# Neutrality and Variability: Two Sides of Evolvability in Linear Genetic Programming

Ting Hu
Department of Computer Science
Memorial University
St. John's, Canada
tingh@mun.ca

Wolfgang Banzhaf
Department of Computer Science
Memorial University
St. John's, Canada
banzhaf@mun.ca

## ABSTRACT

The notion of *evolvability* has been put forward to describe the "core mechanism" of natural and artificial evolution. Recently, studies have revealed the influence of the environment upon a system's evolvability. In this contribution, we study the evolvability of a system in various environmental situations. We consider neutrality and variability as two sides of evolvability. The former makes a system tolerant to mutations and provides a hidden staging ground for future phenotypic changes. The latter produces explorative variations yielding phenotypic improvements. Which of the two dominates is influenced by the environment. We adopt two tools for this study of evolvability: i) the rate of adaptive evolution, which captures the observable adaptive variations driven by evolvability; and ii) the variability of individuals, which measures the potential of an individual to vary functionally. We apply these tools to a Linear Genetic Programming system and observe that evolvability is able to exploit its two sides in different environmental situations.

## Categories and Subject Descriptors

I.2.2 [**Artificial Intelligence**]: Automatic Programming, Program Synthesis

## General Terms

Experimentation, Measurement, Performance

## Keywords

Evolvability, Rate of Evolution, Neutrality, Variability

## 1. INTRODUCTION

In the process of evolution, genetic variation explores new evolutionary material, the corresponding phenotypic variation provides adaptive characteristics, and stabilization operators like recombination and selection preserve these improvements over the previous generations. It is the interactions among these operations that allow evolution to work. Thus, the evolvability of an evolutionary system is constituted by the capability to coordinate these operations in order to yield phenotypic improvements. A growing number of efforts have been dedicated to understanding [15, 21] and enhancing [8, 25] evolvability.

Evolvability is the potential of a population to evolve. While the concept of evolvability is still very much under discussion, we want to venture to propose a definition that is equally applicable to natural and artificial systems:

*Definition 1.* Evolvability — the capability of a system to generate adaptive phenotypic variation under certain environmental conditions and to transmit it via an evolutionary process.

Altenberg [1] describes evolvability from a viewpoint of Evolutionary Computation (EC) as the ability of a genetic operator or representation scheme to produce offspring fitter than their parents. In Biology, Kirschner and Gerhart [14] suggest that evolvability should be understood as an organism's capacity to generate heritable and selectable phenotypic variation. An explicit comparison between evolvability of biological and computational systems has been performed by Wagner and Altenberg [25]. In their view, evolvability should be considered as the ability of random variants to produce occasional improvements, which depends critically on the plasticity of the genotype-phenotype mapping. Marrow [16] suggests that evolvability means the capability to evolve, and this characteristic should be relevant to both natural and artificial evolutionary systems. Nehaniv [18] proposes the perspective of using evolutionary system complexity to describe evolvability. Recently, a growing number of evolutionary biologists and computer scientists have shown interests in this topic. In an evolutionary system, many properties of a population are considered related to evolvability, including adaptive representation [22], facilitation of extra-dimensional bypass and robustness against genetic variability [7, 23], redundancy and flexibility during developmental processes [14], mutation rate adaptation [3], and balancing between robustness and innovation [12].

It would be overly optimistic to expect a formula to describe evolvability mathematically, due to the complexity of organisms, the dynamics of populations, and the influence of the environment. The most striking feature of evolvability is its capability to generate adaptive phenotypic variation from random genetic changes. Neutrality and variability are the two sides of evolvability important to control random genetic changes. Genetic changes do not necessarily result in any observable phenotypic variation. This "neutrality" has


*GECCO'09,* July 8–12, 2009, Montréal Québec,Canada.

two functions, i) it improves a system's robustness against deleterious genetic changes, and ii) neutrality also provides variation potential through exploiting neutral networks [2, 10]. In contrast, "variability" generates observable phenotypic variations for adaptation to the environment. These two sides of evolvability may appear contradictory at first sight. However, they closely cooperate to facilitate evolution. Moreover, it is the environment that dictates which side dominates at which stage of an evolutionary process.

In this contribution, we investigate the two sides of evolvability under various environmental scenarios in the context of a Linear GP system, focusing on a polynomial symbolic regression problem. The rest of the paper is organized as follows. In Section 2, we discuss the importance of the environment in evolvability research and review some of the relevant recent discoveries in the biological community. We discuss two metrics adopted in this work to investigate evolvability in Section 3. First, the nonsynonymous to synonymous substitution ratio $k_\mathrm{a}/k_\mathrm{s}$ [11] captures the adaptive and silent substitution rates of a population. Second, for a closer look at a population's structure, the variability of an individual is defined based on its connectivity in neutral networks [2, 24]. Section 4 presents our simulated studies in a Linear GP system. Three typical environmental scenarios are modeled: i) random evolution (RE), ii) fixed target (FT), and iii) moving target (MT) evolution. Our concluding remarks and future work will be discussed in Section 5.

## 2. EVOLVABILITY AND ENVIRONMENT

It has been well accepted that evolution can be understood in general with three fundamental elements: variation, selection and inheritance. It is impossible to study a system's evolutionary capability without consideration of its environment, which is the selection force that preserves variations. The detection and investigation of evolvability are non-trivial and intriguing problems. Phenotypic fitness is directly observable and serves as a selection criterion. However, as a *potential* to generate better fitness and a *capability* for adaptive evolution, evolvability is more difficult to observe and to select for. Therefore, some empirical methods have been proposed in the literature to investigate evolvability "indirectly" in various environments.

Orr [19] analyzes the acceptance of mutations in a system moving toward a stationary optimum, and suggests that the effects of accepted mutations are decreasing. That is, towards a fixed target, rapid phenotypic variations can be observed at the beginning, but the rate of observable adaptive evolution will slow down later on. Further, Collins et al. [6] study adaptive walks in dynamic environments, and report that the rate of environmental change has a systematic effect on adaptive walks. Gradual changes allow more small-effect genetic substitutions than sudden changes, and favor more robust individuals that do not behave poorly in any intermediate environment. Thus, a large drop in fitness almost never happens in a gradually changing environment. Earl and Deem [9] suggest that evolvability can be selected for by varying the environment. By observing genetic changes in protein evolution, they find that rapid or dramatic environmental changes generate strong selection pressure for evolvability. Thus, high evolvability can be detected and favored by such selection pressure. Meyer et al. [17] state that fluctuating environments can drive populations towards the edge of a neutral network. Kashtan et al. [13, 20] report that varying environments, especially in a mode that prefers modular changes, can facilitate rapid adaptive phenotypic variations.

In the GP literature, evolvability has also emerged as a very important research topic. Ebner et al. [10] incorporate redundant mapping from genotype to phenotype in evolutionary computation models as a form of neutrality and show how neutral networks can influence evolvability. Further, Banzhaf and Leier [2] examine the behavior of an evolutionary search process in neutral networks using Linear GP for a stationary Boolean search problem. Belle and Ackley [4] design a dynamic environment exploiting modularity among varying goals, and argue that enhancing the search modularity in a changing environment can increase a GP system's evolvability. Yu [27] reports that GP populations exhibit various program distributions under different environmental variation rates.

In this contribution, we adopt interesting current discoveries from Biology to GP. We focus on the influence of the environment on the two sides of evolvability, neutrality and variability. We hypothesize that evolvability of our computational evolutionary system can be different in various environments.

## 3. METHODS

Variation is the driving force of evolution. However, most random genetic variations are well known to be deleterious. Evolutionary systems exploit neutrality and variability, as two opposite strategies, to control random genetic changes in different situations. The core of evolvability is to generate *adaptive* variations at the phenotypic level from random genetic changes. Therefore, we believe that the cooperation of persistence and sensitivity to random genetic changes contributes substantially to evolvability, and that it will be interesting to test whether the dominance of either side is driven by the environment. To do this, we adopt two methods, one emphasizing the temporal aspect and the other emphasizing the spatial aspect of evolvability.

### 3.1 Nonsynonymous to Synonymous Substitution Ratio $k_\mathrm{a}/k_\mathrm{s}$

The *nonsynonymous (amino acid) to synonymous substitution ratio* $k_\mathrm{a}/k_\mathrm{s}$ is a widely accepted measurement of the rate of genetic substitutions in molecular biology [26]. It is known that some mutations of a genetic sequence will lead to amino acid substitutions while others will not, due to the degenerate code employed for translation. Therefore, $k_a/k_s$ measurement compares two homologous protein-coding gene sequences from two related species. The $k_a/k_s$ ratio resulting from measuring the number of nonsynonymous (amino acid) substitutions per nonsynonymous site ($k_a$) to the number of synonymous substitutions per synonymous site ($k_s$) characterizes the adaptive evolution rate between these two sequences. Recently, this metric has been applied to measure the rate of genetic substitutions in tree-based GP by us [11]. Similar to natural systems, in GP evolution, it is known that not all genetic changes are effective [5]. Genetic changes that can modify the encoded function are regarded *nonsynonymous*, while others are regarded as *synonymous* changes. Thus, accepting nonsynonymous genetic changes leads to observable phenotypic variations in a system, and synonymous substitutions contribute to evolution in a "silent" manner. In this measurement, $k_\mathrm{a}$ measures the rate of nonsynonymous

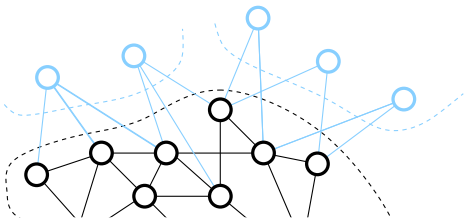

**Figure 1: A simple example of neutral networks. Black lines show the connection within a neutral network, and blue lines mark the connection among different neutral networks.**

genetic substitutions of an evolutionary process, and $k_s$ measures the rate of neutral changes accepted. Practically, $k_a$ captures the adaptive evolution "distance", and $k_s$ provides the background "clock ticks". Thus, the $k_a/k_s$ ratio is regarded as a measure for the adaptive evolution rate [11].

Here, we slightly adapt the $k_a/k_s$ ratio measurement for Linear GP. Note that there can be variants in defining nonsynonymous and synonymous genetic changes. For a strict analogy to biological systems, which often have no explicit fitness, the effects of a genetic change would refer to the influence on its phenotype. However, in a GP system, fitness is explicitly defined and is in most cases the only criterion for selection. Therefore, we make the simplifying assumption that a GP genetic change is nonsynonymous (synonymous resp.) if it changes (maintains resp.) the fitness of an individual. Further, we quantify the rates of nonsynonymous and synonymous substitutions in each generation of evolution in a way different from [11]. Specifically, in generation $t$, we observe all the surviving individuals from the previous generation. We use $N_a(t)$ to count the number of all *attempted* nonsynonymous genetic changes and $M_a(t)$ to count the number of *accepted* nonsynonymous genetic changes when producing generation $t$. Thus, $k_a(t) = M_a(t)/N_a(t)$ measures the rate of nonsynonymous genetic substitutions. The synonymous substitution rate $k_s(t)$ can be defined similarly by dividing the number of accepted synonymous genetic changes $M_s(t)$ by the number of attempted synonymous genetic changes $N_s(t)$. Typically, the ratio $k_a/k_s$ is calculated over time, and it measures the rate of adaptive genetic substitutions relative to a background silent genetic substitution rate as a time series.

## 3.2 Neutral Networks

In genotype space, a neutral network is usually defined as a set of genotypes that map to the same phenotype [2, 10, 17, 24]. Each genotype corresponds to one vertex in the neutral networks. A genotype $G_1$ is linked to another genotype $G_2$ if $G_2$ can be obtained from $G_1$ via a one-step mutation. Note that these links are usually bidirectional due to the reversibility of mutation. Further, a link can exist both within and across neutral networks. We say that a genotype is a "neighbor" of $G_1$ if it is linked with $G_1$. In addition, it is a "neutral neighbor" if it belongs to the same neutral network as $G_1$. Otherwise, it is a "non-neutral neighbor". For a given genotype, we follow Wagner [24] in defining its variability as the fraction of non-neutral neighbors among all of its neighbors. This quantifies the likelihood that a mutation from a given genotype leads to a phenotypic change. Again, as a simplification in this contribution, we assume that two genotypes are in the same neutral network if they have the same fitness, instead of looking at their phenotypes.

Fig. 1 depicts a simple example of three genotype neutral networks. Genotypes with high variability are positioned near the edge of a neutral network, and genotypes more robust against genetic changes are placed close to the center of this network. Therefore, the distribution of individuals in neutral networks can reflect the dominance of either neutrality or variability of a population at a given point in time.

For simple problems, all reachable genotypes can be exhaustively enumerated. However, the genotype space grows exponentially with the complexity of a problem. Thus, we need to sample the genotype space to obtain an approximation for complex problems. That is, for a given genotype, we sample a sufficiently large number of its neighbors to estimate its variability. This is the approach we will adopt here.

**Table 1: LGP Configuration.**

| | |
|---|---|
| Population size | 1000 |
| Initial program length | $5 \sim 15$ |
| Maximum program length | 200 |
| Number of input register | 1 |
| Number of calculation register | 3+1 |
| Constants | 1,2,...,9 |
| Operator set | $+$, $-$, $\times$, protective $\div$ |
| Mutation rate | 1 |
| Sample set | 100 cases in [-1:1] |
| Fitness function | Mean error |
| Truncation selection | Tournament with size 6 |
| Maximum generations | 500 |
| Neutral network space | 1000 sampled neighbors |

# 4. SIMULATED STUDIES

We use Linear Genetic Programming in our experiments. We choose Linear GP over the more commonly studied Tree GP because Linear GP seems to have a better resemblance to biological systems. Moreover, we would like to study a different representation since we have tested the $k_a/k_s$ ratio on a Tree GP system previously [11]. We design a set of varying environmental scenarios, and measure the $k_a/k_s$ ratio and the variability of genotypes in neutral networks in order to investigate evolvability in different environmental situations.

## 4.1 Test Case

Our benchmark is the polynomial symbolic regression problem ($\sum_{i=1}^{n} x^i$, for some $n$). Note that there can be similar patterns within this polynomial. For example, when $n = 4$, $x^4+x^3+x^2+x = x(x+1)(x^2+1) = x^2(x^2+1)+x(x^2+1)$. Also, if we increase $n$, we can design moving targets based on this expression. Here, only mutation is used for genetic changes. Each mutation can take two forms. A micro-mutation limits the change to one element of a specific instruction, i.e., the return register, the operator, or one of the two operand registers. A macro-mutation inserts a randomly generated instruction into the program or deletes one instruction, both at a random location. In particular, the mutation rate of a program is 1, with half of the likelihood happening at the micro level and half at the macro level.

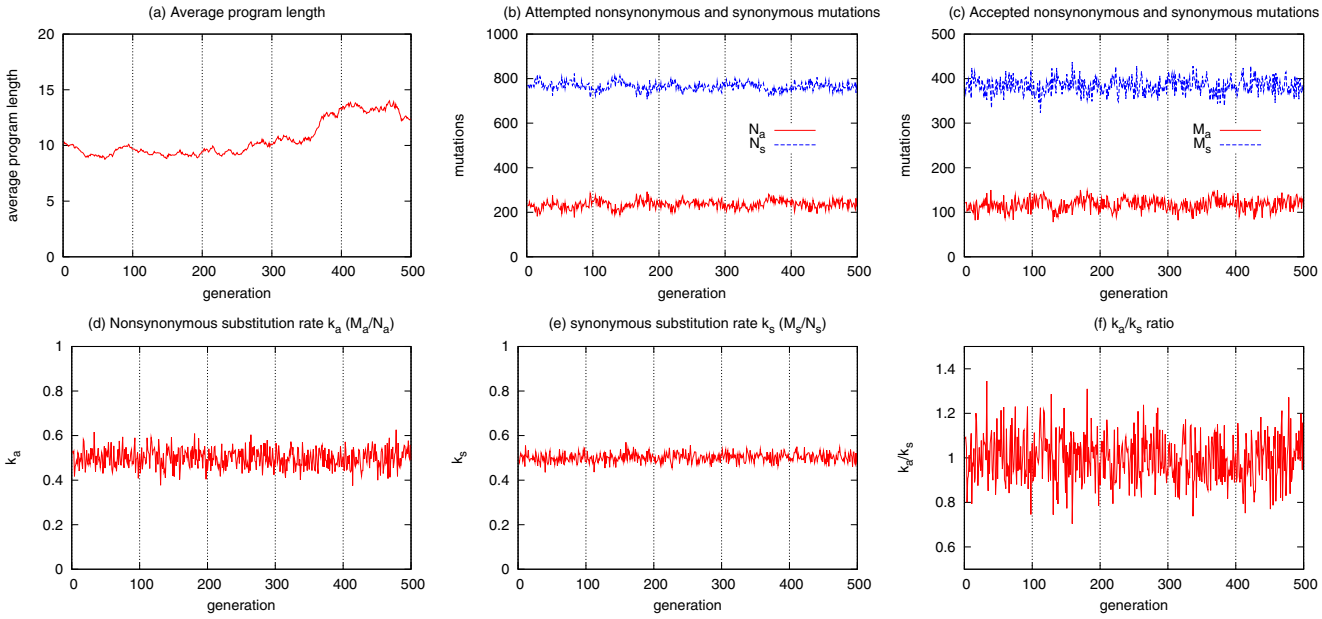


**Figure 2: A typical single run of RE**

When a program adopts a macro-mutation, the instruction insertion and deletion occur with equal probability. We employ a truncation selection scheme where both parent and offspring populations will compete to form the next generation. The configuration is specified in Table 1.

## 4.2 Varying Environmental Scenarios

Here, in the context of a Linear GP system with symbolic regression, we define environment as the target polynomial expression. In this sense, typical environmental scenarios include i) random evolution, where no specific evolution target is defined, ii) fixed target evolution, and iii) moving target evolution. In the following, we study the effects of these scenarios on Linear GP.

### *4.2.1 Random Evolution (RE)*

We implement random evolution by applying random selection when forming a new generation. We plot various measurements of the process in Fig. 2. In Fig. 2(a), we plot the average program length over time and observe that there is no general trend of length. This is distinct from normal Linear GP, where average program length increases. All other metrics in this figure, however, indicate a fair level of stability (Fig. 2(b)∼(f)). In Fig. 2(b), the system presents a constant 20-80 split among the 1000 total mutations between the nonsynonymous ($N_\mathrm{a}$) and synonymous ($N_\mathrm{s}$) changes. The accepted nonsynonymous ($M_\mathrm{a}$) and synonymous ($M_\mathrm{s}$) substitutions remain at half of the level (Fig. 2(c)) because a new generation always starts with the combination of all parents and offspring and half of them survive at random. This means that the $k_\mathrm{a}$ and $k_\mathrm{s}$ rates are both approximately 0.5 (Fig. 2(d) and (e)), with $k_\mathrm{a}$ having slightly higher variance. This further implies that the $k_\mathrm{a}/k_\mathrm{s}$ ratio stays at around 1, i.e., the neutrality and variability apply equal influence in a random evolution system.

In addition to observing the system as it progresses, we are also interested in the variability of all the individuals at time snapshots. In particular, we plot our measures at the beginning and end of the evolution in Fig. 5(a). In the figure, each snapshot corresponds to one plotting. We sort the individuals according to their degree of variability for better readability. Apparently, the random evolution process does not alter the variability composition of the population.

### *4.2.2 Fixed Target (FT) Evolution*

In this experiment, we start out with a simple fixed target of $x^2 + x$. The evolution quickly leads the first individual to optimum at generation 5 and the entire system converges to this optimum at generation 10. After this, the average program length keeps increasing (Fig. 3(a)), which builds more and more redundancy into individual programs. Fig. 3(b) records the number of nonsynonymous mutations $N_\mathrm{a}$ and that of the synonymous mutations $N_\mathrm{s}$ at each generation. These two metrics start with a 20-80 split as with the previous scenario because of the randomness of the initial population composition. As the system progresses the optimum and converges (up to generation 10), $N_\mathrm{a}$ increases as a large number of the mutations are nonsynonymous. After this point, $N_\mathrm{a}$ decreases and approaches 0 due to the increasing robustness in the population. $N_\mathrm{s}$ follows the complementary trend in this process. As in Fig. 3(c), the system starts to completely reject nonsynonymous changes ($M_\mathrm{a}$) after the convergence to the optimum because any such change is deleterious and is not favored by selection. During the process, $M_\mathrm{s}$ remains at about half of the level of $N_\mathrm{s}$ because of the half-half composition of a new generation before selection. The nonsynonymous substitution rate $k_\mathrm{a}$ (Fig. 3(d)) has a positive value until convergence, indicating no phenotypic evolution occurs after this point. The synonymous substitution rate $k_\mathrm{s}$ (Fig. 3(e)) suggests a very active background evolution before system convergence, which stabilizes at approximately 0.5 afterwards. As a result, the $k_\mathrm{a}/k_\mathrm{s}$ is always less than 1, and has a positive value until the system converges. This is the result of the majority of random muta-

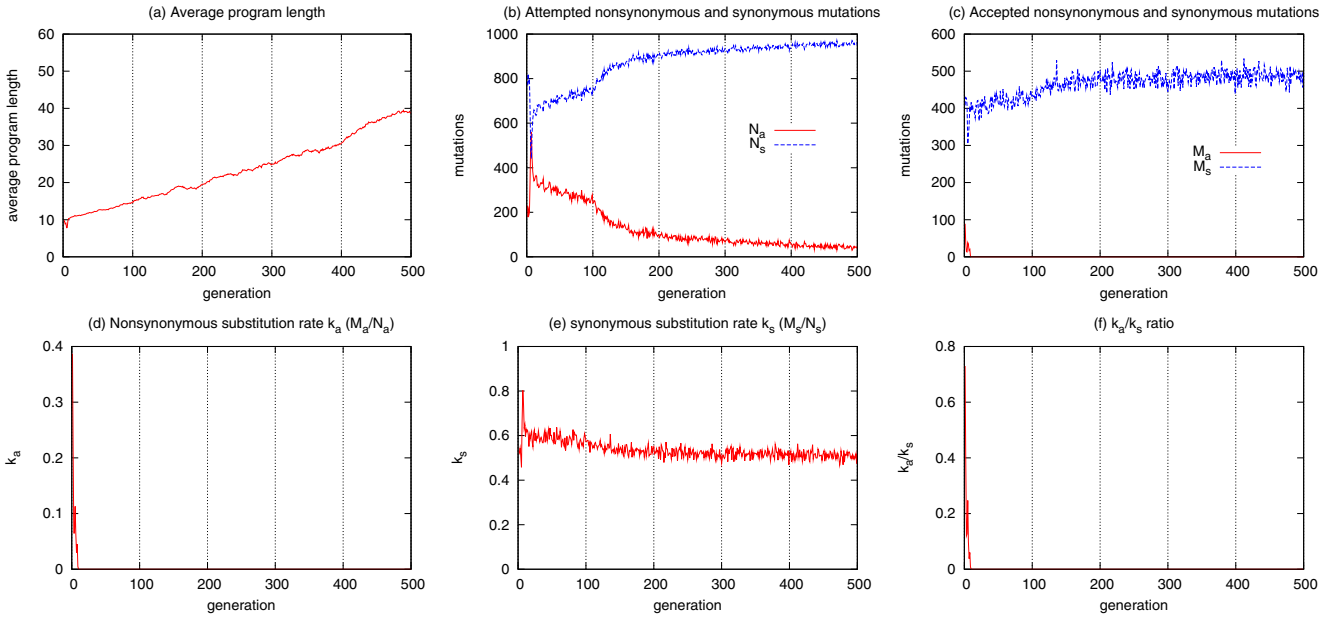


**Figure 3: A typical single run of FT on $x^2 + x$**

tions being deleterious, which is a recognized phenomenon in both biological and artificial evolutionary systems.

We next zoom in to the results of quartic polynomial regression ($x^4 + x^3 + x^2 + x$). Compared to the simpler target of $x^2 + x$, the evolution here takes a longer time to complete, but the general trend of these two runs is the same. The quartic case provides more abundant information to study the process of locating various local and global optima. Here, we plot only the first 200 generations during evolution (Fig. 4). The best fitness and average fitness are plotted in Fig. 4(a), where the fittest individual hits the global optimum at generation 117 and the fitness converges at generation 123. Notice that there is always approximately 5 generations of lag between hitting a local or global optimum and assembling the population to that point. Fig. 4(b) plots the number of individuals that have the same fitness as the fittest individual over time. Observe that the there are 4 periods of frequent replacement of the fittest individual, i.e., generations 10∼15, 70∼75, 85∼90, and 110∼120. As with the previous scenario, we also plot the mutations (Fig. 4(c)), accepted substitutions (Fig. 4(d)), and their relative rates (Fig. 4(e)). In all these measurements, we observe whenever there is frequent replacement of the fittest individual, the system is actively yielding and accepting phenotypic variations. Note that the rate of $k_\mathrm{s}$ remains at approximately a constant level regardless of the system dynamics. However, $k_\mathrm{a}$ faithfully captures the rate at which the system makes observable improvements. Thus, the ratio of $k_\mathrm{a}$ to $k_\mathrm{s}$ also provides a reliable measurement of evolution rate. Consequently, all of the metrics taken in Fig. 4 verify that alternation of the dominance of neutrality and variability is a driving force for evolution throughout time. In addition, the observations we have made here coincide with biological evolution in that i) most random mutations are deleterious so that the $k_\mathrm{a}/k_\mathrm{s}$ ratio is mostly less than 1, and ii) this ratio generally decreases as fitness improvements become finer-grained [19].

In terms of system variability (Fig. 5(b)), we are interested in four points of time during the evolutionary process. That is, at the very beginning, when the fittest individual hits the optimum, when the system converges, and at the end. We observe that the initial population possesses the same high diversity in variability as in random evolution (Fig. 5(a)). As the system evolves, the population has a high overall level of variability but less diverse. When the system converges the optimal fitness and the population starts to possess approximately the same genotypic structure, both the variability and diversity decrease, but the system is still fairly sensitive to mutations. As the evolution progresses to the end of the run and more redundancy accumulates, the entire population has very low variability eventually.

In both the temporal and spatial sense, when a system has a specific target posed by the environment, the coordination between neutrality and variability behaves rather differently from void environmental influences.

### 4.2.3 *Moving Target (MT) Evolution*

We design a moving target by increasing the degree $n$ of the polynomial $\sum_{i=1}^{n} x^i$ periodically. Thanks to the similarity among these targets, there is a good amount of inherent modularity in these environmental changes [13]. In the following experiments, we study how the system responds to such modular changes. At the outset, the system evolves towards the polynomial $x^2 + x$, but we change the target to a higher degree every $c$ generations, called the *switching period*. Specifically, the target is a function of time (or generation) $t$,

$$T(t) = \sum_{i=1}^{\lfloor t/c \rfloor + 2} x^i.$$

Here, we change the target polynomial every 100 generations ($c = 100$), and then allow our Linear GP system to evolve for 500 generations. The target polynomial will increase its degree from 2 to 6. As the polynomial degree increases,

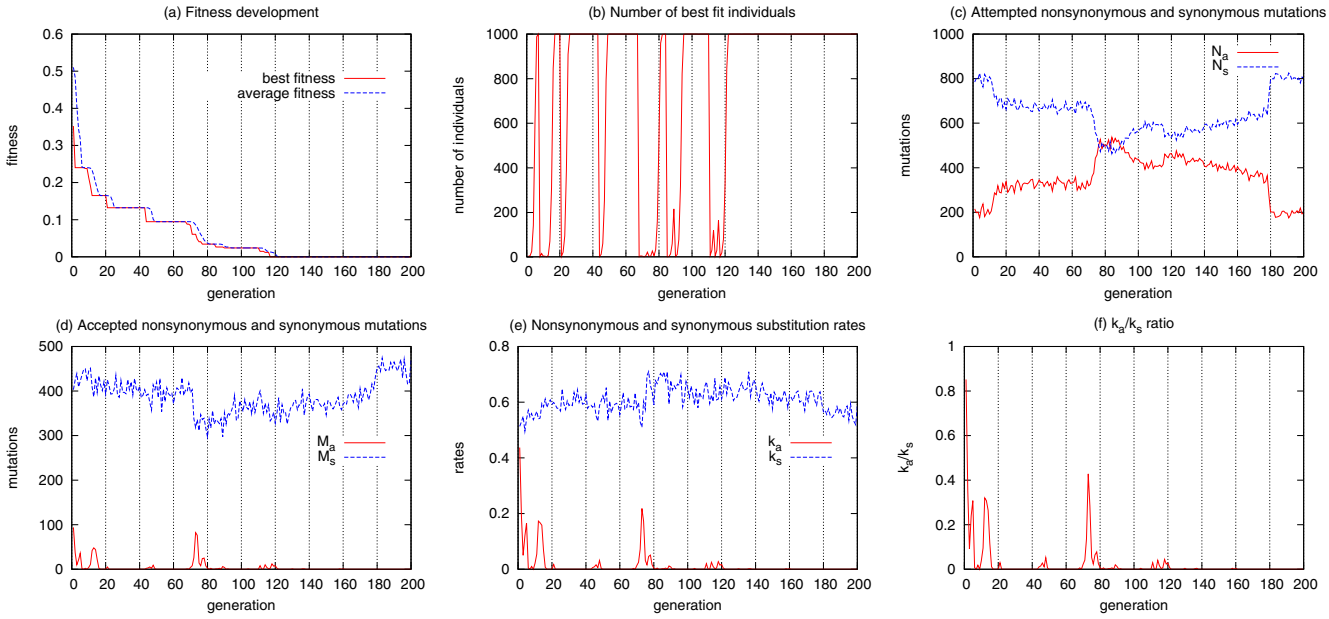


**Figure 4: A typical single run of FT on $x^4 + x^3 + x^2 + x$ during the first 200 generations**

the target takes a more complex form. However, since the target changes in a modular way, there can be many reusable patterns from previous target polynomials, and the search process is expected to learn from history.

Fig. 6 depicts the metrics we looked into also in previous scenarios. Each time the target is switched, we see that the fitness worsens (Fig. 6(a)) and the replacement of the fittest individual happens frequently (Fig. 6(b)). This is similar to the fixed target scenarios, where individuals are becoming sensitive to mutations whenever the system is frequently replacing its currently fittest individual (Fig. 6(c)(d)). It is interesting to see from the chart for the nonsynonymous and synonymous substitution rates (Fig. 6(e)) that, despite the periodic target switching, the synonymous rate $k_s$ still stays fairly stable. This indicates that neutral genetic changes take place and are accepted at a stable rate during the entire evolutionary process, but phenotypic variations can only be observed when the system adapts to its new environment. Again, this results from the close cooperation of neutrality and variability, as two sides of evolvability harnessing random genetic changes to generate adaptive phenotypic variations.

Moreover, carefully designed modular target switching is expected to accelerate evolution. We present three typical runs in Fig. 7 to investigate this. For each case, we plot fitness development and the $k_a/k_s$ ratio. In case 1, the system cannot reach the target in any period before the target moves. As discussed previously, in this case the system is changing very actively. In case 2, the system only finds and converges to the target for the first two periods. In case 3, the system successfully reaches the target by the end of each period. In all of these cases, we observe that, when the target is moved before the system finds and converges to it, the fitness changes are smaller at the target switching point and the $k_a/k_s$ ratio is higher at these points. In contrast, the system is slower to respond to a target change if it has found and converged to a target previously. In terms of neutral networks, as the system finds and converges to a target, the individuals of the system "settle to the center" of the neutral network, and the system becomes more robust. Thus, phenotypic variations start slowly once it is exposed to new environmental challenges. In this case, the individuals need to first move to the edge of the neutral network, i.e., to "pull" the system out of stagnation, before adapting to this new environment. Another observation is that the polynomial target changing in a modular way can improve search efficiency. That is, an evolutionary system can find an ultimate target by following a series of intermediate goals faster than by trying to find it directly. This also suggests some interesting future research on problem modularity and evolvability.

## 5. CONCLUSION AND FUTURE WORK

The most important feature of evolvability is its capability to generate *adaptive* phenotypic variations from random genetic changes. Neutrality and variability are the two sides of evolvability controlling random genetic variations. The environment plays an important role in evolvability to determine which of these two sides is dominant. In this contribution, we employed a Linear GP system as a case study to examine the behavior of evolvability in various environmental situations, by using two tools that can capture evolvability in the temporal and spatial sense. We observed that an evolutionary system actively generates phenotypic variation only when it is adapting to a new environmental challenge. However, this adaptivity is not coming out of void but is the result of constant genetic variations in the background, with the majority being neutral. To cope with environmental fluctuations, a system can improve its phenotypic variation rate without changing its genetic variation rate.

This contribution is preliminary work on understanding the mechanisms in evolvability. It can be extended in the following directions.

- The rate and intensity of environmental fluctuations is expected to have considerable influence on the behavior of evolvability. An environment changing periodically and in a modular fashion is used in the moving target scenario here, with a fixed rate and intensity of the change. More dynamic environments will require further investigation.
- Linear GP as a computation model, the scope of primitive set and genetic operator can also play important roles in the system's variability and neutrality. The impact of these factors will be examined next.
- We expect to implement our methods in other branches of Evolutionary Computation and hope similar conclusions can hold regarding evolvability in other artificial evolutionary systems.
- Our current work focuses on the understanding of evolvability. We hope to be able to explore mechanisms to enhance evolvability when applying EC models in more complex problems.

## 6. ACKNOWLEDGEMENTS

We thank NSERC for support under Discovery Grant RGPIN 283304-07.

## 7. REFERENCES

[1] L. Altenberg. The evolution of evolvability in genetic programming. In *Advances in Genetic Programming*, pages 47–74. MIT Press, Cambridge, MA, USA, 1994.

[2] W. Banzhaf and A. Leier. Evolution on neutral networks in genetic programming. In *Genetic Programming Theory and Practice III*, pages 207–221. Springer, 2006.

[3] M. A. Bedau and N. H. Packard. Evolution of evolvability via adaptation of mutation rates. *BioSystems*, 69(2):143–162, 2003.

[4] T. V. Belle and D. H. Ackley. Code factoring and the evolution of evolvability. In *GECCO '02: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1383–1390. ACM, 2002.

[5] M. Brameier and W. Banzhaf. *Linear Genetic Programming*. Number XVI in Genetic and Evolutionary Computation. Springer, 2007.

[6] S. Collins, J. de Meaux, and C. Acquisti. Adaptive walks toward a moving optimum. *Genetics*, 176(2):1089–1099, 2007.

[7] M. Conrad. The geometry of evolution. *BioSystems*, 24(11):61–81, 1990.

[8] R. Dawkins. The evolution of evolvability. In *Artificial Life: The Quest for a New Creation*, pages 201–220. Addison-Wesley, Reading, MA, USA, 1989.

[9] D. J. Earl and M. W. Deem. Evolvability is a selectable trait. *Proceedings of the National Academy of Sciences*, 101(32):11531–11536, 2004.

[10] M. Ebner, M. Shackleton, and R. Shipman. How neutral networks influence evolvability. *Complexity*, 7(2):19–33, 2002.

[11] T. Hu and W. Banzhaf. Nonsynonymous to synonymous substitution ratio $k_a/k_s$: Measurement for rate of evolution in evolutionary computation. In *Proceedings of the 10th International Conference on Parallel Problem Solving from Nature (PPSN X)*, volume 5199 of *LNCS*, pages 448–457. Springer, 2008.

[12] Y. Jin, R. Gruna, I. Paenke, and B. Sendhoff. Multi-objective optimization of robustness and innovation in redundant genetic representations. In *Proceedings of the 2009 IEEE Symposium on Computational Intelligence in Multicriteria Decision-Making*, pages 38–45. IEEE Press, 2009.

[13] N. Kashtan, E. Noor, and U. Alon. Varying environments can speed up evolution. *Proceedings of the National Academy of Sciences*, 104(34):13713–13716, 2007.

[14] M. W. Kirschner and J. C. Gerhart. Evolvability. *Proceedings of the National Academy of Sciences*, 95(15):8420–8427, 1998.

[15] K. F. Liem. Key evolutionary innovations, differential diversity, and symecomorphosis. In *Evolutioanry Innovations*, pages 147–170. University of Chicago Press, Chicago, IL, USA, 1990.

[16] P. Marrow, M. Heath, and I. I. Re. Evolvability: Evolution, Computation, and Biology. In *GECCO'99: Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop on Evolvability*, pages 30–33. Morgan Kaufmann, 1999.

[17] L. A. Meyers, F. D. Ancel, and M. Lachmann. Evolution of genetic potential. *PLoS Computational Biology*, 1(3):0236–0243, 2005.

[18] C. L. Nehaniv. Measuring evolvability as the rate of complexity increase. In *Proceedings of the Artificial Life VII Workshop*, pages 55–57. MIT Press, 2000.

[19] H. A. Orr. The population genetics of adaptation: the distribution of factors fixed during adaptive evolution. *Evolution*, 52(4):935–949, 1998.

[20] M. Parter, N. Kashtan, and U. Alon. Facilitated variation: How evolution learns from past environments to generalize to new environments. *PLoS Computational Biology*, 4(11):e1000206, 2008.

[21] R. A. Raff. *The Shape of Life: Genes, Development, and the Evolution of Animal Form*. University of Chicago Press, Chicago, IL, USA, 1996.

[22] J. Reisinger and R. Miikkulainen. Selecting for evolvable representations. In *GECCO'06: Proceedings of the 2006 Genetic and Evolutionary Computation Conference*, pages 1297–1304. ACM, 2006.

[23] A. Wagner. Robustness, evolvability, and neutrality. *Federation of European Biochemical Societies Letters*, 579(8):1772–1778, 2005.

[24] A. Wagner. Robustness and evolvability: A paradox resolved. *Proceedings of The Royal Society B*, 275(1630):91–100, 2008.

[25] G. P. Wagner and L. Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976, 1996.

[26] Z. Yang and J. P. Bielawski. Statistical methods for detecting molecular adaptation. *Trends in Ecology and Evolution*, 15(12):496–503, 2000.

[27] T. Yu. Program evolvability under environmental variations and neutrality. In *Proceedings of the 9th European Conference on Advances in Artificial Life (ECAL07)*, pages 835–844. Springer, 2007.

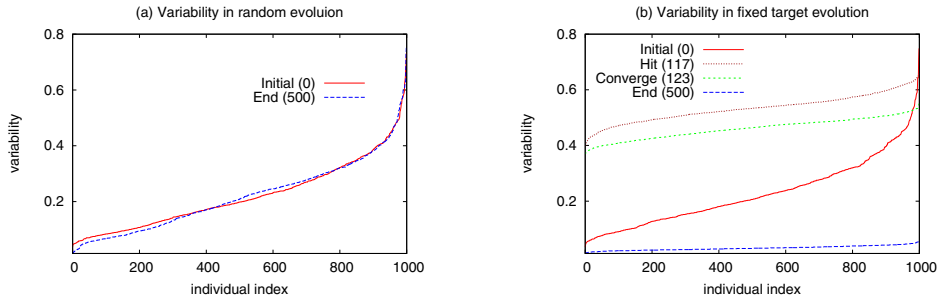


**Figure 5: Variability measurements (numbers in the brackets represent generations)**

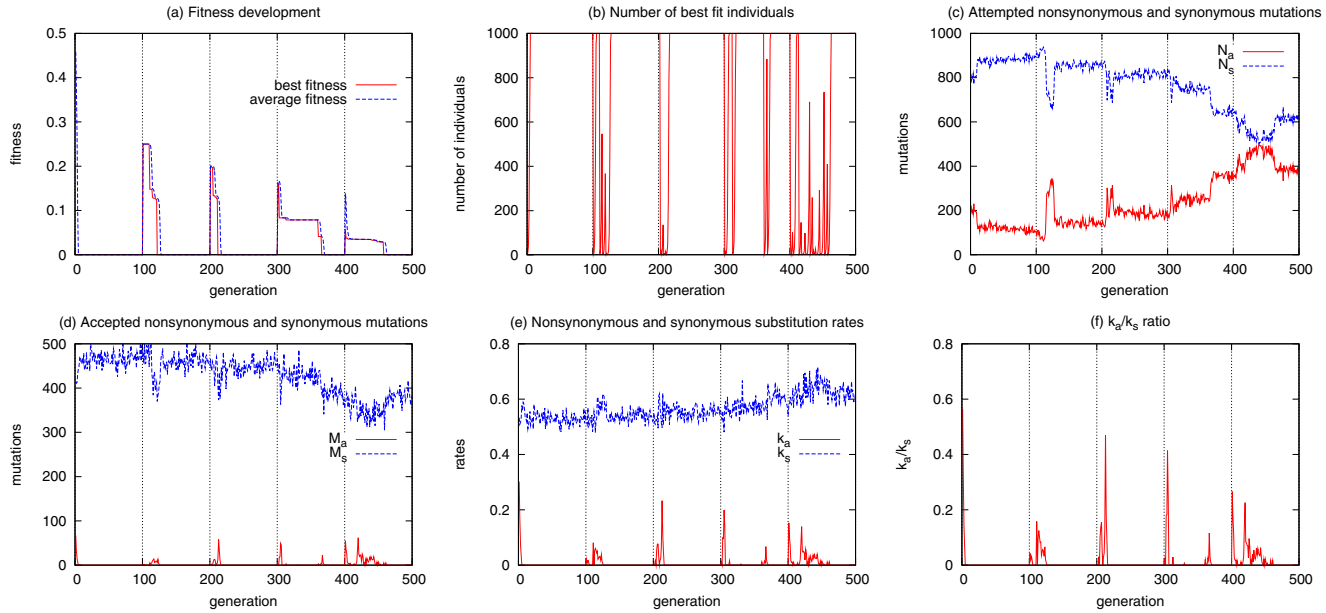


**Figure 6: A typical single run of MT**

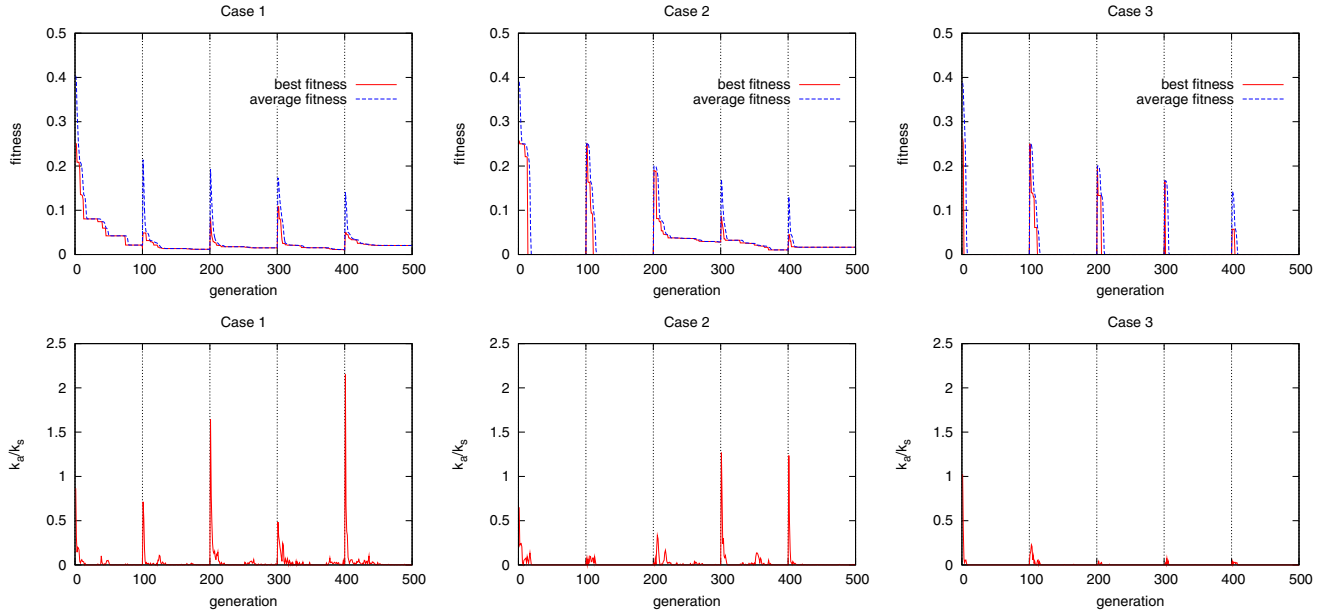


**Figure 7: Three typical cases of MT**