Reconciling Explanations for the Evolution of Evolvability

Bryan Wilder and Kenneth O. Stanley

To Appear In: *Adaptive Behavior* journal. London: SAGE, 2015.

Abstract

Evolution's ability to find innovative phenotypes is an important ingredient in the emergence of complexity in nature. A key factor in this capability is evolvability, or the propensity towards phenotypic variation. Numerous explanations for the origins of evolvability have been proposed, often differing in the role that they attribute to adaptive processes. To provide a new perspective on these explanations, experiments in this paper simulate evolution in gene regulatory networks, revealing that the type of evolvability in question significantly impacts the dynamics that follow. In particular, while adaptive processes result in evolvable *individuals*, processes that are either neutral or that explicitly encourage divergence result in evolvable *populations*. Furthermore, evolvability at the population level proves the most critical factor in the production of evolutionary innovations, suggesting that nonadaptive mechanisms are the most promising avenue for investigating and understanding evolvability. These results reconcile a large body of work across biology and inform attempts to reproduce evolvability in artificial settings.

Keywords: Evolvability, evolutionary computation, gene regulatory networks

Introduction

Something about the structure of biological systems allows evolution to find new innovations. An important question is whether such evolvability (defined roughly as the propensity to introduce novel phenotypic variation) is itself a product of selection. In essence, is evolvability evolvable (Pigliucci, 2008)? Much prior work has attempted to encourage evolvability in artificial systems in the hope of reproducing the open-ended dynamics of biological evolution (Reisinger & Miikkulainen, 2006; Grefenstette, 1999; Bedau et al., 2000; Channon, 2006; Spector, Klein, & Feinstein, 2007; Standish, 2003). Additionally, biological research has investigated conditions that may be favorable for the emergence of evolvable structures and explored various mechanisms that could be responsible for evolvability (Ciliberti, Martin, & Wagner, 2007; Crombach & Hogeweg, 2008; Draghi & Wagner, 2009; Steiner, 2012; Daniels, Chen, Sethna, Gutenkunst, & Myers, 2008; G. P. Wagner & Altenberg, 1996; A. Wagner, 2005, 2008a, 2008b; Kirschner & Gerhart, 1998; Sniegowski & Murphy, 2006). Intriguingly, there is no consensus across this work on whether selection can favor evolvability.

Rather than consensus, a number of different mechanisms for producing evolvability have been proposed. For example, many experiments have suggested that evolvabilty can be increased by selection in changing environments (Earl & Deem, 2004; Kashtan, Noor, & Alon, 2007). A common technique is to vary the target of selection in a modular fashion, leading up to the desired goal (Kashtan & Alon, 2005; Kashtan et al., 2007). Similarly, several studies have proposed incremental fitness functions that focus on one portion of the task at a time (Mouret & Doncieux, 2008; Urzelai & Floreano, 1999; Gomez & Miikkulainen, 1997). These experiments suggest that varying the fitness function in this manner promotes the development of modular building blocks that can be assembled to form better behaviors than would arise under a static fitness function. A large body of biological work has also shown that evolvability can be produced by environments that vary over time or space (Draghi & Wagner, 2009; Crombach & Hogeweg, 2008; Steiner, 2012; Palmer & Feldman, 2011). The main idea is that if selection sets a moving target, individuals will be more likely to introduce variation in their offspring to adapt to an uncertain future; mutations to the genotype will accordingly be more likely to result in phenotypic change.

These results contrast with the success of nonadaptive techniques. For example, the novelty search algorithm (Lehman & Stanley, 2008, 2011) ignores the objective of an evolutionary search entirely and selects only for behavioral diversity. Novelty search has been shown to evolve better neural networks for tasks such as navigating a maze or controlling a biped walker than conventional evolutionary algorithms, despite the fact that the desired goal is never incorporated into selection (Lehman & Stanley, 2011). These results dovetail with biological studies that emphasize neutral processes that work around adaptive pressure. Studies by Wagner and other collaborators (A. Wagner, 2005, 2008a, 2008b; Ciliberti et al., 2007) emphasize the importance of phenotypic robustness in evolvability. Because many genotypes map to the same phenotype, a large amount of genetic variation can build up without changing the phenotype and thereby decreasing fitness. That way, organisms in different parts of the genotype space, while sharing the same phenotype, can access radically different phenotypes in their immediate mutational neighborhood. This genetic diversity then allows the population as a whole to access a wide variety of phenotypes, while avoiding the impact of selection entirely. According to this line of work, lower levels of phenotypic variation actually *promote* evolvability. Ebner, Shackleton, and Shipman (2001) suggest leveraging phenotypic robustness in this manner to improve evolutionary algorithms.

Another set of results independently question the importance of selection in producing evolvability. In particular, the limited capacity of environmental niches may be sufficient to produce highly evolvable individuals. In several such experiments by Lehman and Stanley (2013), there is no selection towards a preferred phenotype: all niches are equally viable. However, when niches have a limited carrying capacity, individuals who find previously undiscovered niches benefit from a founder effect, resulting in an increase in evolvability over time.

These varied perspectives also reflect a lack of clarity on how precisely to define evolvability. The hypothesis considered in this paper is that a comprehensive theory of evolvability requires distinguishing between *evolvable individuals* and *evolvable populations.* On the one hand, evolvable individuals are more likely than others to introduce phenotypic variation in their offspring. On the other hand, in evolvable populations a greater amount of phenotypic variation is accessible to the population as a whole, regardless of how evolvable any individual may be in isolation. This distinction will help to disentangle conflicting results and theories on the relationship between evolvability and selection. In particular, results about selection in variable environments deal with evolvable individuals, while results based on nonadaptive processes have focused on evolvable populations. It is important to note that population-level evolvability is not equal to the sum over individual evolvability because the novel phenotypes contributed by different individuals may be redundant. That is, population-level evolvability measures the extent to which individuals contribute *distinct* innovations to the pool that is available to the group.

To understand this distinction, consider that that adaptive pressures have differing effects on the average individual and the population considered as a whole. In the presence of varying selection, individuals may gain the ability to yield more diversity to cope with the likelihood of change in the future. However, the population as a whole could emerge with less ability to explore alternate genetic paths because adaptive pressure penalizes deviations from the individually optimal strategy. Thus the population might miss the opportunity for innovations that would have been discovered without such adaptive pressure. It is therefore natural that introducing an adaptive regime can result in individuals who produce greater levels of variation, but at a cost to the population. In effect, there is no incentive for the population to maintain diversity; individuals will tend to cluster around locally optimal points in the space.

However, selection can mean more than adaptive pressure towards particular phenotypes. Limited-capacity niches push the population towards greater levels of diversity by rewarding those who introduce novel behaviors. Negative frequency-dependent selection explicitly encourages phenotypic diversity because uncommon behaviors have greater fitness (Fitzpatrick, Feder, Rowe, & Sokolowski, 2007; Gigord, Macnair, & Smithson, 2001). These forms of selection, which we collectively call *divergent selection*, could encourage evolvability at the population level by selecting for individuals that find strategies uncommon in the rest of the population. It should be noted that any form of selection can be viewed as "adaptive" in the general sense that it rewards individuals for meeting some criteria. However, the distinction between selective processes that push towards a particular target behavior and those that simply encourage diversity is still useful because of their differing consequences (as this paper will highlight). Therefore, in this paper, selection is called *adaptive* if it pushes phenotypes towards a particular target behavior. In contrast, in *divergent* forms of selection (such as negative frequency-dependent selection or novelty search), selection does not focus on an externally determined target. Instead, it depends on the composition of the population in a manner that rewards behavioral diversity. This difference in terminology is designed to draw a clear rhetorical distinction between the two forms of selection.

The study in this paper simulates evolution in a gene regulatory network model to demonstrate the impact of divergent and adaptive selection on evolvability. The simulations span different forms of environmental change in adaptive regimes, and also cover negative frequency-dependent selection as an example of divergent selection. The results demonstrate that while adaptive selection produces evolvable *individuals*, only divergent selection results in *populations* that can collectively access the most nearby variation. A key implication of these experiments is that such population-level evolvability is the deeper issue at stake in explaining how evolution can produce complexity. Ultimately, individuals become important to evolution in aggregate through their contribution to the population. The key question then for the ability of the population to continue to innovate is not whether any individual in isolation will introduce variation, but what innovations can be found by the population considered all together. This insight further implies that nonadaptive mechanisms are the more important means for creating complex evolutionary systems because they produce populations oriented towards innovation.

Experiment

To investigate the implications of different forms of evolvability, evolution is simulated with a model of artificial gene regulatory networks (GRNs). GRNs have been widely studied because of their foundational role in determining gene activation and the corresponding expression of phenotypic traits. Several studies in particular have addressed the evolution of evolvability in GRNs, demonstrating both the impact of environmental variation and the role of neutral processes (Draghi & Wagner, 2009; Ciliberti et al., 2007; Steiner, 2012; Crombach & Hogeweg, 2008), thereby establishing their utility in investigating the mechanisms behind evolvability.

Gene regulation model

We adopt a model of gene regulation proposed by Draghi and Wagner (2009), which is similar to a broad range of widely used models (Bergman & Siegal, 2003; Azevedo, Lohaus, Srinivasan, Dang, & Burch, 2006; Siegal, Promislow, & Bergman, 2007; Ciliberti et al., 2007; Steiner, 2012) that originated with the work of A. Wagner (1996). An individual's genome is a directed graph consisting of k vertices, each one representing a transcriptional regulator. Edges between these vertices represent regulatory influence: positive weights are activating influences, while negative weights are repressing. Weights are confined in the interval $[-1, 1]$. The genome is represented by a matrix W, where W_{ij} is the weight of the edge from gene i to gene j. If no such edge is present, $W_{ij} = 0$.

The phenotype is a pattern of gene expression produced by this network. Each gene is either active (1) or inactive (0). A developmental process unfolds wherein the state vector is iterated from an initial condition in which all genes are active until a fixed point is reached. The vector of states at time t is $S(t)$, and the expression level of gene i is $S_i(t)$. Starting from $S(0) = \mathbf{1}$, the state is updated according to

$$S_i(t+1) = \begin{cases} 1 & \text{if } \sum_{j=1}^{k} W_{ji} S_j(t) > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Because each genes expression can take 2 values, at least one configuration must be revisited within 2^k time steps. Because the dynamics are deterministic, once a configuration is repeated, the network will either reach a fixed point $S(\infty)$ where $S(t+1) = S(t)$ or cycle through a trajectory of states. An individual is *viable* if a fixed point is reached, and nonviable otherwise. For computational efficiency, individuals are deemed nonviable if no fixed point is found after 500 iterations.

Evolution is simulated with a population of N individuals. The population begins from identical clones of a single ancestor, and then evolves according to selection and mutation. Each individual's fitness is scored according to one of two methods:

In *adaptive selection*, the individual's phenotype $S(\infty)$ is compared to a target phenotype T. Let d be the Hamming distance between $S(\infty)$ and this target, or the number of genes for which the expression value differs. Then adaptive fitness can be computed as

$$F_A = \frac{1}{(1+s)^d}. \tag{1}$$

Here, s is a parameter that tunes the strength of selection. Nonviable individuals are assigned $d = k+1$, making them less fit than all viable individuals.

In contrast, in *divergent selection*, individuals are penalized according to the number of others who share the same phenotype. Let n_i be the number of individuals with the same phenotype as individual i. The corresponding fitness is

$$F_D = \frac{1}{(1+s)^{n_i}}. \tag{2}$$

The manner of selection is the same as for adaptive selection, but individuals are scored exclusively based on their phenotypic uniqueness instead of comparison to an optimal expression pattern.

Evolutionary simulation

At each generation, N individuals are drawn from the population to reproduce with probability proportional to their fitness. Individuals may be selected multiple times. Each offspring is a clone of its parent, with some probability of mutation. To make comparisons possible, all populations start from the distribution of neutral evolution; a founding individual is generated with random values for W_{ij}

and then 500 mutation steps are performed to allow the weights to equilibrate. In mutation, existing edges are mutated with probability μ by adding a random variable distributed as $\mathcal{N}(0, 0.1)$ to the weight. Topological mutations are made with probability μ_A. That is, each existing connection is removed and each missing connection is added with independent probability μ_A.

Several different modes of selection are investigated:

1. Static adaptive selection: The target phenotype is set at the start of the run and remains constant throughout, with fitness assigned according to Equation (1).
2. Fluctuating adaptive selection: The target phenotype changes every p generations, with fitness assigned according to Equation (1). This mode of selection simulates evolution in a changing environment. Two types of such changes are explored. First, in *random fluctuating selection*, the target is randomly assigned a distinct value. Second, when *modularly varying goals* (MVG) are used, only half of the target is changed each time, alternating between $T_{1:k/2}$ and $T_{(k/2)+1:k}$. Thus, the target is modified in a modular fashion, segmented between the two halves of genes. This scenario replicates the conception of MVG introduced by Kashtan and Alon (2005).
3. Divergent selection: Fitness is assigned according to F_D as described in Equation (2). This regime pressures individuals to achieve previously undiscovered phenotypes.
4. Neutral evolution: No selection is present and the population evolves strictly according to genetic drift.

Together, these forms of selection capture the major hypotheses that have been proposed on evolvability in the literature.

Results

Evolution was simulated for 2,000 generations with a population of $N = 100$ individuals with $k = 10$ genes. For all runs with selection shown here, $s = 1$ (varying s does not substantially impact the results, as shown in Figures S1-S6). Mutation probabilities were fixed at $\mu = 10^{-2}$ and $\mu_A = 10^{-3}$. For fluctuating selection, $p = 100$. Every 200 generations, both the evolvability of individuals and the evolvability of the population as a whole is recorded.

A variety of definitions and quantifications for evolvability have been proposed (Pigliucci, 2008). Of ultimate interest is the ability of an individual or population to discover phenotypic variation. Therefore, a measurement of evolvability should characterize the amount of variability that can be accessed in an individual or population's genetic neighborhood. Following Lehman and Stanley (2013), we measure evolvability by sampling 500 topological mutations for each individual in the population. An individual's evolvability is the number of distinct phenotypes that are discovered by this process, while the evolvability of a population is the number of phenotypes that are discovered by all individuals combined. Topological mutations are specifically considered because substantively new phenotypes more often result from structural changes to the regulatory network, which is typically robust to alterations in the biochemical parameters represented by weights (Von Dassow, Meir, Munro, & Odell, 2000; Ingolia, 2004).

Simulation results confirm that both fluctuating and divergent selection improve individual-level evolvability. Figure 1 shows the mean individual-level evolvability of the population over 2,000 generations. Divergent selection produces the highest individual-level evolvability. Fluctuating selection produces individuals that are the next most evolvable, more so than static selection or neutral evolution. This result is consistent with numerous findings documenting the beneficial impact of environmental variation on evolvability (Earl & Deem, 2004; Kashtan et al., 2007; Draghi & Wagner, 2009; Crombach & Hogeweg, 2008; Steiner, 2012; Palmer & Feldman, 2011). Neutral evolution's low level of evolvability can be attributed to the number of nonviable phenotypes which are encountered: fully

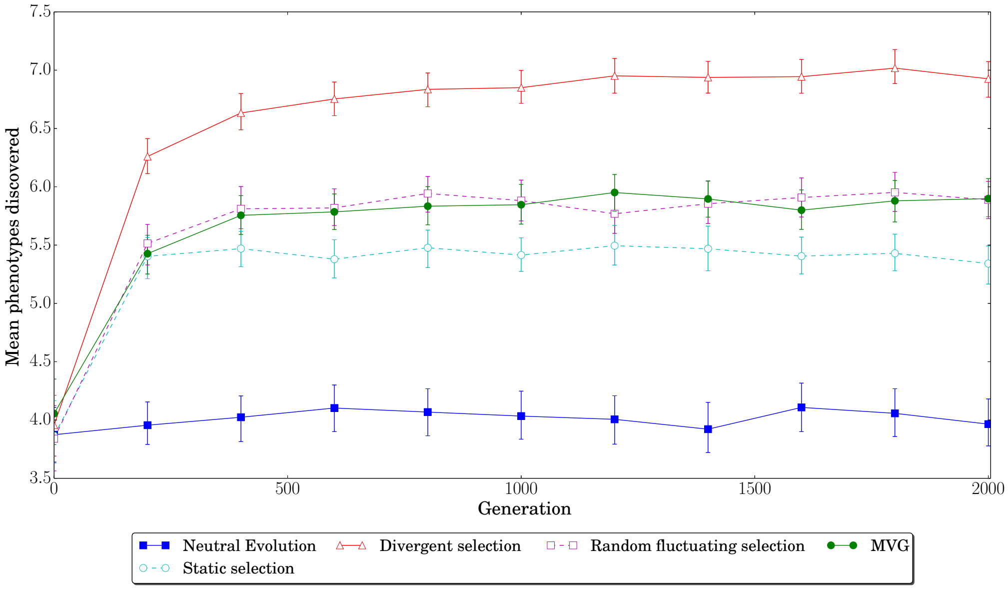

Figure 1: Individual-level evolvability. 500 replicate runs of the simulation are averaged for each mode of selection. Error bars are bootstrapped 99% confidence intervals.

55% of topological mutations of members of the neutrally evolved population resulted in a nonviable phenotype. All other modes of selection guide the population towards regions of the genotype space containing valid networks. As a result, only about 8% of mutations resulted in a nonviable phenotype for all modes except neutral evolution.

However, the results are dramatically different for population-level evolvability. As shown in Figure 2, divergent selection leads to the highest degree of evolvability again, though now by a higher margin. However, fluctuating selection actually results in equal or marginally *lower* levels of evolvability than neutral evolution (the difference is not statistically significant in all generations). Even when the population is placed in a changing environment, pressure towards particular targets prevents the population from accessing some phenotypes. That is, even though the average individual subject to fluctuating selection will tend to produce a greater degree of phenotypic variation than one subject only to genetic

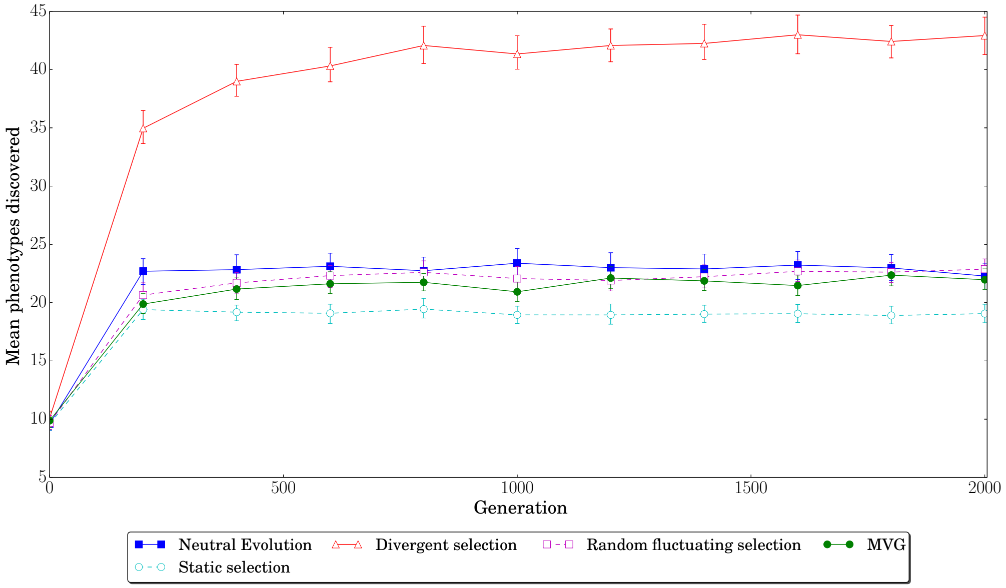

Figure 2: Population-level evolvability. 500 replicate runs of the simulation are averaged for each mode of selection. Error bars are bootstrapped 99% confidence intervals.

drift, the population as a whole has fewer innovations available within its mutational neighborhood. Static selection unsurprisingly displays the lowest degree of population-level evolvability, reflecting the pressure for the population to remain within a single region.

This population-level difference can also be seen in the level of genetic diversity within the population under different modes of selection. Figure 3 shows the mean Euclidean distance between individuals in the genotype space. The distance between two individuals with weight matrices A and B is

$$\sqrt{\sum_{i=1}^{k}\sum_{j=1}^{k}(A_{ij} - B_{ij})^2} \tag{3}$$

and Figure 3 plots the mean of this quantity over all pairs of individuals in the population. Adaptive selection, either static or fluctuating, results in a low level of diversity. Fluctuating selection, despite producing much more evolvable individuals and populations, results in slightly less diversity than static selection. That is, the difference in diversity is very minor, and so static selection's higher diversity is not enough to contribute a substantial number of new phenotypes. For both cases, the restriction imposed on the population by the presence of adaptive selective pressure in the first place prevents a level of diversity that is even comparable to neutral evolution. Even though each individual in fluctuating evolution has access to a greater number of phenotypes than under neutral evolution, the population as a whole tends to be clustered around more similar areas of the genotype space that give the optimal ability for each individual's descendants to adapt to changing circumstances. However, this redundancy in the distribution of genotypes prevents the population from accessing areas that are less lucrative for the individual but that nevertheless contribute new phenotypes to the set accessible by the population. Results from Ciliberti et al. (2007) confirm that accessing the majority of possible phenotypes requires individuals to be located in widely varying portions of the genotype space. Therefore, it is natural that adaptive selection pressure should limit the evolvability of the population.

Divergent selection produces a higher level of diversity than any mode of adaptive selection. Interestingly though, the relative gap between divergent selection and other modes of evolution is narrower in terms of genetic diversity than population-level evolvability. That is, evolvability is not directly proportional to diversity. The implication is that the population under divergent selection is steered towards particular regions of the space from which a large number of phenotypes are accessible. While individuals are still more dispersed than under adaptive selection, the success of divergent selection does not rely on individuals inhabiting wildly different areas of the genotype space. Instead, divergent selection maintains diversity that is concentrated in promising areas. This observation accords with findings by Lehman, Stanley, and Miikkulainen (2013) that maintaining genetic diversity alone is not sufficient to increase performance in deceptive domains such as maze navigation. That is, genetic diversity alone does not translate directly into an evolvable population. It is crucial that the population

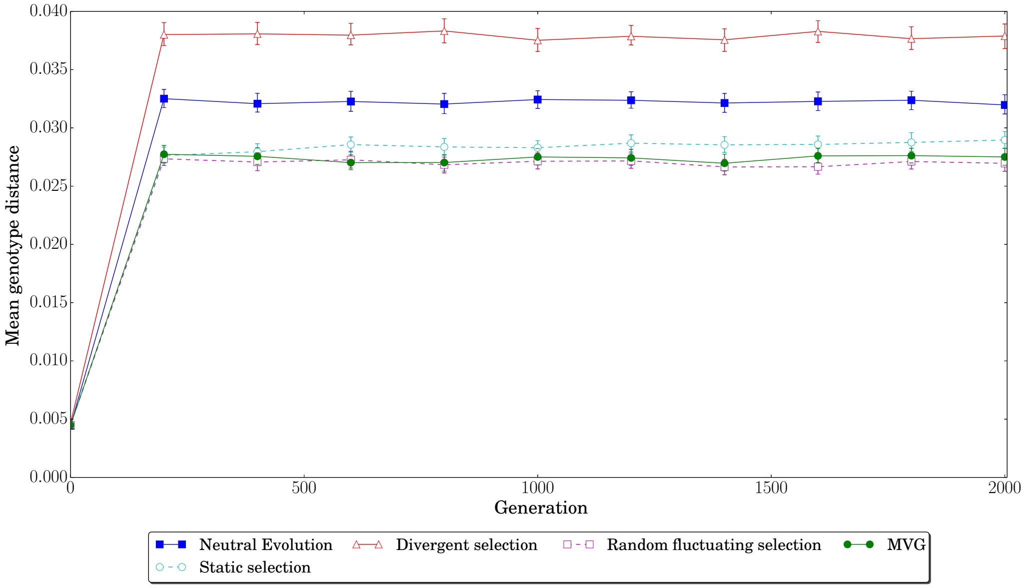

Figure 3: Genetic diversity in the population. Values represent the mean Euclidean distance between individuals in the genotype space. 500 replicate runs of the simulation are averaged for each mode of selection. Error bars are bootstrapped 99% confidence intervals.

also be located in genetic regions that lend themselves to the discovery of new phenotypes.

Discussion

The experimental results highlight the need to distinguish carefully between different senses of evolvability. Processes that produce evolvable individuals may be counterproductive when evolvable populations are at issue and vice versa. While divergent selection can bring about evolvability at both levels, more traditional forms of adaptive selection are effective only at producing evolvable individuals.

This insight helps reconcile conflicting proposals for reproducing evolvability in artificial systems. A large body of literature (Earl & Deem, 2004; Kashtan et al., 2007; Kashtan & Alon, 2005; Mouret

& Doncieux, 2008; Urzelai & Floreano, 1999; Gomez & Miikkulainen, 1997) has focused on explicitly selecting for greater evolvability through variation in the fitness function. On the face of it, there is a tension between this work and methods that generate behavioral diversity instead of setting a single objective for the entire population. Such divergent processes fundamentally allow the population to access more variation, because it is the group, not a single individual, that diffuses through the set of genotypes and discovers widely varying phenotypes in different portions of the space. However, the tension disappears once it is made clear which sense of evolvability is under consideration. Populations become evolvable when greater exploration is promoted, and individuals become evolvable when selection is for phenotypic variation.

These observations highlight the importance of population-level evolvability. Evolution is inherently a process that operates at the level of the population, while the evolvability of an individual in isolation affects only the fitness of its descendants. Questions about how computational systems can emulate the open-ended nature of biological evolution are thus addressed more comprehensively by the characteristics of the population as a whole. Even if each individual in a population has a number of phenotypes accessible via mutation, the potential for innovation is limited if these new phenotypes are largely the same. Diversity is important to evolution insofar as different individuals contribute different possibilities for future development. One implication is that research on producing evolutionary complexity can benefit from looking past traditional modes of selection. Other processes with less focus on adaptation are essential to innovation because new discoveries require the population to explore new possibilities instead of being directed narrowly by the constraints of fitness.

Our understanding of divergent selection continues to improve. Future work can investigate how the genotype-phenotype mapping interacts with selective pressures. It is important to understand the mechanisms through which divergent selection locates portions of the genotype space with access to a wide variety of phenotypes. This work could help characterize the properties of such regions, for example by identifying topological features that are common in highly evolvable networks. Beyond divergent selection as considered here, another promising direction for future work is to identify additional nonadaptive mechanisms that shape the emergence of evolvability.

Conclusion

This paper compared evolvability at the level of the individual versus the population, a distinction that underlies disagreements in previous work on creating evolvable systems. Experiments with simulated evolution in gene regulatory networks showed that evolvable individuals are produced by adaptive selection in a changing environment, while evolvable populations are produced by divergent selection, which encourages phenotypic diversity.

These results reconcile conflicting notions of how evolvability can be replicated in artificial settings. While varying the target of selection may be an improvement over a static environment, the population as a whole is most evolvable when selection explicitly promotes exploration instead of focusing on a specific objective. This result suggests that future efforts to create open-ended evolutionary complexity are best served by making use of divergent selection as opposed to adaptive measures.

Funding statement: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

Azevedo, R. B., Lohaus, R., Srinivasan, S., Dang, K. K., & Burch, C. L. (2006). Sexual reproduction selects for robustness and negative epistasis in artificial gene networks. *Nature*, *440*(7080), 87–90.

Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., ... Ray, T. S. (2000). Open problems in artificial life. *Artificial life*, *6*(4), 363–376.

Bergman, A., & Siegal, M. L. (2003). Evolutionary capacitance as a general feature of complex gene networks. *Nature*, *424*(6948), 549–552.

Channon, A. (2006). Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines*, *7*(3), 253–281.

Ciliberti, S., Martin, O. C., & Wagner, A. (2007). Innovation and robustness in complex regulatory gene networks. *Proceedings of the National Academy of Sciences*, *104*(34), 13591–13596.

Crombach, A., & Hogeweg, P. (2008). Evolution of evolvability in gene regulatory networks. *PLoS computational biology*, *4*(7), e1000112.

Daniels, B. C., Chen, Y.-J., Sethna, J. P., Gutenkunst, R. N., & Myers, C. R. (2008). Sloppiness, robustness, and evolvability in systems biology. *Current opinion in biotechnology*, *19*(4), 389–395.

Draghi, J., & Wagner, G. (2009). The evolutionary dynamics of evolvability in a gene network model. *Journal of evolutionary biology*, *22*(3), 599–611.

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

Ebner, M., Shackleton, M., & Shipman, R. (2001). How neutral networks influence evolvability. *Complexity*, *7*(2), 19–33.

Fitzpatrick, M. J., Feder, E., Rowe, L., & Sokolowski, M. B. (2007). Maintaining a behaviour polymorphism by frequency-dependent selection on a single gene. *Nature*, *447*(7141), 210–212.

Gigord, L. D., Macnair, M. R., & Smithson, A. (2001). Negative frequency-dependent selection maintains a dramatic flower color polymorphism in the rewardless orchid dactylorhiza sambucina (l.) soo. *Proceedings of the National Academy of Sciences*, *98*(11), 6253–6255.

Gomez, F., & Miikkulainen, R. (1997). Incremental evolution of complex general behavior. *Adaptive Behavior*, *5*(3-4), 317–342.

Grefenstette, J. J. (1999). Evolvability in dynamic fitness landscapes: A genetic algorithm approach. In *Evolutionary computation, 1999. cec 99. proceedings of the 1999 congress on* (Vol. 3).

Ingolia, N. T. (2004). Topology and robustness in the drosophila segment polarity network. *PLoS biology*, *2*(6), e123.

Kashtan, N., & Alon, U. (2005). Spontaneous evolution of modularity and network motifs. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(39), 13773–13778.

Kashtan, N., Noor, E., & Alon, U. (2007). Varying environments can speed up evolution. *Proceedings of the National Academy of Sciences*, *104*(34), 13711–13716.

Kirschner, M., & Gerhart, J. (1998). Evolvability. *Proceedings of the National Academy of Sciences*, *95*(15), 8420–8427.

Lehman, J., & Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *Alife* (pp. 329–336).

Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189–223.

Lehman, J., & Stanley, K. O. (2013). Evolvability is inevitable: Increasing evolvability without the pressure to adapt. *PLOS ONE*, *8*(4), e62186.

Lehman, J., Stanley, K. O., & Miikkulainen, R. (2013). Effective diversity maintenance in deceptive domains. In *Proceeding of the fifteenth annual conference on genetic and evolutionary computation conference* (pp. 215–222).

Mouret, J.-B., & Doncieux, S. (2008). Incremental evolution of animats behaviors as a multi-objective optimization. In *From animals to animats 10* (pp. 210–219). Springer.

Palmer, M. E., & Feldman, M. W. (2011). Spatial environmental variation can select for evolvability. *Evolution*, *65*(8), 2345–2356.

Pigliucci, M. (2008). Is evolvability evolvable? *Nature Reviews Genetics*, *9*(1), 75–82.

Reisinger, J., & Miikkulainen, R. (2006). Selecting for evolvable representations. In *Proceedings of the 8th annual conference on genetic and evolutionary computation* (pp. 1297–1304).

Siegal, M. L., Promislow, D. E., & Bergman, A. (2007). Functional and evolutionary inference in gene networks: does topology matter? *Genetica*, *129*(1), 83–103.

Sniegowski, P. D., & Murphy, H. A. (2006). Evolvability. *Current Biology*, *16*(19), R831–R834.

Spector, L., Klein, J., & Feinstein, M. (2007). Division blocks and the open-ended evolution of development, form, and behavior. In *Proceedings of the 9th annual conference on genetic and evolutionary computation* (pp. 316–323).

Standish, R. K. (2003). Open-ended artificial evolution. *International Journal of Computational Intelligence and Applications*, *3*(02), 167–175.

Steiner, C. F. (2012). Environmental noise, genetic diversity and the evolution of evolvability and robustness in model gene networks. *PloS one*, *7*(12), e52204.

Urzelai, J., & Floreano, D. (1999). Incremental evolution with minimal resources. *Proceedings of IKW99*, 796–803.

Von Dassow, G., Meir, E., Munro, E. M., & Odell, G. M. (2000). The segment polarity network is a robust developmental module. *Nature*, *406*(6792), 188–192.

Wagner, A. (1996). Does evolutionary plasticity evolve? *Evolution*, 1008–1023.

Wagner, A. (2005). Robustness, evolvability, and neutrality. *FEBS letters*, *579*(8), 1772–1778.

Wagner, A. (2008a). Neutralism and selectionism: a network-based reconciliation. *Nature Reviews Genetics*, *9*(12), 965–974.

Wagner, A. (2008b). Robustness and evolvability: a paradox resolved. *Proceedings of the Royal Society B: Biological Sciences*, *275*(1630), 91–100.

Wagner, G. P., & Altenberg, L. (1996). Perspective: Complex adaptations and the evolution of evolvability. *Evolution*, 967–976.