# Evolvability and Static vs. Dynamic Fitness

**Matthew R. Glickman** and **Katia Sycara**

School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213-3891

**Abstract**

Evolvability is key to the adaptive potential of an evolving system. Evolvability itself may be subject to evolution, and this process can greatly enhance or substantially inhibit a system's adaptive potential. In artificial systems, this process can often occur irrespective of the designer's awareness or intention. Sustaining evolvability is of particular importance for extended and possibly open-ended evolution. We recap some of the research done on evolvability and its potential to evolve, and identify a common hazard of this process, the premature stagnation of search via the evolutionary suppression of variability. We describe some approaches to forestalling this problem in the context of static fitness regimes. We then propose that dynamic fitness regimes are a possible key to sustaining variability, and offer some candidates for the form of such regimes.

## Introduction: Evolvability and Variation

Although clearly of central importance to investigators concerned with the potential of evolutionary systems, evolvability is a somewhat difficult property to define precisely. Nonetheless, whether evolvability is sufficiently present in a system is usually quite clear, *i.e.* "you know it when you see it."

Viewed from the standpoint of Evolutionary Computation (EC) or Artificial Life, biological organisms in their extraordinary diversity and range of fine adaptation represent the height of evolvability. On the other hand, constructing an artificial system that demonstrates a profound lack of evolvability is quite straightforward.

For instance, to evolve a computer program to perform some given task, we might choose to represent the programs in some common programming language (making the encoded program an individual's *genotype*). Then, at the point when we've selected the "fittest" programs in the population to reproduce to form the next generation, we might choose to induce variation in their offspring by randomly mutating a few characters in their ASCII representation. Such a scheme would appear destined for failure–in a program of any appreciable fitness, practically any conceivable change would simply degrade its behavior (or *phenotype*). In fact, automated program evolution didn't attract large-scale interest until evolvability was improved via the identification of representations (*e.g.* LISP-like, functional languages) and mechanisms for stochastic variation (*e.g.* subtree-crossover and language-token mutation) that could yield new adaptive variants with sufficient probability.

This example illustrates the dependence of evolvability upon phenotypic *variability* (Wagner and Altenberg 1996). In this case, variability (as distinct from variation) refers to the range of phenotypic variation that is produced when the genotype is stochastically perturbed. In the example above, phenotypic variability refers to the the expected range of behaviors produced when, given a program with some particular behavior, you mutate a few of characters of its ASCII representation.

As pointed out in (Altenberg 1994b), the specific nature of this process of parent-to-offspring phenotypic *transmission* is determined by two elements of the system: (1) the *representation*–the mapping from genotype-to-phenotype–and (2) the *search operators*, or *variation mechanism*, which stochastically modify genotypes upon reproduction.

Our general concern with evolvability in this paper has to do with its own potential to be shaped by evolution. Because evolvability is dependent on phenotypic variability being of some proper form, the next section focuses on how traits having to do with variability may evolve. The rest of the paper then generally addresses when the evolution of variability may or may not lead to evolvability. We come to focus on the particular hazard of premature stagnation of evolutionary search via the evolutionary suppression of variability. We present some of the approaches to this problem in the common EC context of static fitness, and then propose that dynamic fitness may provide a natural path to sustaining variability. We suggest a few models upon which such fitness schemes may be based.

## Selection for Variability

Although we may speak of the evolvability inherent in particular choices of representation and variation mechanisms by virtue of their general effect on phenotypic variability, there also exists variation in variability at the individual- and/or population-level. For instance, when two phenotypically similar individuals are each given the opportunity to reproduce, one may tend to yield offspring with phenotypes that closely resemble that of their parent(s), while the other may yield offspring that

vary more widely. This variance in variability opens up the possibility for selection to shape variability and thus, for evolvability to itself evolve.

As it turns out, the action of selection on variability is less than straightforward. Selection only discriminates on the basis of traits with direct impact on the fitness function, and an individual's phenotypic variability is only observed once it has been selected and has the opportunity to reproduce. Therefore, the action of selection on traits which affect variability occurs indirectly: Individuals with high fitness–those selected to reproduce–will *tend* to have been those which were produced by parents with superior variability characteristics, *i.e.* parents that are more evolvable. Dependent upon the heritability of such characteristics, the selected individuals will themselves tend to be more evolvable.

At the same time, this *a posteriori* correlation between high direct fitness and high evolvability will not always hold; sometimes (*e.g.* near a local optimum) directly measurable fitness and evolvability will be negatively correlated. In such a case, selection will favor direct fitness; an individual, $A$, with high fitness but low evolvability will be favored by selection over another individual, $B$, with lower fitness but higher evolvability.

For this reason, selection's ability to adapt evolvability is enhanced by decoupling traits that affect variability from those which affect direct fitness and allowing them to vary independently. One way to provide for this capacity is via the representation. Representations which support this capacity provide for multiple genotypes which map to the same direct phenotype, but which differ in variability. For instance, given such a representation, there would exist the potential for the individual, $A$, in the example above to produce offspring that inherit its fitness, but are also more evolvable. Examples of such representations include the use of self-adapting mutation rates (Schwefel 1981; Fogel 1991), in which the representation is explicitly extended to include variability-related traits (mutation rates), and the canonical LISP-like representation of Genetic Programming (GP, (Koza 1992)), the nature of which implicitly provides for adapting variability (Altenberg 1994b).

It is important to note, however, that the ability of selection to operate on variability-related traits is not strictly dependent on such traits being decoupled from direct fitness. Even without a many-to-one genotype-to-phenotype map, in the example of $A$ and $B$ above, the descendants of $B$ may ultimately prevail over those of $A$ as long as (1) the fitness scores of $B$ and its descendants remain high enough that they are never entirely excluded from reproduction until (2) $B$'s superior variability characteristics ultimately result in some of its descendants achieving fitness scores higher than those of $A$.

## The Utility of Adapting Variability

Situations in which the adaptation of variability may be key to reasonable search performance are not difficult to imagine. Say, for instance, that two interdependent traits important to fitness differ vastly in their sensitivity to variation. To be able to simultaneously tune these two characters, it may well be necessary for variability to evolve so as to dampen the effects of variation on the sensitive trait and/or magnify them on the insensitive trait.

Nevertheless, the evolution of variability-related traits does not necessarily always promote evolvability. Experiments we conducted using self-adapting mutation rates on a very simple bit-string problem (Glickman and Sycara 1998) showed that allowing mutation rates to evolve on a per-individual basis often yielded poorer performance than using a reasonable fixed mutation rate. It was only when the problem was sufficiently scaled up–the bit-strings were extended to a sufficient length–that evolution of mutation rates began to provide a benefit. In another study (Glickman and Sycara 2000), we found that adding an explicit variability-related trait (mutation rates, once again) to a representation that already had an implicit capacity for independent adaptation of variability rapidly suppressed the evolvability of the system. Such a premature loss of evolvability is in fact a general hazard of evolving variability, and is discussed in more detail in the next section.

## Evolvability and Static Fitness

### Premature Stagnation of Search

In EC, it is common that the procedure used to compute the fitness score remains unchanged during the entire course of evolution. Such a *static fitness regime*, when coupled with a representation that provides for relatively facile adaptation of variability independent of objective fitness, can lead to the premature stagnation of evolutionary search.

To see why, note that the more fit the phenotype of an individual, the more it stands to lose from being stochastically perturbed; as fitness increases, variation becomes increasingly costly. We thus observe that as the population increases in fitness, selection will (to the extent to which it is possible to do so) tend to favor individuals that increasingly conserve their fitness and exhibit less and less phenotypic variation.

From the standpoint of the designer, this is not necessarily a negative property; it seems reasonable for search to shift from coarse to increasingly fine-grained variation over time. However, in EC the goal of the designer is at the population level, *e.g.* an evolving population which ultimately leads to a population or individuals of maximal fitness. Selection, however, operates at the level of the individual, simply favoring those individuals that maximize their *reproductive success* (*i.e.* yield the maximum number of offspring that survive and reproduce). There can come a point where, configurations of higher-fitness occur so rarely that reproductive success nearly completely correlates with simply preserving parental fitness, via the suppression of phenotypic variability. At this point, effective search comes to a halt as quickly as possible.

Altenberg (1994a; 1994b) discusses the evolutionary suppression of variability via the phenomenon of "code bloat" in Genetic Programming.

## Methods to Prolong Exploration

There are a number of ways to prolong exploration in the face of static fitness, perhaps the simplest of which is to simply increase the selection strength.

With stronger selection, a smaller number of individuals get to reproduce, but those that do get more opportunities to do so. At the same time, parents are under increased pressure to produce highly fit offspring, *i.e.* offspring sufficiently fit to be one of the few that will reproduce in the next generation. This state of affairs favors risking more disruption of the parental phenotype for the chance of producing one or more offspring of exceptional fitness (*e.g.* see (Glickman and Sycara 1998; 1999); Blickle and Thiele (1994) show a relationship between selection strength and code bloat in GP).

However, because a smaller number of individuals account for a greater fraction of the reproduction, population diversity may also be reduced. Because the effectiveness of recombination for search in sexually reproducing populations is dependent upon population diversity, there is thus a tendency for practitioners to minimize selection strength. Recognition that weaker selection strength may also suppress phenotypic variability indicates a danger of this tendency.

Altenberg (1994b) proposed a mechanism called *soft brood selection* designed to counteract the tendency toward individuals of overly conservative variability in Genetic Programming. With soft brood selection, the number of offspring produced by parents is magnified. Selection is then first applied among offspring of the same parents, and those which are selected in this round then go on to compete at the full population level.

The intuition behind this procedure is that if only the best of a parent's offspring will survive anyway, then (as with stronger selection) the parent can afford to produce some poorer quality offspring for the chance of producing a few of significantly higher fitness.

A third approach to prolonging exploration when evolving variability is exemplified by the technique of *Nested Evolution Strategies*[1] (Herdy 1992). As discussed above, the problem of premature stagnation of search can be seen as a conflict between the population-level goal of the designer and the nature of selection, which acts (primarily) at the level of the individual. One approach to evolving a population-level property is to promote *group selection*. Group selection occurs between groups rather than individuals, and is generally held to be a significantly weaker force in nature than individual-level selection. This is because it requires that groups are sufficiently isolated that they, as is commonly the case with individuals, will independently suffer the consequences of, or derive the benefits from, their own traits. At the same time, highly fit populations must then be able to "reproduce", *i.e.* pass their traits on to other populations, which remain otherwise isolated.

Group selection appears to be what nested evolution strategies achieve by copying the primary population to form a number of isolated sub-populations, which then undergo independent evolution. Individual sub-populations differ in their variability characteristics, and over time, those sub-populations with superior evolvability will tend to produce higher fitness individuals. Periodically, the sub-populations are themselves subjected to selection, resulting in the best one reproducing to form a new set of sub-populations. Experimenters have reported (Rechenberg 1999) that such a scheme derives a much greater performance benefit from self-adapting mutation rates than the traditional, single-population method.

[1] A description is available via the WWW at http://www.bionik.tu-berlin.de/intseit2/xs2evost.html

## Evolvability and Dynamic Fitness

In nature, selection takes place within an ecology, with organisms' fitness greatly dependent upon their mutual interactions. Fitness in nature can thus be a much more dynamic affair than is common in evolutionary computation.

We argue that the implications of dynamically changing fitness for the evolution of evolvability are potentially great. The problem of evolutionary stagnation via loss of variability becomes much less troubling because when the sense of fitness changes with sufficient frequency, the capacity to vary becomes a critical trait (*e.g.* (Maley 1995) found that co-evolutionary pressure favored higher mutation rates). Dynamic fitness *focuses* selection on evolvability.

For examples of the form of dynamic fitness that may promote the evolution of evolvability, we may perhaps look to hypotheses proposed to explain the origin and maintenance of sexual reproduction (a clearly variability-related trait) in evolving populations. One fitness regime is the well-known, contrarian dynamic of the Red Queen, in which the environment grows to disfavor whichever traits or types were previously most common. Such a dynamic is characteristic in the presence of antagonistic, co-evolving populations such as parasites.

A classic example of such co-evolution in the artificial domain is that of (Hillis 1992) involving the evolution of sorting networks. In Hillis' experiments, the fitness evaluation of each individual network is based on a series of tests using a set of *fitness cases* (vectors of numbers to be sorted, in this study). Such a scheme allows the establishment of an antagonistic, co-evolving population of fitness cases whose fitness is determined by how challenging they are for the original population.

A distinct but related fitness dynamic proposed to explain selection for sexual reproduction is that of the Tangled Bank, proposed by Bell (1982). The environment of the tangled bank is partitioned into a large number of ecological niches, each of which can support a limited number of individuals of a particular phenotype. Because such an environment might possibly admit a stable array of independent micro-populations–which would clearly disfavor phenotypic variability–it is necessary that niches periodically vanish with some frequency for at least a brief interval. This situation would then favor individuals to produce a variety of offspring phenotypes which may then come to occupy a variety of niches. Related methodologies in EC include schemes

for *niching* or *fitness sharing* (Goldberg and Richardson 1987).

An alternative to the co-evolving population of fitness cases may be to instead co-evolve a population that draws from a family of problems related to the original fitness function. The value of such a scheme would be the potential to evolve populations that adapt to the structure common to a family of tasks, and which may then prove particularly effective when applied to other, previously unseen tasks in the same family.

The potential implementation approaches discussed above are oriented toward explicitly fostering particular fitness dynamics to improve evolvability with the goal of solving given problems in optimization and/or design. A potentially much richer forum for investigating the interaction between dynamic fitness and evolvability might be found in more artificial life-style simulations, such as Ray's Tierra (Ray 1991), in which such dynamics arise implicitly.

## Summary

While evolvability is key to the adaptive potential of an evolutionary system, the evolution of evolvability via the action of selection on traits influencing phenotypic variability is does not always enhance this potential. In particular, with a static notion of fitness, selection may favor the suppression of variability before the potential for adaptation is exhausted. Approaches to forestalling this problem include increasing selection strength, soft brood selection, and promoting group selection. We've further proposed that dynamic fitness regimes may significantly obviate this problem by consistently rewarding variability. Because reproducing sexually is clearly a variability-related trait, models proposed to explain the origin and maintenance of sexual reproduction in evolving populations may provide a source of ideas about the forms of dynamic fitness that may promote sustained variability.

## Acknowledgements

We would like to thank Frank Dellaert, Michael Mateas, and two anonymous reviewers for their very helpful comments on earlier drafts of this document. Our research is partially supported by the Office of Naval Research, Contracts N00014-95-1-0591 and N00014-96-1-1222.

## References

Lee Altenberg. 1994. Emergent phenomena in genetic programming. In Anthony V. Sebald and Lawrence J. Fogel, editors, *Proceedings of the Third Annual Conference on Evolutionary Programming*, pages 233–241. World Scientific.

Lee Altenberg. 1994. The evolution of evolvability in genetic programming. In K. Kinnear, editor, *Advances in Genetic Programming*. MIT Press

Graham Bell. 1982. *The Masterpiece of Nature: The Evolution and Genetics of Sexuality*. University of California Press, Berkeley.

T. Blickle and L. Thiele. 1994. Genetic programming and redundancy. In J. Hopf, editor, *Genetic Algorithms within the Framework of Evolutionary Computation (Workshop at KI-94, Saarbrcken)*. Max-Planck-Institut fr Informatik (MPI-I-94-241).

D. B. Fogel, L. J. Fogel, and J. W. Atmar. 1991. Meta-evolutionary programming. In R. R. Chen, editor, *Proceedings of 25th Asilomar Conference on Signals, Systems and Computers*, pages 540–545, Pacific Grove, California.

Matthew Glickman and Katia Sycara. 1998. Evolutionary algorithms: Exploring the dynamics of self-adaptation. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, San Francisco, CA. Morgan Kaufmann.

Matthew Glickman and Katia Sycara. 1999. Comparing mechanisms for evolving evolvability. In A. S. Wu, editor, *GECCO-99: Workshop Proceedings of the Genetic and Evolutionary Computation Conference.*

Matthew R. Glickman and Katia Sycara. 2000. Reasons for premature convergence of self-adapting mutation rates. In *CEC2000: Proceedings of the Congress on Evolutionary Computation.*

David E. Goldberg and Jon Richardson. 1987. Genetic algorithms with sharing for multimodal function optimization. In J. Grefenstette, editor, *Proceedings of the Second International Conference on Genetic Algorithms and their Applications*, pages 41–49, Hillsdale, New Jersey. Lawrence Erlbaum Associates.

M. Herdy. 1992 Reproductive Isolation as Strategy Parameter in Hierarchically Organized Evolution Strategies. In R. Männer and B. Manderick, editors, *Parallel Problem Solving from Nature, 2*, Amsterdam. Elsevier.

Daniel Hillis. 1992 Co-evolving parasites improve simulated evolution as an optimization procedure. In *Artificial Life II.* Addison Wesley.

John R. Koza. 1992 *Genetic Programming: On the Programming of Computers by the Means of Natural Selection.* MIT Press.

Carlo Maley. 1995. The coevolution of mutation rates. In M. Moran and C. Morelo, editors, *Advances in Artificial Life*, Berlin. Springer-Verlag.

Thomas Ray. 1991. An approach to the synthesis of life. In *Artificial Life II.* Addison Wesley.

Ingo Rechenberg. 1999. Artificial evolution in action. Invited talk, Genetic and Evolutionary Computation Conference, July 16, 1999.

H. P. Schwefel. 1981. *Numerical Optimization of Computer Models.* Wiley, Chichester.

Günter P. Wagner and Lee Altenberg. 1996. Complex adaptations and the evolution of evolvability. *Evolution*, 50:967–976.