A Simple Method for Detecting Domino Convergence and Identifying Salient Genes Within a Genetic Algorithm

Hal Stringer
School of Electrical Engineering and Computer Science
University of Central Florida, Orlando, FL 32816
stringer@cs.ucf.edu

Annie S. Wu
School of Electrical Engineering and Computer Science
University of Central Florida, Orlando, FL 32816
aswu@cs.ucf.edu

Abstract

Within a genetic algorithm, all genes may not be created equal. This concept is the central idea explored in this paper. A second and equally important idea is that this inequality in gene importance or salience can be detected and identified within a GA. To support these ideas, a technique for directly measuring genetic diversity within a GA population and thereby indirectly measuring gene-specific importance is provided. Diversity graphs are offered as a powerful technique for visualizing measurement results. Our theories, metrics and tools are tested on GAs for two problem classes and four different selection methods.

1 INTRODUCTION

Within a genetic algorithm (GA), all genes may not be created equal. Anecdotal evidence of this can be obtained from any student of genetic algorithms who has attempted to solve a symbolic regression problem. For example, consider a GA which finds the coefficients for the following equation:

$$y = ax^6 + bx^5 + cx^4 + dx^3 + ex^2 + fx + g + h\cos(x) \quad (1)$$

Intuitively, we expect genes representing the variables *a* through *h* to have varying impacts on fitness evaluation due to differences in the exponential order associated with each term. We would further expect the *a*-gene and *b*-gene to be the most important in determining fitness of an individual. The values for genes *c* through *h* would be largely irrelevant in terms of raw fitness until these first two genes converged to some local optimum.

The idea of a gene's importance or temporal-salience has already been described in (Thierens, Goldberg & Pereira, 1998). A side effect of this property is the phenomenon of "domino" convergence introduced in (Rudnick, 1992). A GA with non-uniformly salient genes converges serially over time starting with the more important genes and moving to less salient genes similar to the way a row dominos falls. Domino convergence and variations in gene importance have been shown to occur in genetic algorithms attempting to solve exponentially scaled fitness problems.

In subsequent works, (Goldberg, 1999) and (Srivastava & Goldberg 2001) explored how gene salience and domino convergence can be used to develop GAs with a serial mode of processing. A serial GA consists of small populations and short epochal runs. During each epoch different salient genes converge to their respective optimal values. Between each epoch, a continuation operator is activated to rejuvenate the diversity of less salient genes while leaving more important (and previously converged) genes alone.

Without continuation operators, GAs for exponentially scaled problems tend to converge around highly salient genes. The GA may then drift and stall at a less than optimal solution due to lack of diversity in less salient genes.

The use of epochs and continuation operators was found unproductive for problems with uniformly salient genes (e.g., OneMax). For these types of problems, the traditional GA's implicit parallelism, larger populations, and single long epoch were found to still be the most productive method of processing.

We believe that the idea of gene-specific temporal salience provides a valuable insight into how a GA functions. In the case of exponentially scaled problems, the concept opens up new opportunities for developing continuation operators to fine tune GA performance. But in order to use this approach, we must first find an effective method to determine if a problem includes genes with non-uniform salience and if so, a method for identifying those genes that are more important than others. This is particularly important in problems with very large numbers of genes where *a priori* knowledge of gene salience is less likely. In this work, we present a simple method for detecting domino convergence and identifying genes with high levels of importance. We show how tracking gene diversity within a GA population can provide the information we need to obtain a measurement of gene salience.

Our measurement technique focuses on two metrics. The first is the variation in unique alleles associated with each gene in a population. Unique allele counts plotted over time (generations) constitutes a convergence profile for a given problem and selection method. This profile clearly indicates the presence or absence of domino convergence. Our second metric consists of a ratio of unique sub-genotypes to alleles and assigns a numeric salience value to each gene. Graphically presented, this ratio gives us a salience profile identifying genes of higher importance.

In this work we describe the general nature of the experiments we performed to test the use of diversity as a salience indicator. Experiments include GAs for different problem classes and various selection methods. This range of problem classes and selection methods allow us to validate our method against previous theoretical work performed by other researchers.

2 MEASURING GENE SALIENCE THROUGH GENETIC DIVERSITY

The concept of gene salience or importance is all around us. For example, normal human beings are born with two eyes. Yet there exist numerous variations in eye color within the population. On a simplistic level, we can assume that the genes which affect the number of eyes in one's head are more important than those affecting eye color. We can also assume that a lack of diversity in the number_of_eyes genes relative to eye_color genes indicates that the first is more salient than the others.

The same concept applies to genetic algorithms with non-uniform gene importance. Over time, diversity of salient genes diminishes faster than that of non-salient genes. Less salient genes are not subject to the same selection pressures due to their low fitness impact. The diversity of alleles for each gene in a population relative to other genes provides a good indication of gene salience. The less diverse, the more important.

Using this idea we began investigating various ways to measure genetic diversity (or lack there of) within a GA. Initial experiments looked at uniqueness of entire chromosomes within a population. It was assumed that this method would provide a good showing of genetic diversity and illustrate how a population converges toward a small number of similar individuals over the course of multiple generations. This method was tested but found to be unsatisfactory. Looking at entire chromosomes did not single out specific genes nor indicate their specific importance. Nor did this method clearly show the presence or absence of domino convergence. We also investigated convergence to fitness values as a way of tracking convergence and diversity. This also proved to be less than satisfactory in identifying salient genes.

Throughout these initial experiments, we notice that there appeared to be a strong correlation between gene-salience and diversity of alleles within a single gene and also within partial chromosomes ("sub-genotypes"). The final version of our measurement methods used this idea and are described below.

2.1 UNIQUE ALLELES

The starting point for our method of determining genetic diversity within a GA is to count the number of unique alleles for each gene within a population at a given time. An allele can be thought of as a single representation instance of a gene. For example, using bit strings to represent a 9-bit gene allows for 2^9 different alleles.

For notational purposes a single gene location within a chromosome will be identified as G_i where $1 \leq i \leq n$ and n equals the total number of genes which make up a single chromosome. Two additional subscripts t and j are added to further specify a gene. t indicates a specific time or generation. j identifies an individual chromosome where $1 \leq j \leq p$ and p equals the population size. For example, $G_{3,\,100,\,12}$ denotes the third gene located on individual 12's chromosome at generation 100.

$U(G_{i,t})$ will be used to denote the count of unique alleles for G_i within the total population at the start of generation t.:

$$U(G_{i,t}) = \left| \{ G_{i,t,j} \mid G_{i,t,j} \neq G_{i,t,k} \text{ where } 1 \leq j,k \leq p \} \right|$$

To illustrate, assume that at the start of generation 54 during a GA's run, the third gene on all chromosomes contained bit representations (genotypes) for one of the following numbers: 12, -47, 178 or 3 (phenotypes). The population has evolved to contain chromosomes with only four unique alleles in the third position. In this example, $U(G_{3,54}) = 4$. Note that we are not concerned with how many genes contain a given allele, only the number of unique alleles within the population. $U(G_{i,t})$ provides a measure of the diversity of G_i within the population at the start of generation t.

Interesting results were obtained by following the behavior of a population using this measure. A low $U(G_{i,t})$ for a given gene relative to other genes in a chromosome indicates that the population is converging towards a few select alleles thus towards some local optimum. Unfortunately, the difference between $U(G_{i,t})$ for all genes within a GA was sometimes very small. This limited our ability to draw any firm conclusions regarding a specific gene's level of importance. Nor did this single statistic provide a total picture of what was occurring within the GA as a whole. Additional information was required.

2.2 UNIQUE SUB-GENOTYPES

Counting unique alleles gave us a way to track convergence of a given gene. But what about the rest of the genetic material within a chromosome?

To answer this question, we have developed the idea of a partial chromosome or "sub-genotype". A sub-genotype is the entire chromosome excluding a single gene. For notational purposes, $S_{i,t}$ will represent a chromosome's sub-genotype with respect to G_i at the start of generation t. The sub-genotype for a specific gene consists of the concatenation of all genetic material in the chromosome excluding the gene itself.

The example below illustrates how allele representations and sub-genotypes are derived from a hypothetical five-gene chromosome associated with individual 9 at generation 60:

Original Chromosome #9 at Start of Generation 60:

Gene:	#1	#2	#3	#4	#5
Value:	1010	1111	0011	0000	1101

Derived Gene Values and Sub-Genotypes:

$G_{1,60,9}$ = 1010, $S_{1,60,9}$ = 1111 0011 0000 1101
$G_{2,60,9}$ = 1111, $S_{2,60,9}$ = 1010 0011 0000 1101
$G_{3,60,9}$ = 0011, $S_{3,60,9}$ = 1010 1111 0000 1101
$G_{4,60,9}$ = 0000, $S_{4,60,9}$ = 1010 1111 0011 1101
$G_{5,60,9}$ = 1101, $S_{5,60,9}$ = 1010 1111 0011 0000

$U(S_{i,t})$ will be used to denote the count of unique sub-genotypes within the total population at generation t:

$$U(S_{i,t}) = \left| \{ S_{i,t,j} \mid S_{i,t,j} \neq S_{i,t,k} \text{ where } 1 \leq j,k \leq p \} \right|$$

2.3 RATIO OF SUB-GENOTYPES TO ALLELES

As a final measure of diversity, we also looked at the ratio of sub-genotype counts to the count of unique alleles. This ratio (R) is equal to the sub-genotype count divided by the unique allele count and can be shown as follows:

$$R_{i,t} = U(S_{i,t}) \,/\, U(G_{i,t})$$

Examples illustrating the importance of this relationship will be given later. For now, it is sufficient to say that this ratio "amplifies" the measurement of gene-specific salience and provides an better indicator of this important characteristic.

3 EXPERIMENT DESIGN

Many experiments were performed to capture the metrics described in Section 2. The purpose of these experiments was to test our ability to detect non-uniform salience and identify the salient order of genes within a chromosome. Experiments involved calculating and then graphing $U(G_{i,t})$, $U(S_{i,t})$ and $R_{i,t}$ for a variety of problem classes and selection methods. An analysis of the data obtained from the experiments supports our proposal that genetic diversity can reveal gene-specific salience in a GA.

Two different problem classes were tested and included in this paper: Symbolic Regression and OneMax. It was our expectation that gene-specific salience would be found in the symbolic regression problem. Based on the work researchers previously cited, we should find no important genes in the OneMax problem.

Experiments were conducted as follows:

1. A GA was executed for 50 runs of 100 generations each. All runs were initialized with a different random number seed.
2. All unique alleles and associated sub-genotypes were counted for each gene during each generation.
3. The allele and sub-genotype counts from step 2 were averaged across all 50 runs.
4. A ratio of the values from step 3 was calculated for each generation. Ratios were summed and divided by 100 for an average ratio across all generations.
5. The results from 3 and 4 were plotted for each problem as a set of six diversity graphs.

3.1 GA PARAMETERS AND SETTINGS

Our experiments used one of four selection methods: Fitness Proportional, Tournament, Rank Proportional and Random. Features and parameters incorporated into our GA for all experiments included the following: Population Size = 200 Individuals, Representation Method = Bit String, Number of Genes per Chromosome = 8, Number of Bits per Gene = 9, Crossover Type = 2-Point, and Crossover Rate = 100%.

With the exception of one experiment, mutation was not employed in any of our experiments. Our diversity metrics are based on counts of unique alleles and sub-genotypes. Mutation has the effect of increasing overall diversity in a population and tended to obscure though not hide our results. Leaving out mutation allows us to remove its effects from our measurements and focus on the evolution of individuals using only genetic material available from the initial population. One can think of the results of our mutation-less experiments as providing a baseline measure of gene salience and selection pressure within a GA.

3.2 COUNTING UNIQUE ALLELES AND SUB-GENOTYPES

The method proposed in this paper for identifying gene-specific salience requires that the number of unique alleles and sub-genotypes be determined for each gene in each generation. There are many different methods that can be used for such a counting function, some more efficient than others. The method employed for our experiments was simple though not necessarily the most efficient computationally.

All genes consisted of 9-bit binary strings representing integer values from –255 to +255. During fitness evaluation, these genotypic strings were converted to their phenotypic decimal equivalents. Genes were left in their original string format for counting purposes.

At the beginning of each experiment an *m* x *n* array (*count*) was constructed for storing unique allele counts where $m = 100$ was the number of generations in a run and $n = 9$ was the number of genes in each chromosome. All array elements were initialized to 0.

A hash table was used to keep track of unique alleles. The table was queried for the existence of each allele during the counting process. A gene value not found in the hash table was considered to be a new unique allele – the first of its kind. The corresponding element in *count* was incremented by 1 and the gene was added into the hashtable. If an allele was found to already exist in the hash table, no action was taken. The uniqueness of the allele had already been noted and added to the count for that gene during that generation. The following pseudo-code further illustrates this process:

```
for (i=1; i<=number_of_genes; i++){
  clear hash table;
  for (j=1; j<=population_size; j++){
     extract gene G_i from chromosome;
     if (G_i not in hashtable){
        add 1 to count[generation][i];
        add G_i to hashtable;
     }
     else no action necessary;
  }
}
```

A similar process was utilized to count unique sub-genotypes associated with each gene. It should be noted that the counting method described here is based primarily on the number of genes in a chromosome and is therefore usable with both small scale GAs and GAs with larger gene sizes (number of bits) or populations.

3.3 VISUALIZATION OF DIVERSITY

Results were written from the *count* array to a comma-delimited text file at the end of each experiment. The file contained the count of unique gene values and sub-genotypes for all 100 generations. Using data from this file, two graphs were plotted for each term described previously ($U(G_{i,t})$, $U(S_{i,t})$ and $R_{i,t}$.) One graph shows the change in the term over time (by generation) giving us an online view. The second graph shows an offline average value for each term for the entire GA run. Thus for each experiment, a suite of six graphs was prepared which, when viewed as a set, provided an excellent picture of the changing genetic diversity within a GA. Examples of these diversity graphs are provided throughout the remainder of this work along with our analysis.

4 FINDING TEMPORAL SALIENCE

Given the introduction to this work, it is fitting that symbolic regression be the first problem used to test our diversity measurement technique. Predetermined *x* and *y* values were provided as input to the GA's fitness function. The GA's task was to find optimal values for coefficients *a* through *h*. Positionally, these coefficients corresponded to genes 1 through 8 on a chromosome.

Intuition and knowledge of the problem lets us know that the first gene (G_1) will be the most salient and have the greatest impact on fitness evaluations due to its association with the term ax^6. The population should converge around this one gene before all others. G_2 representing the coefficient for bx^5 would be next in importance followed by G_3, G_4, G_5 and so on.

Experiments were run per the design in Section 3 using tournament selection. Unique alleles and sub-genotypes for all runs were counted, averaged and plotted on a set of diversity graphs (see Figure 1).

Figure 1(a) shows the convergence profile for this problem/selection method combination. Allele diversity for the gene associated with the *a* coefficient - $U(G_1)$ - drops fastest followed by $U(G_2)$ and the other genes. By generation 41, only one allele for G_1 exists in the population. The gene's salience caused a single value to quickly take over this gene in the entire population. This graph also shows a similar but delayed behavior for G_2 though G_8 over the course of 100 generations. The result is a staggered look to the graph clearly indicating the domino convergence occurring in this experiment.

Figure 1(b) shows the diversity of sub-genotypes for this problem. Diversity for all sub-genotypes drops decreases over time as the GA converges to a single result.

Figures 1(d) and 1(e) provide an offline view of allele and sub-genotype counts. Both of these graphs contain the average number of unique alleles or sub-genotypes over 100 generations. For our test problem, G_1 and G_2 have the lowest average unique number of alleles. On average, only 7.7 different values for G_1 existed during each generation due to this gene's salience. Although hard to tell from the graph, G_1 has the greatest sub-genotype diversity. On average, any given individual in the population will include one of 105 G_2-through-G_8 gene combinations regardless of G_1's value. It was found that generally the lower the unique allele count, the higher the sub-genotype count within the GA.

From these first four graphs, we can see that G_1 and G_2 or coefficients *a* and *b* respectively, are the more important genes and exert a higher degree of selection pressure than the other genes in this GA. But these graphs alone may not be enough to clearly indicate gene salience. A more reliable indicator has proven to be the ratio of sub-genotypes to gene values ($U(S_{i,t})$ / $U(G_{i,t})$). These ratios are plotted for our same test problem in Figures 1(c) and 1(f).

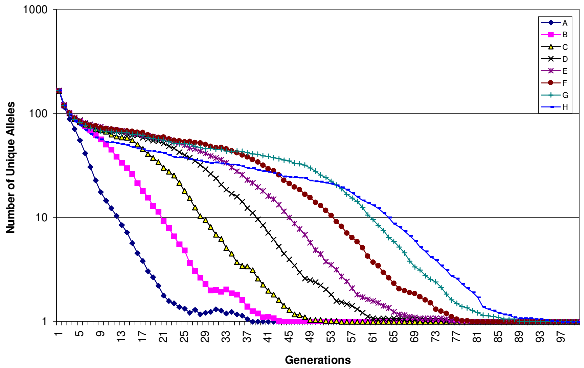

(a) Unique Allele Counts $U(G_i)$ over Time (Convergence Profile)

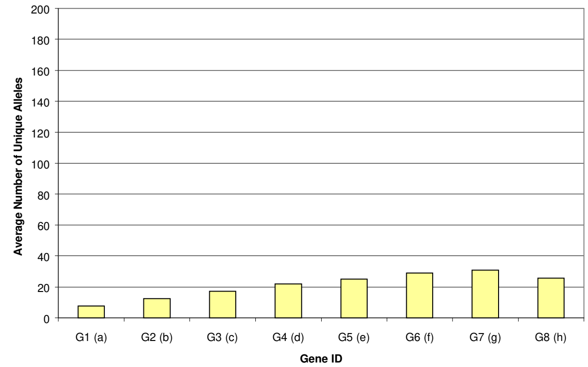

(d) Avg. Unique Alleles per Gene

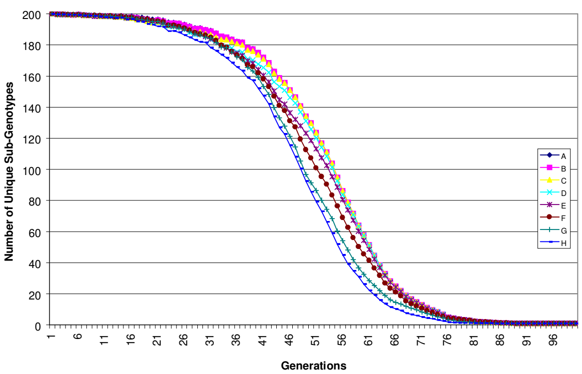

(b) Unique Sub-Genotype Counts $U(S_i)$ over Time

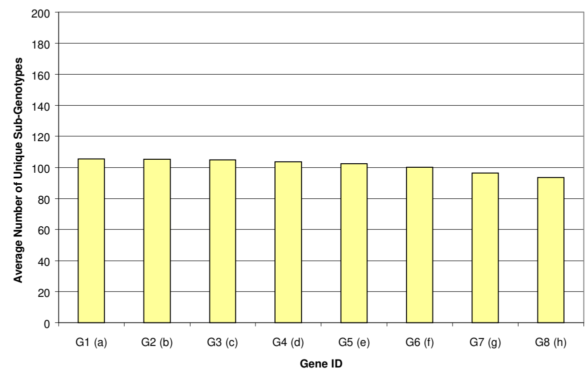

(e) Avg. Unique Sub-Genotypes per Gene

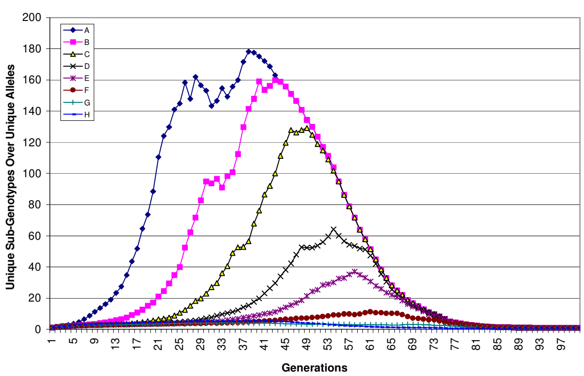

(c) Ratio (R_i) of Sub-Genotypes to Alleles over Time

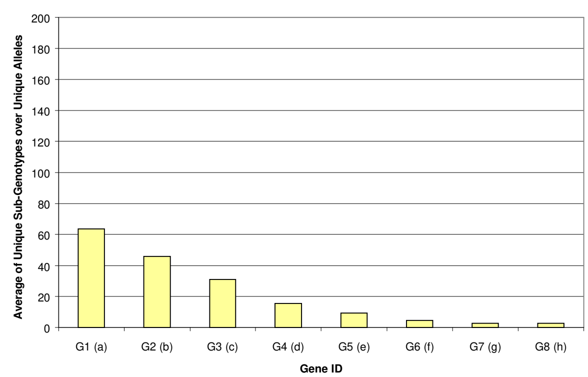

(f) Avg Ratio (R_i) of Sub-Genotypes to Alleles (Salience Profile)

Figure 1: Diversity Graphs for 8-Term Symbolic Regression Problem Using Tournament Selection

As mentioned earlier, this ratio $R(G_{i,t})$ tends to amplify our ability to detect gene-specific importance and make it easier to pick out the genes with greatest salience. Figure 1(f) is most important to us and we have called this type of plot a "salience profile." From Figure 1(f) it is very clear which genes in our GA are more salient than others.

A variation on the preceding symbolic regression problem was developed to check the previous results. In this second problem, the positional order of terms was mixed. The resulting equation is:

$$y = ax + bx^4 + c\cos(x) + dx^5 + e + fx^2 + gx^6 + hx^3$$

Assuming our proposal is correct, G_7, G_4, and G_2 should exhibit behavior that typifies genes of higher importance. Figure 2 shows the salience profile for this reordered problem. As expected, $G_{7,}$ G_4 and G_2 had the highest average ratio of unique sub-genotypes to alleles $R(G_{i,t})$ of the eight genes.

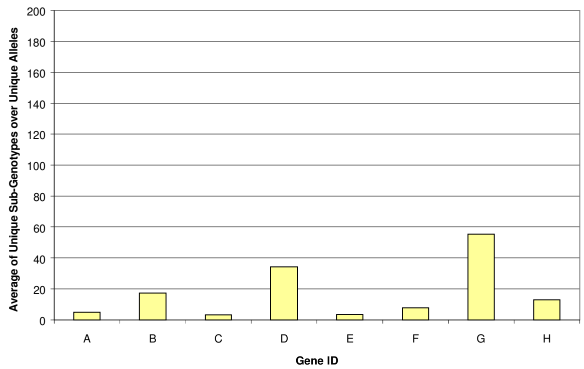

Figure 2: Saleince Profile for Symbolic Regression Problem with Reordered Terms

This second experiment confirms that our measurement technique can identify salient genes regardless of their position within a chromosome.

As mentioned in Section 3, most of our experiments were run without mutation. For sake of completeness, we incorporated bit mutation at a rate of 0.01 in a third experiment using equation (1). Figure 3 shows the salience profile for this GA. A comparison of Figure 3 with Figure 1(f) shows that mutation reduced but did not eliminate the indication of gene-specific salience calculated from $R(G_i)$. Using a magnified *y*-axis, the stair step pattern indicating the presence of domino convergence is still apparent.

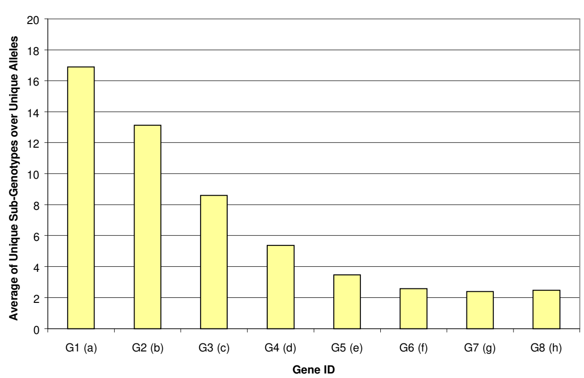

Figure 3: Salience Profile for Symbolic Regression Probliem with Mutation = .01

5 OTHER SELECTION METHODS

Graphs in Figures 1, 2 and 3 were associated with GAs using tournament selection. How well does our measurement technique work with other selection methods? To answer this question we present convergence and diversity profiles for GAs solving equation (1) using random and fitness proportional selection (Figures 4 and 5 respectively). Space does not allow a detailed description of the results. However, a few points should be noted.

The plots for random selection show that lack of directed selection pressure leads only to drift in gene diversity.

Gene-specific salience also appears in GAs run with fitness proportional selection. The exponential effect of the selection method causes the GA to converge very rapidly around highly salient genes. As a result we do not see the stair stepped or staggered type of convergence profile found in Figure 1. The salience profile is stronger for genes of higher importance.

The important concept to be seen in these graphs is the impact of selection method. Rank or tournament selection is best for detecting domino convergence and identifying the gene order in terms of salience. However, fitness proportional selection provides a very clear indication of the importance of the most salient genes in a chromosome. As a result, fitness proportional selection may be most useful when results for other selection methods are less clear.

Experiments testing our method on rank fitness were also performed. We do not include the diversity and salience profiles for these experiments as they were very similar to those of tournament selection. When combined with a OneMax problem, profiles for binary tournament and rank selection were nearly identical as was expected based on showings in (Blickle & Thiele, 1995).

6 OTHER PROBLEM CLASSES

It appears that we have found a simple method for identifying domino convergence and gene-specific salience in a GA. But what about detecting a lack of gene importance? Random selection results in the elimination of gene salience regardless of the problem type. Can we also show that a problem class in and of itself lacks salient genes. To further test our technique, we ran a GA using a OneMax problem. The fitness function merely counts the number of ones in the entire chromosome. This problem has been shown to have uniform salience across all genes. Figure 6(a) and (b) contain convergence and salience profiles for this problem using tournament selection. We can see from these plots that no gene is more salient than another.

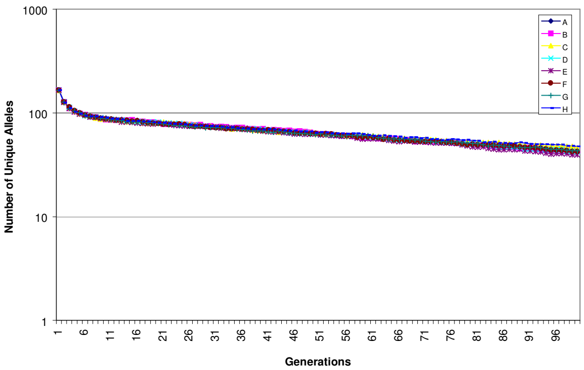

Figure 4(a): Convergence Profile for Symbolic Regression Problem with Random Selection

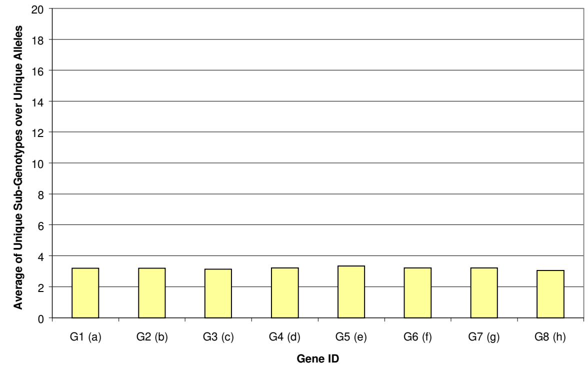

Figure 4(b): Salience Profile for Symbolic Regression Problem with Random Selection

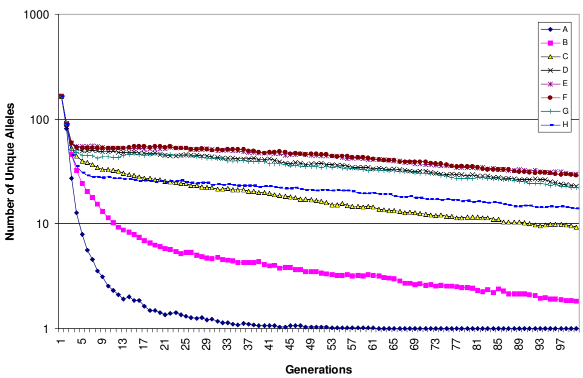

Figure 5(a): Convergence Profile for Symbolic Regression with Fitness Proportional Selection

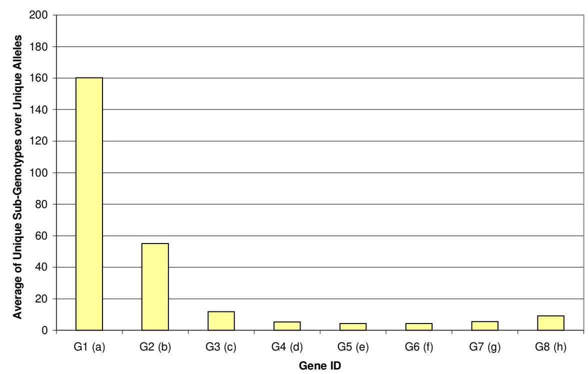

Figure 5(b): Salience Profile for Symbolic Regression with Fitness Proportional Selection

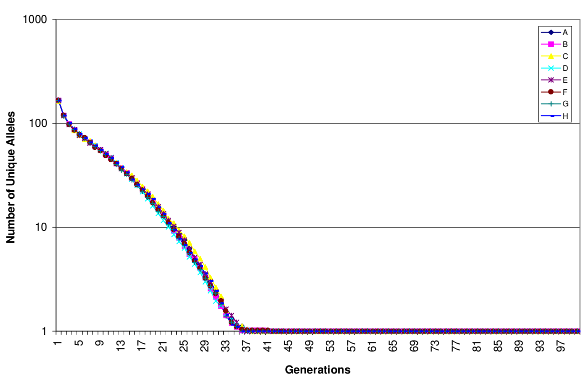

Figure 6(a): Convergence Profile for One-Max Problem with Tournament Selection

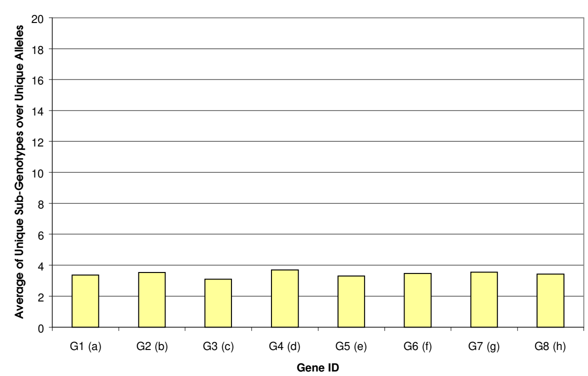

Figure 6(b): Salience Profile for OneMax Problem with Tournament Selection

7 CONCLUSIONS

In this work we have presented a simple but useful method for detecting domino convergence and gene-specific salience within a given problem. It is not uncommon for certain regions of GA individuals to consistently converge early. Those regions are typically expected to be regions that have high impact on the fitness function. The ability to detect high impact regions would allow practitioners to potentially develop operators that may improve GA performance on a particular problem. Our detection method is based on a count of the unique gene alleles and unique sub-genotypes that occur within a short run. While both of these counts provide some indication of gene salience, it is the ratio of the sub-genotype count to the unique allele count that appears to give the clearest picture as to which genes have the strongest impact on the GA search process.

We tested our method for detecting salient genes on problems in which genes are and are not expected to have varying impact. From results we are clearly able to detect salient genes when they exist, regardless of their position within a chromosome.

A comparison of salience profiles for varying selection methods indicate that choice of selection method can enhance or diminish gene-specific salience depending on the desires of the GA researcher/developer. Our experimental evidence shows that fitness proportional selection magnifies a gene's selection pressure. Tournament or rank fitness selection methods reduce that pressure and allow the temporal salient nature of more genes to shine through.

Choice of selection method is an example of how gene-salience can be manipulated on a chromosomal- or problem-wide scale. But can we manipulate selection pressure at the gene level? The use of continuation operators is a step in that direction. We believe that the ability to identify salient genes within a GA will help researchers in those development efforts.

While the information presented here is of value we do recognize that our methods have their limitations. Specifically, our methods were designed for GAs with fixed gene positions and would not be directly applicable to locus-variable situations such as messy GAs or GAs with variable length chromosomes. We believe the development of methods for detecting gene salience in these other GA categories will be a productive area for future research.

In addition, our research focused on gross numerical counts of unique allele values and sub-genotypes This approach can suffer from scalability issues which may be addressed by taking measurements on restricted GA runs (e.g., short duration or small populations.) These restricted runs can reduce processing time while still providing information about the problem. Such gross numerical counts also make no attempt to evaluate genes or sub-genotypes qualitatively. Further research in these areas are expected to provide a greater understanding of genetic diversity within a GA.

Despite these limitations, we feel the knowledge gained from our research has immediate value. We can now detect domino convergence within a GA and thus non-uniform gene salience. In addition, we can identify important genes within these GAs and begin to use this knowledge towards development of better control mechanisms.

In terms of immediate applications, our method may be helpful in a number of ways. The programming effort required to extract our measurements (V(Gi,t), S(Gi,t) and R(Gi,t)) is relatively small. A few lines of code added to any fixed position GA would allow a quick view of any gene-specific temporal salience that the GA might encounter.

We feel our diversity graphs will be useful in development and evaluation of new genetic operators and selection methods. We have already mentioned the strikingly similar results found for rank and tournament selection which concur with theoretical studies. Diversity graphs would show where new selection methods are similar to existing methods and where they differ.

References

Blickle, T. & Thiele, L. (1995). A Comparison of Selection Schemes Used in Genetic Algorithms. *TIK-Report No. 11*, Computer Engineering and Communication Networks Lab (TIK), Swiss Federal Institute of Technology (ETH), Zurich, Switzerland.

Goldberg, D. E. (1999). Using Time Efficiently: Genetic-Evolutionary Algorithms and the Continuation Problem, In *Proceedings of the Genetic and Evolutionary Computation Conference: GECCO 1999,* Volume 1 pp. 212-219, Morgan Kaufman Publishers: San Francisco.

Rudnick, W. Michael (1992). *Genetic Algorithms and Fitness Variance with an Application to the Automated Design of Artificial Neural Networks*. Unpublished doctoral dissertation, Oregon Graduate Institute of Science and Technology.

Srivastava, R.P. & Goldberg D.E. (2001). Verification of the Theory of Genetic and Evolutionary Continuation, *In Proceedings of the Genetic and Evolutionary Computation Conference: GECCO 2001*, Morgan Kaufman Publishers: San Francisco.

Thierens, D., Goldberg, D. E., & Pereira, A. B. (1998). Domino Convergence, Drift and the Temporal-Salience Structure of Problems, In *The 1998 IEEE International Conference on Evolutionary Computation Proceedings*, pp. 535-540, IEEE Press: New York, NY.