Function Finding and the Creation of Numerical Constants in Gene Expression Programming

Cândida Ferreira

Gepsoft, 73 Elmtree Drive,
Bristol BS13 8NA, UK
candidaf@gepsoft.com
http://www.gene-expression-programming.com/author.asp

Gene expression programming is a genotype/phenotype system that evolves computer programs of different sizes and shapes (the phenotype) encoded in linear chromosomes of fixed length (the genotype). The chromosomes are composed of multiple genes, each gene encoding a smaller sub-program. Furthermore, the structural and functional organization of the linear chromosomes allows the unconstrained operation of important genetic operators such as mutation, transposition, and recombination. In this work, three function finding problems, including a high dimensional time series prediction task, are analyzed in an attempt to discuss the question of constant creation in evolutionary computation by comparing two different approaches to the problem of constant creation. The first algorithm involves a facility to manipulate random numerical constants, whereas the second finds the numerical constants on its own or invents new ways of representing them. The results presented here show that evolutionary algorithms perform considerably worse if numerical constants are explicitly used.

1. Introduction

Genetic programming (GP) evolves computer programs by genetically modifying nonlinear entities with different sizes and shapes (Koza 1992). These nonlinear entities can be represented as diagrams or trees. Gene expression programming (GEP) is an extension to GP that also evolves computer programs of different sizes and shapes, but the programs are encoded in a linear chromosome of fixed length (Ferreira 2001). One strength of the GEP approach is that the creation of genetic diversity is extremely simplified as genetic operators work at the chromosome level. Indeed, due to the structural organization of GEP chromosomes, the implementation of high-performing search operators is extremely simplified, as any modification made in the genome always results in valid programs. Another strength of GEP consists of its unique, multigenic nature which allows the evolution of complex programs composed of several simpler sub-programs.

It is assumed that the creation of floating-point constants is necessary to do symbolic regression in general (see, e.g., Banzhaf 1994 and Koza 1992). Genetic programming solved the problem of constant creation by using a special terminal named "ephemeral random constant" (Koza 1992). For each ephemeral random constant used in the trees of the initial population, a random number of a special data type in a specified range is generated. Then these random constants are moved around from tree to tree by the crossover operator.

Gene expression programming solves the problem of constant creation differently (Ferreira 2001). GEP uses an extra terminal "?" and an extra domain Dc composed of the symbols chosen to represent the random constants. For each gene, the random constants are generated during the inception of the initial population and kept in an array. The values of each random constant are only assigned during gene expression. Furthermore, a special operator is used to introduce genetic variation in the available pool of random constants by mutating the random constants directly. In addition, the usual operators of GEP plus a Dc specific transposition guarantee the effective circulation of the numerical constants in the population. Indeed, with this scheme of constants manipulation, the appropriate diversity of numerical constants can be generated at the beginning of a run and maintained easily afterwards by the genetic operators.

Notwithstanding, in this work it is shown that evolutionary algorithms do symbolic regression more efficiently if the problem of constant creation is handled by the algorithm itself. In other words, the special facilities for manipulating random constants are indeed unnecessary to solve problems of symbolic regression.

2. Genetic algorithms with tree representations

All genetic algorithms use populations of individuals, select individuals according to fitness, and introduce genetic variation using one or more genetic operators (see, e.g., Mitchell 1996). In recent years different systems have been developed so that this powerful algorithm inspired in natural evolution could be applied to a wide spectrum of problem domains (see, e.g., Mitchell 1996 for a review of recent work on genetic algorithms and Banzhaf *et al*. 1998 for a review of recent work on GP).

Structurally, genetic algorithms can be subdivided in three fundamental groups: i) Genetic algorithms with individuals consisting of linear chromosomes of fixed length devoid of complex expression. In these systems, replicators (chromosomes) survive by virtue of their own properties. The algorithm invented by Holland (1975) belongs to this group and is known as genetic algorithm or GA; ii) Genetic algorithms with individuals consisting of ramified structures of different sizes and shapes and, therefore, capable of assuming a richer number of functionalities. In these systems, replicators (ramified structures) also survive by virtue of their own properties. The algorithm invented by Cramer (1985) and later developed by Koza (1992) belongs to this group and is known as genetic programming or GP; iii) Genetic algorithms with individuals encoded in linear chromosomes of fixed length which are afterwards expressed as ramified structures of different sizes and shapes. In these systems, replicators (chromosomes) survive by virtue of causal effects on the phenotype (ramified structures). The algorithm invented by myself (Ferreira 2001) belongs to this group and is known as gene expression programming or GEP.

GEP shares with GP the same kind of ramified structure and, therefore, can be applied to the same problem domains. However, the logistics of both systems differ significantly and the existence of a real genotype in GEP allows the unprecedented manipulation and exploration of more complex systems. Below are briefly highlighted some of the differences between GEP and GP.

2.1. Genetic programming

As simple replicators, the ramified structures of GP are tied up in their own complexity: on the one hand, bigger, more complex structures are more difficult to handle and, on the other, the introduction of genetic variation can only be done at the tree level and, therefore, must be done carefully so that valid structures are created. A special kind of tree crossover is practically the only source of genetic variation used in GP for it allows the exchange of sub-trees and, therefore, always produces valid structures. Indeed, the implementation of high-performing operators, like the equivalent of natural point mutation, is unproductive as most mutations would have resulted in syntactically invalid structures. Understandingly, the other genetic operators described by Koza (1992) – mutation and permutation – also operate at the tree level.

2.2. Gene expression programming

The phenotype of GEP individuals consists of the same kind of diagram representation used by GP. However, these complex phenotypes are encoded in simpler, linear structures of fixed length – the chromosomes. Thus, the main players in GEP are the chromosomes and the ramified structures or expression trees (ETs), the latter being the expression of the genetic information encoded in the former. The decoding of GEP genes implies obviously a kind of code and a set of rules. The genetic code is very simple: a one-to-one relationship between the symbols of the chromosome and the functions or terminals they represent. The rules are also very simple: they determine the spatial organization of the functions and terminals in the ETs and the type of interaction between sub-ETs in multigenic systems.

In GEP there are therefore two languages: the language of the genes and the language of ETs. However, thanks to the simple rules that determine the structure of ETs and their interactions, it is possible to infer immediately the phenotype given the sequence of the genotype, and *vice versa*. This bilingual and unequivocal system is called *Karva* language. The details of this new language are given in Ferreira (2001).

3. Two approaches to the problem of constant creation

In this section the problem of constant creation is discussed by comparing the performance of two different algorithms. The first manipulates explicitly the numerical constants and the second solves the problem of constant creation in symbolic regression by creating constants from scratch or by inventing new ways of representing them.

3.1. Setting the system

The comparison between the two approaches (with and without the facility to manipulate random constants) was made on three different problems. The first is a problem of sequence induction requiring integer constants. In this case the following test sequence was chosen:

$$a_n = 5n^4 + 4n^3 + 3n^2 + 2n + 1 \qquad (3.1)$$

where n consists of the nonnegative integers. This sequence was chosen because it can be exactly solved and therefore can provide an accurate measure of performance in terms of success rate.

The second is a problem of function finding requiring floating-point constants. In this case, the following "V" shaped function was chosen:

$$y = 4.251a^2 + \ln(a^2) + 7.243e^a \qquad (3.2)$$

where a is the independent variable and e is the irrational number 2.71828183. Problems of this kind cannot be exactly solved by evolutionary algorithms and, therefore, the perform-

ance of both approaches is compared in terms of average best-of-run fitness and average best-of-run R-square.

The third is the well-studied benchmark problem of predicting sunspots (Weigend *et al*. 1992). In this case, 100 observations of the Wolfer sunspots series were used (Table 1) with an embedding dimension of 10 and a delay time of one. Again, the performance of both approaches is compared in terms of average best-of-run fitness and R-square.

Table 1
Wolfer sunspots series (read by rows).

101	82	66	35	31	7	20	92
154	125	85	68	38	23	10	24
83	132	131	118	90	67	60	47
41	21	16	6	4	7	14	34
45	43	48	42	28	10	8	2
0	1	5	12	14	35	46	41
30	24	16	7	4	2	8	17
36	50	62	67	71	48	28	8
13	57	122	138	103	86	63	37
24	11	15	40	62	98	124	96
66	64	54	39	21	7	4	23
55	94	96	77	59	44	47	30
16	7	37	74				

For the sequence induction problem, the first 10 positive integers n and their corresponding term a_n were used as fitness cases. The fitness function was based on the relative error with a selection range of 20% and maximum precision (0% error), giving maximum fitness $f_{max} = 200$ (Ferreira 2001).

For the "V" shaped function problem, a set of 20 random fitness cases chosen from the interval [-1, 1] was used. The fitness function used was also based on the relative error but in this case a selection range of 100% was used, giving $f_{max} = 2,000$.

For the time series prediction problem, using an embedding dimension of 10 and a delay time of one, the sunspots series presented in Table 1 result in 90 fitness cases. In this case, a wider selection range of 1,000% was chosen, giving $f_{max} = 90,000$.

In all the experiments, the selection was made by roulette-wheel sampling coupled with simple elitism and the performance was evaluated over 100 independent runs. The six experiments are summarized in Table 2.

3.2. First approach: Direct manipulation of numerical constants

To solve the sequence induction problem using random numerical constants, F = {+, -, *}, T = {a, ?}, the set of integer random constants R = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9}, and "?" ranged over the integers 0, 1, 2, and 3. The parameters used per run are shown in the first column of Table 2. In this experiment, the first perfect solution was found in generation

45 of run 9 (the contribution of each sub-ET is indicated in square brackets):

$$y = [a^2] + [a] + [2a^4 + 4a^3] +$$
$$+ [0] + [2a^2] + [1 + a] + [3a^4] \qquad (3.3)$$

which corresponds to the target sequence (3.1).

As shown in the first column of Table 2, the probability of success for this problem is 16%, considerably lower than the 81% of the second approach (see Table 2, column 2). It is worth emphasizing that only the prior knowledge of the solution enabled us, in this case, to choose correctly the type and the range of the random constants.

To find the "V" shaped function using random constants F = {+, -, *, /, L, E, K, ~, S, C} ("L" represents the natural logarithm, "E" represents e^x, "K" represents the logarithm of base 10, "~" represents 10^x, "S" represents the sine function, and "C" represents the cosine) and T = {a, ?}. The set of rational random constants R = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9}, and "?" ranged over the interval [-1, 1]. The parameters used per run are shown in the third column of Table 2. The best solution, found in run 50 after 4584 generations, is shown below (the contribution of each sub-ET is indicated in square brackets):

$$y = \left[\ln\left(0.99782a^2\right)\right] + \left[10^{\sin(1.27278a)}\right] +$$
$$+ \left[10^{0.929a}\right] + \left[0.77631 - 2.80112a^3\right] +$$
$$+ \left[2.45714 + e^{0.981a} + e^a\right] \qquad (3.4)$$

It has a fitness of 1989.566 and an R-square of 0.9997001 evaluated over the set of 20 fitness cases and an R-square of 0.9997185 evaluated against a testing set of 100 random points also chosen from the interval [-1, 1].

It is worth noticing that the algorithm does in fact integrate constants in the evolved solutions, but the constants are very different from the expected ones. Indeed, GEP (and I believe, all genetic algorithms with tree representations) can find the expected constants with a precision to the third or fourth decimal place when the target functions are simple polynomial functions with rational coefficients and/or when it is possible to guess pretty accurately the function set, otherwise a very creative solution would be found.

To predict sunspots using random numerical constants, the set of functions F = {4+, 4-, 4*, 4/} and the set of terminals T = {a, b, c, d, e, f, g, h, i, j, ?}. The set of rational random constants R = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9}, and "?" ranged over the interval [-1, 1]. The parameters used per run are shown in the fifth column of Table 2. The best solution, found in run 92 after 4759 generations, is shown below:

$$y = \frac{2j^2}{h + i + j} + \frac{a + b + g}{0.995 + 0.847c + e} + \frac{1.903j + j^2}{i^2 + j} \qquad (3.5)$$

Table 2
General settings used in the sequence induction (SI), the "V" function, and sunspots (SS) problems. The "*" indicates the explicit use of random constants.

	SI*	SI	V*	V	SS*	SS
Number of runs	100	100	100	100	100	100
Number of generations	100	100	5000	5000	5000	5000
Population size	100	100	100	100	100	100
Number of fitness cases	10	10	20	20	90	90
Function set	+ - * /	+ - * /	+ - * / L E K ~ S C	+ - * / L E K ~ S C	4 (+ - * /)	4 (+ - * /)
Terminal set	a, ?	a	a, ?	a	a - j, ?	a - j
Random constants array length	10	--	10	--	10	--
Random constants range	{0, 1, 2, 3}	--	[-1,1]	--	[-1,1]	--
Head length	6	6	6	6	8	8
Number of genes	7	7	5	5	3	3
Linking function	+	+	+	+	+	+
Chromosome length	140	91	100	65	78	51
Mutation rate	0.044	0.044	0.044	0.044	0.044	0.044
One-point recombination rate	0.3	0.3	0.3	0.3	0.3	0.3
Two-point recombination rate	0.3	0.3	0.3	0.3	0.3	0.3
Gene recombination rate	0.1	0.1	0.1	0.1	0.1	0.1
IS transposition rate	0.1	0.1	0.1	0.1	0.1	0.1
IS elements length	1,2,3	1,2,3	1,2,3	1,2,3	1,2,3	1,2,3
RIS transposition rate	0.1	0.1	0.1	0.1	0.1	0.1
RIS elements length	1,2,3	1,2,3	1,2,3	1,2,3	1,2,3	1,2,3
Gene transposition rate	0.1	0.1	0.1	0.1	0.1	0.1
Random constants mutation rate	0.01	--	0.01	--	0.01	--
Dc specific transposition rate	0.1	--	0.1	--	0.1	--
Dc specific IS elements length	1,2,3	--	1,2,3	--	1,2,3	--
Selection range	20%	20%	100%	100%	1000%	1000%
Precision	0%	0%	0%	0%	0%	0%
Average best-of-run fitness	179.827	197.232	1914.8	1931.84	86215.27	89033.29
Average best-of-run R-square	0.977612	0.999345	0.957255	0.995340	0.713365	0.811863
Success rate	16%	81%	--	--	--	--

It has a fitness of 86603.2 and an R-square of 0.833714 evaluated over the set of 90 fitness cases.

3.3. Second approach: Creation of numerical constants from scratch

To solve the sequence induction problem without the facility to manipulate numerical constants, the function set was exactly the same as in the experiment with random constants. The terminal set consisted obviously of the independent variable alone.

As shown in the second column of Table 2, the probability of success using this approach is 81%, considerably higher than the 16% obtained using the facility to manipulate random constants. In this experiment, the first perfect solution was found in generation 44 of run 0 (the contribution of each sub-ET is indicated in square brackets):

$$y = [2a] + [a^3 + a] + [a^4 + 3a^3 + 2a^2] +$$
$$+ [4a^4] + [1] + [a^2 - a] + [0] \qquad (3.6)$$

which is equivalent to the target sequence (3.1). In this case

the algorithm created all the necessary constants from scratch by performing simple mathematical operations.

To find the "V" shaped function without using random constants, the function set is exactly the same as in the first approach. With this collection of functions, most of which extraneous, the algorithm is equipped with different tools for evolving highly accurate models without using numerical constants. The parameters used per run are shown in the fourth column of Table 2. In this experiment of 100 identical runs, the best solution was found in generation 4679 of run 10 (the contribution of each sub-ET is indicated in square brackets):

$$y = \left[\ln(2a^2) + 10^{\sin a}\right] + \left[2a + \sin a + a^2\right] +$$
$$+ \left[\cos(\cos(2a)) + e^{a^2}\right] + \left[e^{\sin a}\right] + \left[1 + e^a + e^{a^2}\right] \qquad (3.7)$$

It has a fitness of 1990.023 and an R-square of 0.9999313 evaluated over the set of 20 fitness cases and an R-square of 0.9998606 evaluated against the same testing set used in the first approach, and thus is better than the model (3.4) evolved with the facility for the manipulation of random constants.

To predict sunspots without using random numerical con-

stants, the function set is exactly the same as in the first approach. The parameters used per run are shown in the sixth column of Table 2. In this experiment of 100 identical runs, the best solution was found in generation 2273 of run 57:

$$y = j + \frac{d - i + 3j}{b + e} + \frac{d + bj - ij}{a + 2i} \qquad (3.8)$$

It has a fitness of 89176.61 and an R-square of 0.882831 evaluated over the set of 90 fitness cases, and thus is better than the model (3.5) evolved with the facility for the manipulation of random constants.

It is instructive to compare the results obtained in both approaches. In all the experiments the explicit use of random constants resulted in a worse performance. In the sequence induction problem, success rates of 81% against 16% were obtained; in the "V" function problem average best-of-run fitnesses of 1931.84 versus 1914.80 and average best-of-run R-squares of 0.995340 versus 0.957255 were obtained; and in the sunspots prediction problem average best-of-run fitnesses of 89033.29 versus 86215.27 and average best-of-run R-squares of 0.811863 versus 0.713365 were obtained (see Table 2). Thus, in real-world applications where complex realities are modeled, of which nothing is known concerning neither the type nor the range of the numerical constants, and where most of the times it is impossible to guess the exact function set, it is more appropriate to let the system model the reality on its own without explicitly using random constants. Not only the results will be better but also the complexity of the system will be much smaller.

4. Conclusions

Gene expression programming is the most recent development on artificial evolutionary systems and one that brings about a considerable increase in performance due to the crossing of the phenotype threshold. In practical terms, the crossing of the phenotype threshold allows the unconstrained exploration of the search space because all modifications are made on the genome and because all modifications always result in valid phenotypes or programs. In addition, the genotype/phenotype representation of GEP not only simplifies but also invites the creation of more complexity. The elegant mechanism developed to deal with random constants is a good example of this.

In this work, the question of constant creation in symbolic regression was discussed comparing two different approaches to solve this problem: one with the explicit use of numerical constants, and another without them. The results presented here suggest that the latter is more efficient, not only in terms of the accuracy of the best evolved models and overall performance, but also because the search space is much smaller, reducing greatly the complexity of the system and, consequently, the precious CPU time.

Finally, the results presented in this work also suggest that, apparently, the term "constant" is just another word for mathematical expression and that evolutionary algorithms are particularly good at finding these expressions because the search is totally unbiased.

Bibliography

Banzhaf, W., Genotype-Phenotype-Mapping and Neutral Variation – A Case Study in Genetic Programming. In Y. Davidor, H.-P. Schwefel, and R. Männer, eds., *Parallel Problem Solving from Nature III*, *Lecture Notes in Computer Science,* 866: 322-332, Springer-Verlag, 1994.

Banzhaf, W., P. Nordin, R. E. Keller, and F. D. Francone, *Genetic Programming: An Introduction: On the Automatic Evolution of Computer Programs and its Applications*. Morgan Kaufmann, 1998.

Cramer, N. L., A Representation for the Adaptive Generation of Simple Sequential Programs. In J. J. Grefenstette, ed., *Proceedings of the First International Conference on Genetic Algorithms and Their Applications*, Erlbaum, 1985.

Ferreira, C., 2001. Gene Expression Programming: A New Adaptive Algorithm for Solving Problems. *Complex Systems*, 13 (2): 87-129.

Holland, J. H., *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. University of Michigan Press, 1975 (second edition: MIT Press, 1992).

Koza, J. R., *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA, MIT Press, 1992.

Mitchell, M., *An Introduction to Genetic Algorithms*. MIT Press, 1996.

Weigend, A.S., B. A. Huberman, and D. E. Rumelhart, Predicting Sunspots and Exchange Rates with Connectionist Networks. In S. Eubank and M. Casdagli, eds., *Nonlinear Modeling and Forecasting*, pages 395-432, Redwood City, CA, Addison-Wesley, 1992.