated our individuals using the ramped half-and-half method, where the initial and final depths used are the same as when generating the population (see Table I).

The experiments were conducted using a steady state approach. The parameters used in this study are shown in Table I. To obtain meaningful results, we performed an extensive empirical experimentation ($50 * 2 * 2 * 4$ runs in total)¹.

V. ANALYSIS OF RESULTS

A. Performance

Let us first focus our attention on the average of the best fitness values before a change takes place. This is shown in Table III. The statistical significance for these results was computed using a T-test at 90% level of significance, where significant values are highlighted in boldface.

These results are encouraging in terms of adopting the Kendall tau distance on determining the number of individuals that should be replaced when a change occurs: in six out of eight cases, our proposed approach yield better results compared to an arbitrarily model (random) frequently used by EAs practitioners in DOPs. The latter approach yields slightly better results only on the last problem (f_4 , g_4).

If we now focus our attention on the percentage of success rate achieved by both approaches, shown in Table IV, we can see a similar trend: our Kendall tau distance approach achieves higher success rates. In fact, in none of the problems nor in any type of change, the success rate achieved by our approach is lower compared to the arbitrary model.

B. Average on Fitness Values and Length on Individuals

We now proceed to analyse the average fitness of both, the elitist individual and the entire population, and also, the average length of individuals in the presence of a fixed change, when using an arbitrary model and our adoption of the Kendall

¹50 independent runs, 2 types of replacement of individuals (arbitrary, Kendall tau distance-based), 2 type of changes (fixed change at every 50 generations, random change between $[1, 49]$), 4 problems (shown in Table II).

TABLE III
AVERAGE OF BEST FITNESS VALUES BEFORE A CHANGE TAKES PLACE OVER 50 INDEPENDENT RUNS.

Type of Change	Replacement type							
	Kendall tau distance				Arbitrary			
	f_1, g_1	f_2, g_2	f_3, g_3	f_4, g_4	f_1, g_1	f_2, g_2	f_3, g_3	f_4, g_4
Fixed at every 50 gens.	0.8764	0.4866	0.3343	0.2705	0.8273	0.4608	0.3339	0.2738
Random in the range of [1,50)	0.7964	0.3694	0.3174	0.2425	0.7447	0.3662	0.2945	0.2525

TABLE IV
PERCENTAGE OF SUCCESS RATE USING TWO TYPE OF REPLACEMENTS ON FOUR DIFFERENT SYMBOLIC REGRESSION FUNCTIONS.

Type of Change	Replacement type							
	Kendall tau distance				Arbitrary			
	f_1, g_1	f_2, g_2	f_3, g_3	f_4, g_4	f_1, g_1	f_2, g_2	f_3, g_3	f_4, g_4
Fixed at every 50 gens.	84.0%	27.0%	3.5%	1.0%	76.5%	23.0%	3.5%	1.0%
Random in the range of [1,50)	73.81%	12.54%	4.18%	0.36%	66.72%	12.18%	2.0%	0.36%

tau distance. This is depicted in Figure 1. Let us first focus our attention on the fitness values, shown in the left and centre-hand side of the figure, for the elitist individual and the entire population, respectively. It is clear that the Kendall tau distance (indicated by a dashed line) yields higher values compared to the arbitrary model (indicated by a solid line) for these problems modelled as maximisation problems. This is in correspondence to the performance discussed above and as shown in Tables III and IV. The tendency is stronger in rather easy problems (f_1, g_1 and f_2, g_2) compared to more complex problems (f_3, g_3 and f_4, g_4 - see Table II). If we continue our inspection and turn now our attention to the average length of individuals when using either our proposed approach or an arbitrary model. We can now observe mixed results. For instance, for f_1, g_1 the average length of the entire population when using the Kendall distance is lower compared when using the arbitrary model. This, however, is the opposite when using f_2, g_2 , where the Kendall distance tend to produce slightly longer individuals compared to the arbitrary model. This mixed tendency continues for both f_3, g_3 and f_4, g_4 .

Now let us analyse the average of fitness values, for both the elitist individual and the fitness values of the population, and the average length of individuals in the presence of a random change when using either an arbitrary model or our proposed approach, shown in Figure 2. First, let us discuss the fitness values trend. In most of the cases, the Kendall tau distance yields higher values in each of the four symbolic regression functions used in this work compared to an arbitrary model, which is in correspondence to the performance achieved by these two approaches, as discussed earlier (see Tables III and IV). As for the average length of individuals, we can see the same tendency as before: there are mixed results on this. That is, in some cases, the length of individuals tends to be smaller when using our proposed approach based on the Kendall tau distance compared to an arbitrary model, for example, see the results when using f_1, g_1 shown in the top-right hand side of Figure 2, whereas a mirror image can be seen when using f_2, g_2 shown below the length of f_1, g_1 .

C. Analysis of Successful Runs via Generations Taken to Solve a Problem and Number of Created Individuals

The previous figures helped us to understand some general aspects in terms how a GP system behaves, using either an arbitrary approach or our proposed approach to determine the number of individuals that need to be created in the presence of a change to promote diversity. However, the previous analysis does not inform us how these two approaches behave when solving a given problem.

To this end, we focus our attention only on those runs that were able to solve a problem and keep track on the number of generations necessary to solve a problem, and also, the number of individuals that were created by either the arbitrary approach or by our approach. This is shown in Figure 3, using only the fixed change for clarity purposes.

At the top of the figure, we see the average number of generations that were necessary to solve a problem. It is clear to see that, in general, for our proposed approach (black-filled rectangle), it took a fewer number of generations to solve a function compared to an arbitrary model (white-filled rectangle). The standard deviation also indicates that there is less variation when using our approach compared to the arbitrary model.

If we, now, turn our attention to the average number of individuals created in the presence of a change, we can see the same tendency as before, our proposed approach created a fewer number of individuals (black-filled rectangle), and also, it shows less variance compared to the arbitrary model (white-filled rectangle). Notice that the lack of information in the first generations [0, 50) at the bottom of the figure is to be expected since a change takes place after this. Blanks in other parts at the bottom of this figure indicate that a problem was not solved in correspondence to what it is observed at the top of the figure.

D. Measuring Structural Changes

To further understand the impact that the arbitrary model and our proposed approach has in DOPs, we analyse how abrupt a change is, and to do so, we compute the tree-edit distance between the elitist individual and the target

expression, each of these shown in Table II. The tree-edit distance is integer-valued and reflective of a very intuitive notion of the distance between a pair of trees, based on the number of edits required to transform one into the other. Three types of edits are allowed: insertion, deletion, and substitution. This distance measure has been studied outside EAs, with algorithms given by [14] and used in GP e.g., [4].

This average tree-edit distance is shown in Figure 4. At the top of the figure we can see how abrupt a change is in the presence of a fixed change, where a change is depicted by a vertical line. Our proposed approach, indicated by a dashed line, is less abrupt in terms of a structural change compared to an arbitrary model, depicted by a solid line, except for the last problem (f_4 , g_4), where a structural change is more abrupt using our proposed approach compared to the arbitrary model. This can be explained if we consider what we discussed in the previous paragraphs regarding the number of individuals created by each of these two methods (see Figure 3). The same tendency can be observed in the presence of a random change, shown in the bottom of Figure 4, where a change is depicted by a vertical line.

VI. CONCLUSIONS

It is well-known that diversity helps evolutionary search in dynamic optimisation problems (DOPs). Many approaches have been proposed to promote diversity in DOPs using e.g., GP. In particular, the adoption of structural diversity has been popular given their e.g., positive impact and its simplicity to adopt it in a GP system. One particular element when promoting structural diversity in GP is the adoption of a rather arbitrary (random) model in terms of the number of individuals that need to be replaced by new created individuals. In this approach, we have proposed an approach based on the Kendall tau distance to make a more informed decision.

We have shown how the proposed approach is able to yield better results on the continuous fitness-value functions used in this work and used a variety of tools (e.g., tree-edit distance, number of created individuals, number of generations required to solve the problem) to explain why this approach behaves better than the arbitrary model commonly used in EAs, in general, and in GP, in particular.

To fully see the potential of the proposed approach is necessary to account for some of memory, either implicit or explicit, that has been reported to be necessary when tackling a DOP. It is also necessary to study the behaviour of our proposed approach in other more challenging problems (e.g., real-world DOPs [1], [3]).

ACKNOWLEDGEMENTS

EGL's research is funded by an ELEVATE Fellowship, the Irish Research Council's Career Development Fellowship co-funded by Marie Curie Actions. EGL would like to thank the TAO group at INRIA Saclay France for hosting him during the outgoing phase of the fellowship. The authors would like to thank the reviewers for their comments that helped us to improve our work.

REFERENCES

- [1] E. Galvan, C. Harris, I. Dusparic, S. Clarke, and V. Cahill. Reducing electricity costs in a dynamic pricing environment. In *Proc. Third IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 169 – 174, Tainan, Taiwan, november 2012. IEEE Press.
- [2] E. Galván-López, B. Cody-Kenny, L. Trujillo, and A. Kattan. Using semantics in the selection mechanism in genetic programming: A simple method for promoting semantic diversity. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2013, Cancun, Mexico, June 20-23, 2013*, pages 2972–2979. IEEE, 2013.
- [3] E. Galván-López, T. Curran, J. McDermott, and P. Carroll. Design of an autonomous intelligent demand-side management system using stochastic optimisation evolutionary algorithms. *Neurocomputing*, 170:270–285, 2015.
- [4] E. Galván-López, J. McDermott, M. O'Neill, and A. Brabazon. Defining locality as a problem difficulty measure in genetic programming. *Genetic Programming and Evolvable Machines*, 12(4):365–401, 2011.
- [5] E. Galván-López, E. Mezura-Montes, O. Ait ElHara, and M. Schoenauer. On the use of semantics in multi-objective genetic programming. In J. Handl et al., editors, *Parallel Problem Solving from Nature – PPSN XIV: 14th International Conference, Edinburgh, UK, September 17-21, 2016, Proceedings*, pages 353–363, Cham, 2016. Springer International Publishing.
- [6] T. Jones and S. Forrest. Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 184–192. Morgan Kaufmann, 1995.
- [7] J. R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA, 1992.
- [8] J. Macedo, E. Costa, and L. Marques. *Genetic Programming Algorithms for Dynamic Environments*, pages 280–295. Springer International Publishing, Cham, 2016.
- [9] M. Mavrovouniotis and S. Yang. Population-based incremental learning with immigrants schemes in changing environments. In *IEEE Symposium Series on Computational Intelligence, SSCI 2015, Cape Town, South Africa, December 7-10, 2015*, pages 1444–1451. IEEE, 2015.
- [10] M. Mavrovouniotis and S. Yang. Direct memory schemes for population-based incremental learning in cyclically changing environments. In G. Squillero and P. Burelli, editors, *Applications of Evolutionary Computation - 19th European Conference, EvoApplications 2016, Porto, Portugal, March 30 - April 1, 2016, Proceedings, Part II*, volume 9598 of *Lecture Notes in Computer Science*, pages 233–247. Springer, 2016.
- [11] N. F. McPhee, B. Ohs, and T. Hutchison. Semantic building blocks in genetic programming. In *Proceedings of the 11th European conference on Genetic programming, EuroGP'08*, pages 134–145, Berlin, Heidelberg, 2008. Springer-Verlag.
- [12] T. T. Nguyen, S. Yang, and J. Branke. Evolutionary dynamic optimization: A survey of the state of the art. *Swarm and Evolutionary Computation*, 6:1 – 24, 2012.
- [13] M. Rieckert, K. M. Malan, and A. P. Engelbrecht. Adaptive genetic programming for dynamic classification problems. In *Proceedings of the Eleventh Conference on Congress on Evolutionary Computation, CEC'09*, pages 674–681, Piscataway, NJ, USA, 2009. IEEE Press.
- [14] M. Tacker, F. P. Stadler, G. E. Bornberg-Bauer, L. I. Hofacker, and P. Schuster. Algorithm independent properties of rna secondary structure predictions. *European Biophysics Journal*, 25(2):115–130, 1996.
- [15] L. Vanneschi and G. Cuccu. A study of genetic programming variable population size for dynamic optimization problems. In *IJCCI*, pages 119–126, 2009.
- [16] N. Wagner, Z. Michalewicz, M. Khouja, and R. R. McGregor. Time series forecasting for dynamic environments: The dyfor genetic program model. *IEEE Transactions on Evolutionary Computation*, 11(4):433–452, Aug 2007.
- [17] W. Yan and C. D. Clack. Behavioural GP diversity for dynamic environments: an application in hedge fund investment. In M. Keijzer et al., editors, *GECCO 2006: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, volume 2, pages 1817–1824, Seattle, Washington, USA, 8-12 July 2006. ACM Press.
- [18] S. Yang. Evolutionary computation for dynamic optimization problems. In S. Silva and A. I. Esparcia-Alcázar, editors, *Genetic and Evolutionary Computation Conference, GECCO 2015, Madrid, Spain, July 11-15, 2015, Companion Material Proceedings*, pages 629–649. ACM, 2015.

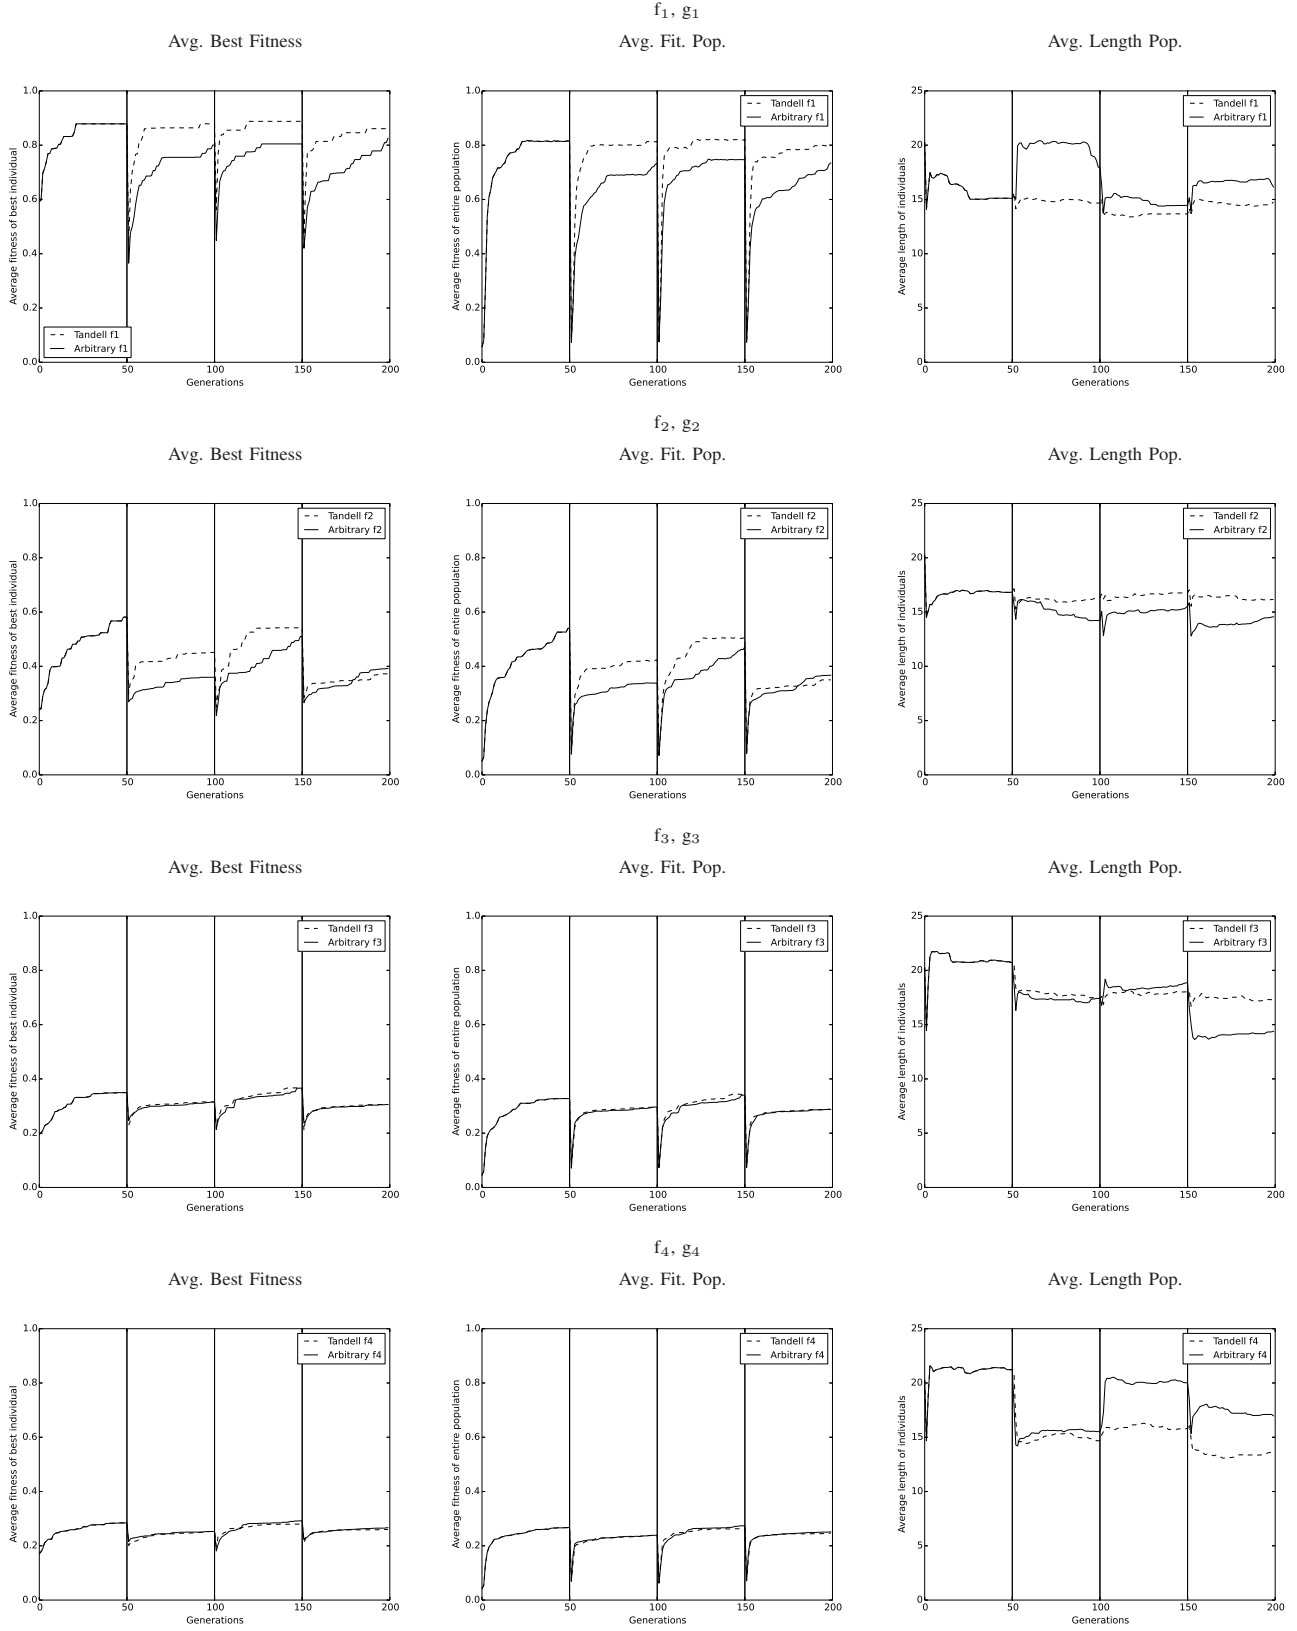


Fig. 1. Average best fitness, average fitness of entire population and average length of entire population computed when using either an arbitrary approach (indicated by a solid line) or the Kendall tau distance (indicated by a dashed line) shown in the left, centre and right-hand side of the figure, respectively, when using a fixed change, where a change is indicated by a vertical line. These are plotted for each of the four problems used in this work (see Table II).




Fig. 2. Average best fitness, average fitness of entire population and average length of entire population computed when using either an arbitrary approach (indicated by a solid line) or the Kendall tau distance (indicated by a dashed line) shown in the left, centre and right-hand side of the figure, respectively, when using a random change, where a change is indicated by a vertical line. These are plotted for each of the four problems used in this work (see Table II).

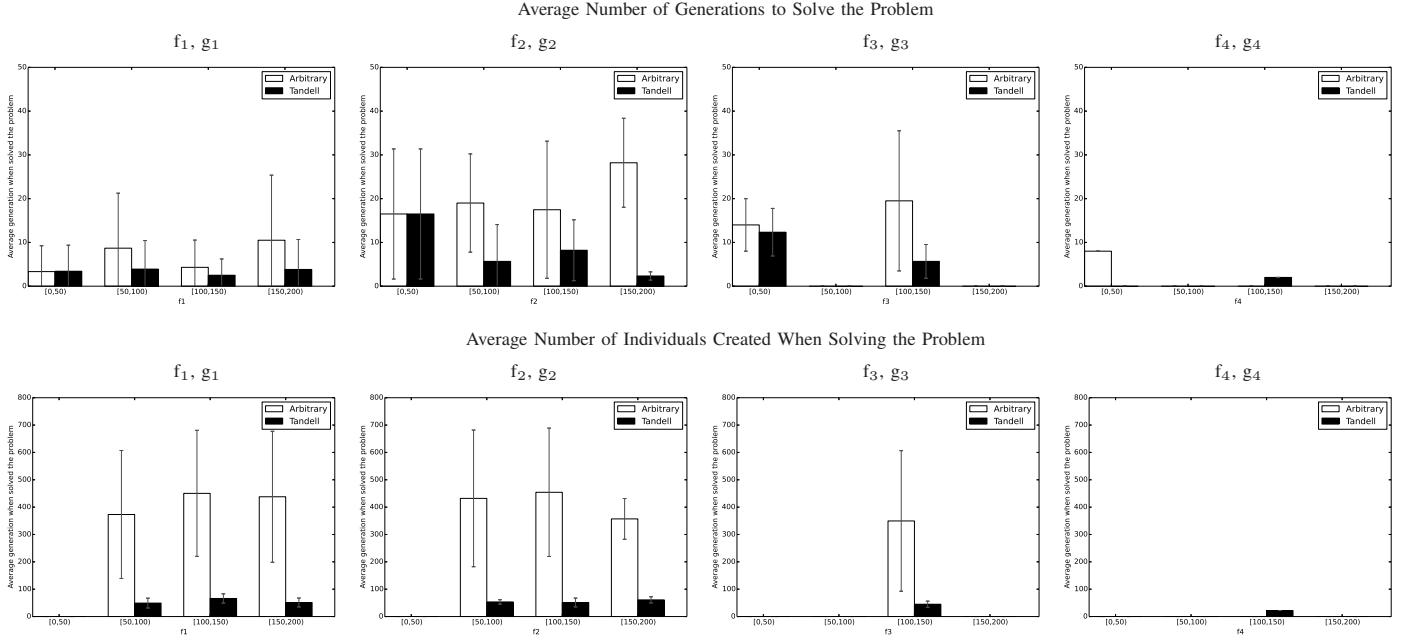


Fig. 3. Average number of generations to solve a problem (top) and average number of created individuals (bottom) along with standard deviation using either the arbitrary approach (white-filled rectangle) or our proposed Kendall tau distance approach (black-filled rectangle) for each of the problems used (see Table II) in the presence of a fixed change. Notice that for the average number of individuals created (bottom), the first generations [0,50), do not show anything given that a change occurs after this. Blanks in other parts at the bottom of this figure indicate that a problem was not solved in correspondence to what can be observed at the top of the figure.

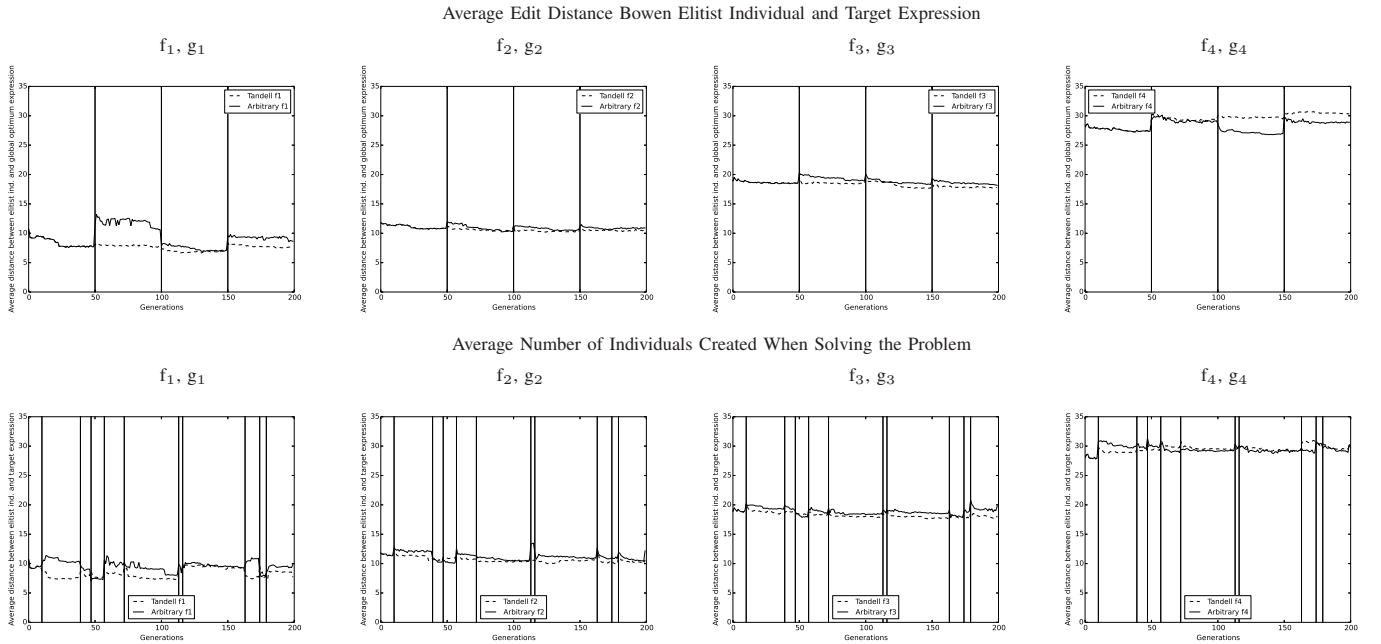


Fig. 4. Average edit distance between elitist individual and target expression calculated when using either the arbitrary approach (indicated by a solid line) or by our proposed approach (indicated by a dashed line), in the presence of a fixed change (top) or a random change (bottom). Vertical lines indicate when a change takes place.