

Compressed Linear Genetic Programming: empirical parameter study on the Even-n-parity problem

Johan Parent

Vrije Universiteit Brussel
IR, ETRO
Pleinlaan 2 1050 Brussel,
BELGIUM
johan@info.vub.ac.be

Ann Nowé

Vrije Universiteit Brussel
Faculty of Science, COMO
Pleinlaan 2 1050 Brussel,
BELGIUM
asnowe@info.vub.ac.be

Anne Defaweux

Vrije Universiteit Brussel
Faculty of Science, COMO
Pleinlaan 2 1050 Brussel,
BELGIUM
adefaweu@vub.ac.be

Kris Steenhaut

Vrije Universiteit Brussel
IR, ETRO
Pleinlaan 2 1050 Brussel,
BELGIUM
kris@info.vub.ac.be

Abstract- This paper presents a parameter study of our Compressed Linear Genetic Programming (cl-GP) using the Even-n-parity problem. A cl-GP system is a linear genetic programming (GP) which uses substring compression as a modularization scheme. Despite the fact that the compression of substrings assumes a tight linkage between alleles, this approach improves the search process. The compression of the genotype, which is a form of linkage learning, provides both a protection mechanism and a form of genetic code reuse. This text presents a study of the different parameters of the cl-GP on Even-n-parity. Experiments indicate that the cl-GP performs best when compressing a small fraction of the population and the length of the substituted substrings is rather short.

1 Introduction

Genetic algorithms (GA) and genetic programming (GP) search the space of possible solutions by manipulating the solution representation using genetic operators like crossover and mutation. Variable length representations allow the structure of the solution to evolve but are still restricted to the genetic primitives used to build the solutions. If GA and GP are to address more complex problems it becomes necessary to automatically adapt the representation to the problem. Modularization adapts the representation by extending the set of genetic primitives with problem specific functionalities. In this text we present a low level modularization technique for linear GP system based on compression. The presented compressed linear GP (cl-GP) attempts to identify usefull combinations of genes in the population and makes these available as new primitives. The search process benefit from these new functionalities to progress more quickly through the search space.

The assumption of tight linkage between individual genes (cfr substrings) was made with the use of a linear representation in mind. The representation of a linear GP system is similar to that used by a GA [4][11]. Following the approach used by [11] the cl-GP loop consists of a standard (generational) GA on top of which a problem specific evaluator is placed to simulate the program execution.

This paper is structured as follows Section 2 describes related work, In Section 3 the compression/substitution scheme is presented in detail. In Section 4 the experiments used to evaluate the impact different parameter settings of our algorithm is presented. Results are presented in Section

5.

2 Related work

The benefits of modularization for the GP search process are well known [9][1][7]. Modularization fosters code reuse on one hand, and on the other hand allows the GP system to identify and use high level functionalities. Different modularization strategies for tree based GP have been explored. Automatically Defined Functions (ADF) is the most prominent approach. It consists of a main tree which evolves together with a predefined number of additional trees. Those additional trees can be called from the main tree and, as such, complement the set of primitives available to the GP system. An alternative method is *encapsulation*. It replaces a subtree by a new symbol, this symbol is added to the primitive set of the system [6][9]. The symbol created in this way corresponds to a terminal/leaf node as it does not have any arguments. A third method, *module acquisition*, works in a similar fashion, but can create both function nodes (modules) and leaf nodes. If the depth of the selected subtree exceeds a certain threshold module acquisition removes the subtrees below this level. In this case a function node will be created by adding an argument for each removed subtree [2][3]. Although modularization has mainly been of interest to the GP community it is of course related to the search of building blocks in a GA. In [5] a *module acquisition algorithm* is presented which exploits modularity and hierarchy. Despite the use of a variable length representation, the problem of modularity is approached purely from a GA point of view. A *module* is defined as a combination of gene values that maximize the fitness. For example, if for the genes g_1 and g_2 the combination $g_1 = v_8$ and $g_2 = v_3$ dominates all other combinations of genes value, then it is considered a module. As in the work of [5] our algorithm relies on a substitution scheme but with notable differences in the substitution strategy. Another difference is that the compressed genotype scheme of the cl-GP was imagined to be used in a linear GP system. The linear encoding of a GP program exhibits a much tighter linkage than is the case for standard GA problems. This assumption is important as it justifies the use of the simple compression scheme which supposes a strong dependency between adjacent values.

```

0. Choose initial population
1. Evaluate each individual's fitness
2. Repeat
3.     Compress individuals
4.     Select best-ranking individuals to reproduce
5.     Mate pairs at random
6.     Apply crossover operator
7.     Apply mutation operator
8.     Decompress individuals
9.     Evaluate each individual's fitness
10. Until terminating condition

```

Figure 1: The pseudo code of a standard GA with 2 extra steps. Step 3 adds the compression of the individual prior to the creation of the next generation. Step 8 decompresses the individuals so that their fitness can be evaluated.

3 Compressed Linear GP

Our cl-GP algorithm consist of the compressed genetic algorithm (cGA) together with problem specific evaluator. This section presents the cGA which consists of a GA extended with a substitution/compression based modularization mechanism. Figure 1 contains the pseudo code for cGA loop. The cGA adds compression and decompression steps in the GA loop. Compression is applied to the genotype of individuals in the population prior to the creation of the next generation. As a result the search process occurs using a compressed representation. After the creation of the next generation the individuals are decompressed so that their fitness can be evaluated.

The schema theory models the probability of disruption by crossover is proportional to the defining length of the building block. Compression can shorten the representation of the building blocks thus reduce the chance of disruption by crossover. The cGA uses a substitution coder to compress the genotype in an attempt to protect substrings that represent building blocks. Protecting the building blocks guarantees the preservation of good allele combinations and is beneficial to the search process. In Subsection 3.1 the basics of the substitution coder used by the cGA are presented. The compression of the individuals is a two stage stochastic process: build a dictionary and select individuals for compression. Subsection 3.2 explains how the dictionary used for the compression is built. Subsection 3.4 describes how compression is applied to the population of the cGA.

3.1 Substitution coder

A substitution coder is a lossless¹ compression algorithm. Compression is obtained by replacing substrings in the input by a shorter reference. A substitution coder uses a *dictionary* which contains the substrings that need to be substituted and associates a placeholder symbol with each entry. Compression involves two phases: 1) building a dictionary and 2) performing the substitutions. Figure 2 represents the

¹Lossless compression means that the original data is restored after compression.

```

0.  $\mathcal{D} = \text{build\_dictionary}()$ ;
1.  $\text{current} = \text{in}$ ; /* in input string */
2.  $\text{out} = ""$ ; /* Empty output */
3. For  $e$  in  $\mathcal{D}$  /* Loop A */
4.   For  $\text{pos} = 1$  to  $|\text{in}|$  /* Loop B */
5.     If  $\text{match}(\text{pos}, \text{current}, e.\text{str})$  Then
6.        $\text{out}[\text{pos}] = \text{current}[\text{pos}]$ ;
7.     Else
8.        $\text{out}[\text{pos}] = e.\text{ref}$ ;
9.        $\text{pos} = \text{pos} + |e.\text{str}|$ ;
10.    End;
11.  End
12.  $\text{current} = \text{out}$ ;
13.End

```

Figure 2: The pseudo code for the substitution step of the coder used in the cGA. Given are the dictionary \mathcal{D} and the input string in . The coder searches for a match for each dictionary entry (loop A). If a match is found the reference is placed in the output, otherwise the original symbol. The inner loop (loop B) can be replaced by the more efficient KMP string matching algorithm with a linear runtime complexity.

pseudo code of a substitution coder. The content of the dictionary is critical for the compression performance. Much of the research in data compression concerns the development of algorithms to build the dictionary. Section 3.2 describes how the dictionary is built in the cGA. The substitution step is computationally expensive as it involves searching for a match for each dictionary entry in the input string. If a match for substring s is found its placeholder symbol is placed in the output instead of the original symbols. The decompression step involves the replacing the placeholder symbols by the original strings. Decompression is much faster since it does not involve any search.

Suppose a dictionary $\mathcal{D} = \{ "101" \rightarrow \alpha, "00" \rightarrow \beta, "11" \rightarrow \gamma \}$ then a substitution coder would compress the individual "1010001011101" to " $\alpha\beta0\alpha1\alpha$ ". One can observe that the order in which the dictionary entries are ordered is important. If the opposite order had been used the result would have been $\{\alpha\beta010\gamma\alpha\}$.

3.2 Building the dictionary

This section explains the stochastic algorithm used in the cGA to build the dictionary as it differs from the algorithm(s) used for *pure* data compression applications. We considered other criteria to build the dictionary since reducing the memory needed to represent the individuals is not our goal. We seek to build a dictionary containing building blocks. The problem of identifying building blocks is recurrent for all the modularization algorithms. In [5] the identification of a set of alleles as a module involves additional fitness evaluations. These are used to determine whether an alternative substring with a better fitness exists. If no such string can be found then the substring will be used as a module. Another approach, used by [10], uses a separate *block fitness function* to evaluate the merits of a subpart of a genetic program. This function is presented as scaled down

version of the fitness function for a smaller version of the original problem. The strategy adopted in this work is to see a GA as the building block discovery process. Since by definition building blocks are contributing in a positive way to the fitness of an individual, they are bound to be present in the better individuals of a population. We believe that this approach is more general as it 1) relies on the information already present in the population, 2) avoids additional computations and 3) does not require additional fitness functions to be engineered. Especially the latter two unacceptable when it comes to real world applications.

The cGA builds its dictionary in two stages. First, a pool of M individuals is selected stochastically from the population. The genotype of these individual will be used to create the dictionary in the second stage. These individuals are selected using fitness proportional selection. The pool consist mainly of above average individuals, yet with a minimum of diversity. Once the pool has been created every substring of length l of each individual is added to the dictionary. The dictionary only contains strings of the same length². For example, suppose a non-binary alphabet $\{a, b, c, d\}$ and the substring length l equal to 3. In this case the individual "abccdac" would add the substring (and their respective placeholder) "abc" $\rightarrow \alpha$, "bcc" $\rightarrow \beta$, "ccd" $\rightarrow \gamma$, "cda" $\rightarrow \delta$ and "dab" $\rightarrow \sigma$ to the dictionary. The cGA dictionary contains next to the substrings and their placeholder symbol a counter of the occurrence of every substring in the pool. This counter is used to order entries of the dictionary. We have chosen to sort the dictionary entries based on the occurrence of each substring, the most frequently occurring substring(s) will be substituted first.

The cGA rebuilds the dictionary for each generation. This allows to update the content of the dictionary with information that reflects the evolution of the population. This differs from [5] where the *module formation algorithm* is applied periodically. This also excludes the presence of dictionary entries containing references to other entries.

3.3 Modularization

Applying compression to the genotype can protect building block. But it does not by itself provide modularization comparable to that of ADFs, encapsulation or module acquisition. Modularization implies that a somehow code reuse should be present. Code reuse is achieved by the cGA by adhering to several conditions. First of all, the representation used by the cGA is adapted during the search. A new symbol is added to the alphabet of the genetic system for each potential building block identified by the cGA. Second, the genetic operator must be unrestricted: no distinction is made between compressed symbols and symbols of the original alphabet. Mutation for instance can replace a compressed symbol by an *original* symbol and vice versa. The third condition is that no distinction is made between individuals with and without compressed genotype. As will be explained in section 3.4, all the individuals are not necessarily subjected to compression. By not discriminating

between original genes and *compressed genes*, the cGA can make full use of the genetic combinations that were identified as possible building blocks. It then becomes possible for the cGA to address problems where repetition is present. Crossover and mutation operators serve as natural mechanism to obtain this form of *translocation* of genetic information. Translocation occurs when a part of a chromosome is detached and reattached at a position different from the its original position in the chromosome.

3.4 Compressing the population

Once the dictionary has been built the substitution coder (Subsection 3.1) can be used to compress the individuals in the population. The cGA applies compression in a stochastic fashion to a part of the population. A fraction κ ($\in [0, 1]$) of the population will be compressed. This fraction κ of individuals are selected using tournament selection. Since the compression setup is repeated anew for each generation this means that the cGA does not systematically compress the same individuals. This makes it possible to explore the benefit of the different modules (dictionary entries) in different contexts (ie. allele combinations).

3.5 Lossless and stochastic

The cGA applies a deterministic transformation, compression, to some of the individuals in the population. Since the compression is lossless the (genetic) information present in the population of the cGA or a GA is exactly the same. Yet, as described above, a stochastic component has been added to the overall population compression process. The creation of the dictionary and the selection of the individuals for compression are non-deterministic processes. This stochastic component is meant to counter balance the effect of the substitutions on the search process. The protection against crossover provided by the compression has a negative impact on the population diversity. Several factors explain this phenomenon. First the search efficiency of the crossover operator is reduced by the compression of the genotype. Second, the compression of the individuals is detrimental for the sampling of schemata. A last reason is that the dictionary entries are not guaranteed to correspond to building blocks. This situation is especially exacerbated during the first generation as the population still needs to discover promising gene combinations. This was illustrated by experiments where the entire population underwent compression. These experiments have invariably lead to premature convergence and consequently suboptimal results. The parameter κ is a way to limit the decrease in the population diversity. The other non-deterministic steps were introduced for the same reason as they maintain a minimum of diversity at the substring level.

3.6 Further changes

Although the cGA is meant to be a minimal extension to the standard GA a few extra modification were required. The biggest difference is an unavoidable transition to a variable

²In traditional substitution coders the dictionary can contain entries of different lengths

length linear representation. This change is due to the compression scheme.

3.6.1 Variable length

Despite the use of a lossless compression algorithm and the fact that the cGA, like a standard GA, starts with a population of individuals of identical size, individuals of different sizes quickly emerge. This phenomenon, due to the use of compression, occurs through several mechanisms:

1. not all the individual are compressed
2. different individuals compress to different sizes
3. the genetic operators can create individuals of very different sizes

Since the cGA does not discriminate between placeholder symbols and *original* symbols, recombination and mutation can add or remove symbols representing a substring and vice versa. This can result in an important change in length. Depending on the situation individuals can exceed the initial individual size or be shorter. Consider the following example. Suppose the dictionary $\mathcal{D} = \{ "xy" \rightarrow \alpha, "zz" \rightarrow \beta \}$ over the alphabet $\{x, y, z\}$. Mutating the second gene of the individual with genotype $"xx"$ from x to α creates $"x\alpha"$. The decompressed version of this individual's genotype is now $"xxy"$. Similar scenarios exists for the crossover.

3.6.2 Recombination and mutation

As illustrated above the genetic operators can create individuals of very different sizes. As a result a maximum individual length has to be enforced to avoid bloat. Individuals created by crossover that exceed the maximum length are truncated. Yet in the case of both mutation and crossover one problem remains. As illustrated in the previous paragraph it is possible for an individual to exceed the maximum once it has been decompressed. In this case the cGA truncates the genotype after decompression. The mutation operator was modified to be able to cope with a variable alphabet sizes.

4 Experiment

The advantage of the modularization of the cl-GP, provided by the cGA, has been illustrated in different GP problems: symbolic regression, classification and a real world data application. This section presents the experiments used to study the impact of the parameters of the cl-GP. We used the Even- n -parity problem for our different experiments. It has been chosen since it is a standard GP problem. Furthermore the difficulty of this problem can be adjusted by modifying the number of input bits n . The cl-GP introduces three new parameters: the size of the pool used to build the dictionary, κ the fraction of the population being compressed and l , the length of the dictionary entries.

The number of individuals used to build the dictionary should influence the quality of the dictionary entries. It is expected at a large pool performs better than a small one.

κ controls the population compression. Compressing a big fraction of the population is expected to have a negative influence on the search. The protection against crossover, the result of the cl-GP compression, reduces the number of possible individuals create by the crossover thus the population diversity.

The last parameter is the length of the dictionary entries. The cl-GP assumes there is a tight linkage in the representation. Using large substrings as dictionary entries is expected to reduce the performance. This is due to the increasing difficulty of identifying substring which corresponds to a schema of length l as the value of l increases.

4.1 Even- n -parity

The Even- n -Parity problem requires the correct classification of bit strings of length n having an even number of 1's. This classification is formulated as a boolean function returning the value *true* for an even number of 1's and *false* otherwise. The terminals and function set are $T = \{b_0, b_1, \dots, b_n\}$ and $F = \{NOOP, AND, OR, NAND, NOR\}$. The NOOP instruction is not executed during the evolution. This instruction allows to represent programmas of variable length using a fixed length representation [11]. We believe the inclusion of this instruction is necessary in order to compare the l-GP, driven by a GA (fixed length), and the cl-GP driven by the cGA (variable length). The raw fitness (f_{raw}) is the ratio of correct classifications over the entire data set: $f_{raw} = 1 - \frac{\#errors}{2^n}$. The standard fitness is equal to the raw fitness for this problem. The evaluator used for this problem is described in the next section.

4.1.1 Boolean evaluator

The Even- n -parity individuals are postfix encoded and evaluated by a stack based virtual machine as described in [8]. Its instruction set provides the four boolean operators³ and five terminals. The operations take their 2 operands from the stack and push their result onto it. The terminal instructions correspond to a push of the individual input bits on the stack. As for the numerical evaluator the result of the program is the value of the top of the stack after evaluation. The individual $[b_0, NOR, NAND, b_0, b_4, OR, AND, b_2, OR, b_2, NOOP, NOR, NAND]$ corresponds to the boolean expression $(b_0 OR b_4) NAND (b_2 NOR b_2)$.

5 Results

All the results presented here are the average of 100 independent runs, using randomly seeded initial populations of 500 individuals. The genotype length is 32. Other settings were: crossover rate 80%, mutation rate 5% and the top 5% of the population was kept at every generation. Every experiment lasted 50 generations. If not mentioned otherwise the following values were used for the runs using the cl-GP:

³Following the observations of [8] the lazy version of the operators was implemented.




Figure 3: An example run for the Even-5-parity problem for population size of 100 and 500. The linear GP system driven by the cGA (cl-GP) outperforms the same system using the GA.

$\kappa = 0.30$ (tournament size 4) and the length of the dictionary entries equals two. A pool of 10 individuals (fitness proportional selected) is used to build the dictionary. Before presenting the results for the different experiments, the next subsection illustrates the benefit of using the cl-GP for two Even- n -parity problem instances. It should be noted that large problem instances ($n \geq 10$) could not be solved with the given genetic primitives.

5.1 cl-GP (cGA) vs l-GP (GA)

This subsection compares the results for the Even-5-parity problem instance using runs with 2 different population sizes. The use of substitution has a positive effect on the performance of the GP system. Figure 3 compares the fitness of the best of population for both algorithms.

In order to further compare the performance of the two algorithms the cumulative probability of success for Even-5-parity has been computed using runs of 200 generations (population 500). The linear GP driven by the GA achieves a 28% probability of success after 200 generations. The same setup but using the cGA, the cl-GP, obtains a probability of success of 74%. For the Even-6-parity problem these numbers become 8% and 28% respectively.

5.2 Population compression, κ

Table 1 contains the raw fitness for the best individuals of the population for different problem instances. The second column, 0% of population selected for compression, serves as the reference since in this case the cl-GP is equivalent to a *normal* l-GP. As was expected high values do not lead to improved performance. This is most noticeable for small problem instances ($n = 5$ for example). Surprisingly, this tendency seems to reverse as the problem size increases.

Table 2 illustrates the effect of κ on the average individual length. Subjecting a larger fraction of the population to compression logically results in a lower average individual size.

n	$\kappa = 0.0$	0.3	0.6	0.9
5	0.891	0.964	0.950	0.949
7	0.680	0.783	0.796	0.769
9	0.5	0.528	0.530	0.536

Table 1: The number of individuals that are selected for compression influence the performance of the system. High values (60% and 90%) reduce the diversity and reduce the overall performance.

n	$\kappa = 0.0$	0.3	0.6	0.9
5	32	20.76	17.85	17.25
7	32	21.99	18.63	17.35
9	32	22.71	19.03	17.63

Table 2: The average individual size after 50 generations. The reduction in size decreases slightly as the problem instances become harder to solve.

5.3 Pool size, N

Table 3 presents the raw fitness of the best of population when using different pool sizes. The use of fitness proportional selection (with overselection) makes the cl-GP relatively insensitive to the pool size. This can be explained by the fact that overselection makes it possible for the some individuals to be present several time in the pool. As a result pools of different sizes can in fact show little differences in diversity.

5.4 Substring length, l

Table 4 contains the fitness of the best individual when using different substring lengths. Using longer dictionary entries decreases the performance. This decreasing trend can be observed for all the problem instances. As the substrings get longer it becomes harder for the cl-GP to protect good schemata. This clearly illustrates that by protecting sub-optimal building blocks the performance of the cl-GP is penalized.

6 Discussion

The cl-GP uses the cGA presented in this paper. The cGA provides a low level modularization mechanism based on compression. Although the linear GP system performs better when using the cGA than the GA, the differences get smaller as the problem instances become harder. This can

n	$N = 10$	20	50	100
5	0.964	0.964	0.964	0.965
7	0.783	0.795	0.785	0.774
9	0.528	0.537	0.531	0.524

Table 3: The raw fitness of the best of population as a function of the pool size used to build the dictionary of the substitution coder. The difference between various pool size is less pronounced compared to influence of κ on the performance.

n	$l = 2$	3	4	5
5	0.964	0.948	0.931	0.909
7	0.783	0.753	0.750	0.717
9	0.528	0.524	0.527	0.522

Table 4: Increasing the length of the dictionary entries makes it less likely for the assumption of tight linkage between gene values to hold. For the problem sizes the raw fitness decrease as the substring length increases.

be explained by the fact that the cGA, being an extension to the standard GA, capitalizes on the GA's ability to find good combinations. When the GA fails to identify building blocks it becomes impossible for the cGA to reuse these combinations to its advantage. In that respect the influence of the substring length l and the population compression ratio κ suggests that it is better to make moderate use of genotype compression. Longer substrings give higher compression ratios but put more strain on the tight linkage assumption of the cGA. Similarly, a high κ indirectly restricts the population diversity and reduces the performance.

7 Conclusions

In this paper we introduce the cl-GP system which uses a GA variant which uses compression as a modularization mechanism. The modularization scheme of the cGA assumes a tight linkage between elements of building blocks. This allows it to replace substrings in the genotype with a shorter reference. This scheme was designed to be used in a linear GP system (cl-GP) where the assumption is more likely to hold. We performed an empirical study of the influence of the different the cl-GP parameters using several Even-n-parity problem instances. Results illustrate the strong influence of the substring length and population compression ratio κ on the performance.

8 Future work

The influence of the parameters of the cl-GP needs to be studied on a wider range of GP problems. Although substrings of length 2 seem to be best at this point, it may be interesting to allow this size to change during the course of evolution. Alternatives to the used of fitness proportional selection for creating the pool of individuals should be explored.

Bibliography

- [1] M. Ahluwalia and L. Bull. Coevolving functions in genetic programming. *Journal of Systems Architecture*, 47(7):573–585, July 2001.
- [2] P. J. Angeline and J. B. Pollack. The evolutionary induction of subroutines. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Bloomington, Indiana, USA, 1992. Lawrence Erlbaum.
- [3] P. J. Angeline and J. B. Pollack. Coevolving high-level representations. In C. G. Langton, editor, *Artificial Life III*, volume XVII of *SFI Studies in the Sciences of Complexity*, pages 55–71, Santa Fe, New Mexico, 15-19 June 1992 1994. Addison-Wesley.
- [4] M. Brameier and W. Banzhaf. Effective linear genetic programming. Technical report, Department of Computer Science, University of Dortmund, 44221 Dortmund, Germany, 2001.
- [5] E. D. de Jong and D. Thierens. Exploiting modularity, hierarchy, and repetition in variable-length problems. In K. Deb, R. Poli, W. Banzhaf, H.-G. Beyer, E. Burke, P. Darwen, D. Dasgupta, D. Floreano, J. Foster, M. Harman, O. Holland, P. L. Lanzi, L. Spector, A. Tettamanzi, D. Thierens, and A. Tyrrell, editors, *Genetic and Evolutionary Computation – GECCO-2004, Part I*, volume 3102 of *Lecture Notes in Computer Science*, pages 1030–1041, Seattle, WA, USA, 26-30 June 2004. Springer-Verlag.
- [6] D. Howard. Modularization by multi-run frequency driven subtree encapsulation. In R. L. Riolo and B. Worzel, editors, *Genetic Programming Theory and Practise*, chapter 10, pages 155–172. Kluwer, 2003.
- [7] J. R. Koza, D. E. Goldberg, D. B. Fogel, and R. L. Riolo, editors. *Genetic Programming 1996: Proceedings of the First Annual Conference*, Stanford University, CA, USA, 28–31 July 1996. MIT Press.
- [8] T. Perkis. Stack-based genetic programming. In *Proceedings of the 1994 IEEE World Congress on Computational Intelligence*, volume 1, pages 148–153, Orlando, Florida, USA, 27-29 June 1994. IEEE Press.
- [9] S. C. Roberts, D. Howard, and J. R. Koza. Evolving modules in genetic programming by subtree encapsulation. In J. F. Miller, M. Tomassini, P. L. Lanzi, C. Ryan, A. G. B. Tettamanzi, and W. B. Langdon, editors, *Genetic Programming, Proceedings of EuroGP'2001*, volume 2038 of *LNCS*, pages 160–175, Lake Como, Italy, 18-20 Apr. 2001. Springer-Verlag.
- [10] J. P. Rosca and D. H. Ballard. Hierarchical self-organization in genetic programming. In *Proceedings of the Eleventh International Conference on Machine Learning*. Morgan Kaufmann, 1994.
- [11] K. Stoffel and L. Spector. High-performance, parallel, stack-based genetic programming. In J. R. Koza, D. E. Goldberg, D. B. Fogel, and R. L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 224–229, Stanford University, CA, USA, 28–31 July 1996. MIT Press.