

Improved Adaptation in Exogenously and Endogenously Changing Environments

Joshua R. Nahum^{1,2}, Jevin West³, Benjamin M. Althouse^{4,5,6}, Luis Zaman^{1,7},
Charles Ofria^{1,2} and Benjamin Kerr^{1,7}

¹BEACON Center for the Study of Evolution in Action, Michigan State University, East Lansing, MI 48824

²Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824

³Information School, University of Washington, Seattle, WA 98195

⁴Institute for Disease Modeling, Bellevue, WA 98005

⁵Santa Fe Institute, Santa Fe, NM 87501

⁶New Mexico State University, Las Cruces, NM 88003

⁷Department of Biology, University of Washington, Seattle, WA 98195
nahumjos@msu.edu

Abstract

Fitness landscapes are visual metaphors that appeal to our intuition for real-world landscapes to help us understand how populations evolve. The object inspiring the metaphor is better described as a networks composed of all possible genotypes, but they are frequently simplified to a surface where the fitness of each genotype is represented by elevation. Selection drives evolving populations to ascend the landscape until they are dominated by genotypes from which no further beneficial mutations are likely, known as a peak. However, by allowing for environmental change, former peaks can vanish, forcing populations to resume adapting. To explore how changing environments affect adaptation, we used the digital evolution platform, Avida, wherein we could manipulate the organisms' environment as they are subject to natural evolutionary forces. We found that transient exposure to alternate environments frequently resulted in more fit genotypes. Negative-frequency-dependent environments, in particular, yielded strong fitness benefits after returning to the original environment. Furthermore, we explored how such environmental change could yield adaptive benefits via valley crossing and how such knowledge could be exploited in systems where improving the rate of adaption is beneficial.

Introduction

On some level, all evolutionary studies involve exploration of how organisms adapt to their environment over some period of time. Adaptation by natural selection requires time in an environment to allow advantageous alleles to appear and rise in frequency. A reasonable expectation would be that longer exposure to an environment should yield increased adaptation, as there is more time for beneficial mutations to arise and spread. However, since the fitness effects of new mutations typically depend on the genetic background in which they occur, a population may be unable to cross through low fitness states to arrive at the genotype of highest fitness, no matter how much time is spent evolving.

Sewall Wrights "adaptive landscape" metaphor can be used to demonstrate a situation where adaptation is not limited solely by time (Wright, 1932). An adaptive landscape relates an organisms genotype to its fitness. Imagine if all individual genotypes of an asexual organism could be placed together on a plane where the distance between two genotypes represents the number of mutations needed to generate one genotype from another. Each genotype is assigned a height directly proportional to its fitness within a specified environment. Each organism can be represented as a point (located at its genotype) and the population as a whole will be a cloud of points. This cloud spreads its range through mutation, but is pruned by selection which drives it to shift its weight uphill (see Figure 1a). Thus the combination of mutation and selection leads to the population "climbing" fitness hills to their "peak," which is a genotype from which all mutations are deleterious (downhill). If the landscape is rugged, (i.e., it has multiple peaks), merely climbing uphill does not guarantee the population reaches the most fit genotype (Kauffman and Levin, 1987). The population may become "trapped" on a sub-optimal peak, a problem Wright addressed with his Shifting Balance Process.

The primary idea behind the Shifting Balance Process is that small populations can escape sub-optimal adaptive peaks through genetic drift and begin climbing in a new place. One assumption made to simplify the model is that the environment remains constant during evolution (the peaks maintain their positions). Interestingly, a *changing environment* may provide an alternative mechanism to escape from sub-optimal peaks (Fisher, 1930; Wright, 1932; Whitlock, 1997; Collins et al., 2007). Different environments, by definition, have differences in the mapping from genotype to fitness (e.g., mutations detrimental in one environment may be beneficial in another). As an environment changes, former adaptive peaks may disappear and new peaks can appear. In this reshaped adaptive landscape, even a large pop-




Figure 1: Adaptive Paths in Hypothetical Changing Landscapes Here we consider two kinds of environmental change. The population is represented as one or more spheres that climb on the adaptive landscape. (a) A population initially stuck on a sub-optimal peak experiences an exogenous environmental shift, resulting in a different adaptive landscape (illustrated in green). On this new landscape, the selection allows the population to move to a new position in genotype space. Upon returning to the original environment, the population can (in this case) climb to a new, higher peak. (b) A population, which was initially stuck on a sub-optimal peak, experiences an endogenously changing environment. Through negative frequency dependence, the fitness of the genotype at the former peak is depressed (illustrated as a green depression in the landscape), allowing selection to favor a diversity of new genotypes (multiple spheres). When the population returns to the original environment, these diverse genotypes initiate trajectories to multiple fitness peaks.

ulation, which evolved to a former peak, can be selected to move to a novel genotypic position. Upon returning to the original environment, the population can climb up to a new, perhaps higher, peak (see Figure 1b).

Changing environments can have profound effects on adaptation (Waxman and Peck, 1999; Ancel Meyers et al., 2005; Parter et al., 2007). Using a variety of different computational systems, Kashtan et al. (Kashtan et al., 2007) found that when the environment changed in specific ways

(i.e., when goals from different environments shared sub-problems), the population was able to more rapidly evolve solutions to complex problems than in a constant environment. Alto et al. (Alto et al., 2013) found that alternating exposure of vesicular stomatitis virus to two different temperatures led to increased performance in both temperatures relative to virus evolved in a constant environment. In these studies, change was externally imposed on the evolving populations. Here we additionally investigate a qualitatively different form of environmental change, where organisms are the causative agents of their environment.

We classify two different types of change in the environment: exogenous and endogenous. Exogenously driven change is brought about by factors outside the influence of the population (e.g., diurnal-nocturnal changes in light or seasonal changes in temperature). In contrast, the evolving population itself produces endogenous change. Actions of organisms in the population, including resource usage, waste production and habitat modification, affect the environment and alter the fitness landscape for themselves and future generations. The process whereby organisms modify their environments has been termed niche construction (Laland et al., 1996; Heino et al., 1998; Odling-Smee et al., 2003; Sterelny, 2005) or ecosystem engineering (Jones et al., 1994). Such niche construction can, of course, feed back to affect the evolution of the population. For example, if organisms reduce the quality of their environment by exploiting certain resources or become susceptible to infectious pathogens, they can depress the fitness of their own, and related, genotypes (see Fig. 1b and (Laland et al., 1999)). This change may select for novel genotypes that differ from those currently constituting the population. Importantly, both exogenous and endogenous environmental change allows populations to leave (formerly) adaptive peaks by selection.

To distinguish the effects of exogenously and endogenously changing environments, we need a system that can exhibit natural evolutionary dynamics in highly controlled settings. Ideally, this system would be simple (to deduce the fitness effects of individual mutations), fast (to allow for evolution across many generations), and tractable (to control the environment, and the effect organisms have on it). For these reasons we chose to perform our experiments with Avida, a computational platform for the study of evolution. Avida provides an ideal system to test the effects of a changing environment and has been used extensively to investigate a wide array of evolutionary questions (Ofria and Wilke, 2004). The organisms within Avida are simple (the mappings between genotype, phenotype, and fitness are easily determined) and fast (generations last less than a second), but the genetic space is still vast—typical organisms have 100 loci with 26 characters per site, allowing for more than 3×10^{141} meaningful combinations. The environment for an evolving population can be measured and manipulated precisely. Most importantly, we can explore changing envi-

ronments: either via exogenous change or by allowing the digital organisms themselves to influence the environment.

System

Avida is an evolution platform, wherein digital organisms (Avidians) can evolve in world with a capacity of 3600 organisms. Each Avidian has a genome composed of a sequence of simple computational instructions. For this experiment, the length of the genome was fixed at 100 instructions. When assembled in particular configurations, these instructions perform functions related to asexual replication or numerical computation. An Avidian's fitness (replication rate) can be improved either by increasing their replication efficiency or by "metabolizing" resources in the environment by performing mathematical tasks specified by the experimenter. Many different globally available resources can be present in the environment, each paired with a particular task. During the replication of an Avidian, the mutation rate is the frequency at which an instruction being copied into an offspring genome is instead substituted with a random instruction. Upon completion of replication, the offspring is placed in the world at a random site, supplanting any previous occupant. Each Avida run was seeded with an identical, self-replicating ancestor that initially could not perform any tasks. The unit of time in Avida is an "update," which is the period for the average organism to perform 30 instructions (for the data presented here, one generation is approximately seven updates).

The environment within Avida is defined by the abundances of available resources and their associated computational tasks. If an organism successfully performs a task, it is rewarded by increasing its "metabolic rate" proportional to the abundance of the resource associated with the task. The metabolic rate of an organism determines how quickly it can execute instructions relative to the other individuals in the population, and thus heavily influences how quickly it can produce offspring. "Rigid" environments have unchanging resource concentrations that are not influenced by the tasks performed. The resources in a "Malleable" environment flow into the world at a fixed rate in a chemostat-like manner and are consumed by organisms when associated tasks are executed. The improvement in metabolic rate associated with task performance is proportional to the amount of the associated resource available; hence, in a Malleable environment, the consumption of resources reduces their availability to other organisms.

In our experiments, every evolutionary run was broken into three equal-length periods, where the environment in each period was either rigid or malleable. The *Fixed* treatment uses the same rigid environment for each period. For all other treatments, the first and third periods are the same rigid environment (termed the "reference") as the *Fixed* runs, while the second period is a different environment (termed the "alternative"). The middle period in the *Flipped*

treatment is a rigid environment, but with a different set of resources (i.e., where different tasks are rewarded). The middle period of the *Negative Frequency Dependent (NFD)* treatment is a malleable environment, where the resources available in the reference environment become consumed when their task is performed. All populations (regardless of treatment) were evolved for 100,000 updates (approximately 12,000 generations).

Environment

For the reference environment, the resources associated with the tasks Not, Nand, And, Nor, Xor and Equals are present in essentially infinite amounts, leading to no measurable depletion when the associated task is performed (see Ofria and Wilke (2004) for details regarding tasks and resources). In the same manner, the alternate environment (the middle third of the *Flipped* treatment) contains only the resources associated with the tasks: OrNot, Or, and AndNot. The fitness reward (merit) for successfully completing each task is provided only the first time an organism performs it and is proportional to the complexity of the task (the number of nand instructions needed to compose the logical operation). During the middle third of frequency-dependent runs, the rewarded resources are the same as the reference, however, the resources have an inflow (100 units per update) and outflow rates (0.01 proportion of concentration per update). The inflow and outflow rates are determined such that if an organism is the only one capable of performing a task, it will receive the same reward as the reference environment. The reward of each task is proportional to the resource concentration in a Michaelis-Menten manner. For the middle third of the *Negative Frequency Dependence* runs, performing a task consumed one unit of the resource, whereas in the *Positive Frequency Dependence* runs, performing a task increased the abundance of the resource associated with the task by 1.

Results

Evolution in Static and Dynamic Worlds

To assess the effect of a changing environment on evolution, we need a baseline for adaptation in an unchanging environment. To obtain this baseline, we evolved populations in a single fixed environment (where rewards for resources did not change). At the conclusion of each evolutionary run in this *Fixed* treatment, we selected the most abundant genotype and determined its line of descent to the ancestral genotype. The fitness trajectory along such a line of descent is shown in Figure 2a for an example *Fixed* population. For contrast, we include an example fitness trajectory from a population evolved in a changing environment (the *Flipped* treatment) in Figure 2b. Here the shaded middle third of the evolutionary run represents exposure to an alternate environment (whereas the population evolves in the reference environment from Fig. 2a for the first and last




Figure 2: Fitness in Static and Exogenously Changed Environments (a) The fitness trajectory of the line of descent over the course of a single *Fixed* treatment run. The black line denotes the fitness of the genotype in the reference environment. (b) A fitness trajectory from a single *Flipped* treatment run, with the time in the alternate environment during the middle third shaded in blue. The blue line denotes the fitness of the line of descent in the alternate environment, whereas the black line gives fitness in the reference environment. Averaging 60 replicates, we next show mean fitness from the *Fixed* treatment (c) and the *Flipped* treatment (d). (e) The final mean fitness from *Fixed* and *Flipped* treatments, as measured in the reference environment. Ribbon and error whiskers denote standard error of the mean.

third of the run). The blue trajectory during the middle third represents the fitness of the genotypes in the alternate environment (whereas the black trajectory gives fitness in the reference environment). In both static and dynamic worlds, nearly all mutations in the line of descent are beneficial with respect to the present environment.

However, some mutations that are favored in the alternate environment (a lift in the blue line) would have been detrimental in the reference environment (a drop in the black line). The average of 60 runs of *Fixed* and *Flipped* treatments (Fig. 2c and 2d, respectively) demonstrate that fitness gains decrease over time to a plateau. However, populations exposed to an alternate environment (*Flipped* treatment) reach significantly higher fitness in the reference environment at the end of the run, regardless of the specific resources being rewarded (Fig. 2e, Mann-Whitney test, $p=0.03$).

One possible advantage that populations evolving in the *Flipped* treatment had relative to those in the *Fixed* treatment, is a greater availability of beneficial trajectories over the course of a run. If populations were exhausting potential beneficial mutations early in the run (leading to constraints in adaptive potential), organisms in the reference environment should have few possible beneficial mutations available, but more possible beneficial mutations with respect to the alternate environment. To investigate this hypothesis, we extracted the organism in the lineage that existed at the conclusion of the first third of the run. We then constructed every possible single mutation at every genome position and evaluated the fitness of each genotype in both the reference and alternate environment (Figure 3). After excluding mutations that were detrimental in both environments or nearly

neutral in at least one environment, we found that mutations beneficial in the reference environment constituted a small minority (0.65%, Quadrants 1 and 4) of the possible mutations, while a much greater fraction of mutations were detrimental in the reference environment but beneficial in the alternate environment (99.35%, Quadrant 2). Thus, mutations beneficial in the alternate environment are often detrimental in the reference environment. This result implies that the evolution in the alternate environment yields more selectively beneficial mutational steps and may lead to genotypic movement that could not have taken place in the reference environment. The improved fitness outcome in the *Flipped* treatment over the *Fixed* treatment can be explained by a number of explanations, including: greater breadth of search of the adaptive landscape, potential weak positive correlation between reference and alternative landscapes, potential correspondence between height of peak and breadth of basin, among others.

Evolution in Frequency-Dependent Environments

Although populations can experience environmental change that is exogenous in origin (as in the *Flipped* treatment), they can also be the actors of change. To model such a situation, we constructed an environment where resources were finite and consumable (whereas the reference environment had an infinite abundance of such resources). A small amount of each resource is continually flowing into the world, but when the resource is consumed (when an organism performs its associated task) its availability is reduced and the fitness reward for further performance of its task decreases. In such an environment, when a phenotype increases in frequency, its relative fitness decreases because less resource is avail-




Figure 3: Mutational Distribution at the Environment Transition We isolated a single genotype from the line of descent from a *Fixed* run that existed at update 33,000 immediately before an exogenous shift would have occurred in the *Flipped* treatment. All possible single mutations were introduced into this genotype and the fitness repercussions in both the reference and alternate environment were measured. In panel b, each point corresponds to a mutation and is positioned according to its fitness consequences in each of the environments. Note, mutations neutral in both environments would fall on the origin. Red and blue points denote mutations beneficial and detrimental in the reference environment, respectively. Mutations that are detrimental in both environments or have a fitness effect less than 5% in either environment are not shown. Panel a shows the transformed density distribution ($\log(10 * \text{abundance} + 1)$) according to *ndr0* kernel smoothing) of non-excluded mutations according to fitness in the reference environment.

able to any one individual, a situation termed *Negative Frequency Dependence (NFD)*. To examine the effect of endogenously driven change we mirrored the structure of the first set of runs. Specifically, we applied the *NFD* environment to the middle third (leaving the first and last third as the rigid reference environment). Populations evolved in the *NFD* treatment were better adapted to the reference environment than populations evolved in the *Fixed* treatment (Figure 4a and 4b, Mann-Whitney, $p=0.0075$).

From the results in the *Flipped* treatment, we know that environmental change can improve fitness. Was the effect of *NFD* merely a consequence of environmental change that accompanies the resource consumption? To address this question, we ran an additional treatment (*Paired Transplant*) where a population evolved in the environment generated by

a separate *NFD* run. Specifically, we first evolved one population in the *NFD* treatment, where it affected its resources during the middle third of the run. We then evolved a second population (the focus of the *Paired Transplant* treatment) using the precise resource levels available in the changing environment of the first. While this second population fully experienced the fluctuating resources of the first population, the latter population was completely unable to alter the environment itself. This treatment isolates the effect of change alone (i.e., without the feedback between the population and environment). The mean final fitness of the *Paired Transplant* runs was significantly lower than the *NFD* runs (Fig. 4c, Mann-Whitney, $p<0.001$), indicating that environmental change alone does not account for the adaptive benefits conferred by negative frequency dependence. Thus the interchange between the environment and populations is necessary for the increased adaptation in the *NFD* treatment.

As feedback between an evolving population and its environment influences the degree of adaptation, we next examined whether the exact nature of the feedback was important. We created another treatment where instead of depleting resources in the environment (as in *NFD*) organisms *increased* the abundance of a resource when its associated task was performed, an example of positive frequency dependence. This *Positive Frequency Dependence (PFD)* treatment is also characterized by feedback between the environment and the population during the middle third of the run. Despite the presence of feedback, populations in the *PFD* treatment had significantly reduced mean final fitness relative to *NFD* and a trend toward reduced mean fitness relative to *Fixed* treatments (Mann-Whitney, $p<0.001$ and $p=0.08$ respectively), implying that the negative environmental feedback of *NFD* is necessary for enhanced adaptation.

One reason *NFD* may facilitate adaptation is that the population can become more diverse during the middle third of the run, i.e., it occupies a larger number of positions in the adaptive landscape. This increased distribution could yield an improved outcome when the reference is revisited due to greater accessibility of adaptive peaks. As expected, during the middle third of runs, *NFD* treatments experienced an increase in genotype diversity (measured as Shannon Entropy across genotypes), while *PFD* runs decreased in diversity relative to the *Fixed* treatment (Fig 4j, Mann-Whitney, $p<0.001$ and $p<0.001$ respectively). With negative frequency dependence, prevalent genotypes reduce their own fitness, flattening the landscape and allowing for otherwise less fit genotypes to coexist.

Discussion

We observed that populations in an exogenously changing environment evolved to a higher fitness relative to populations in unchanging environments. We surmise that populations became constrained in genotype space during the first third of their evolutionary trajectory, as only a few muta-

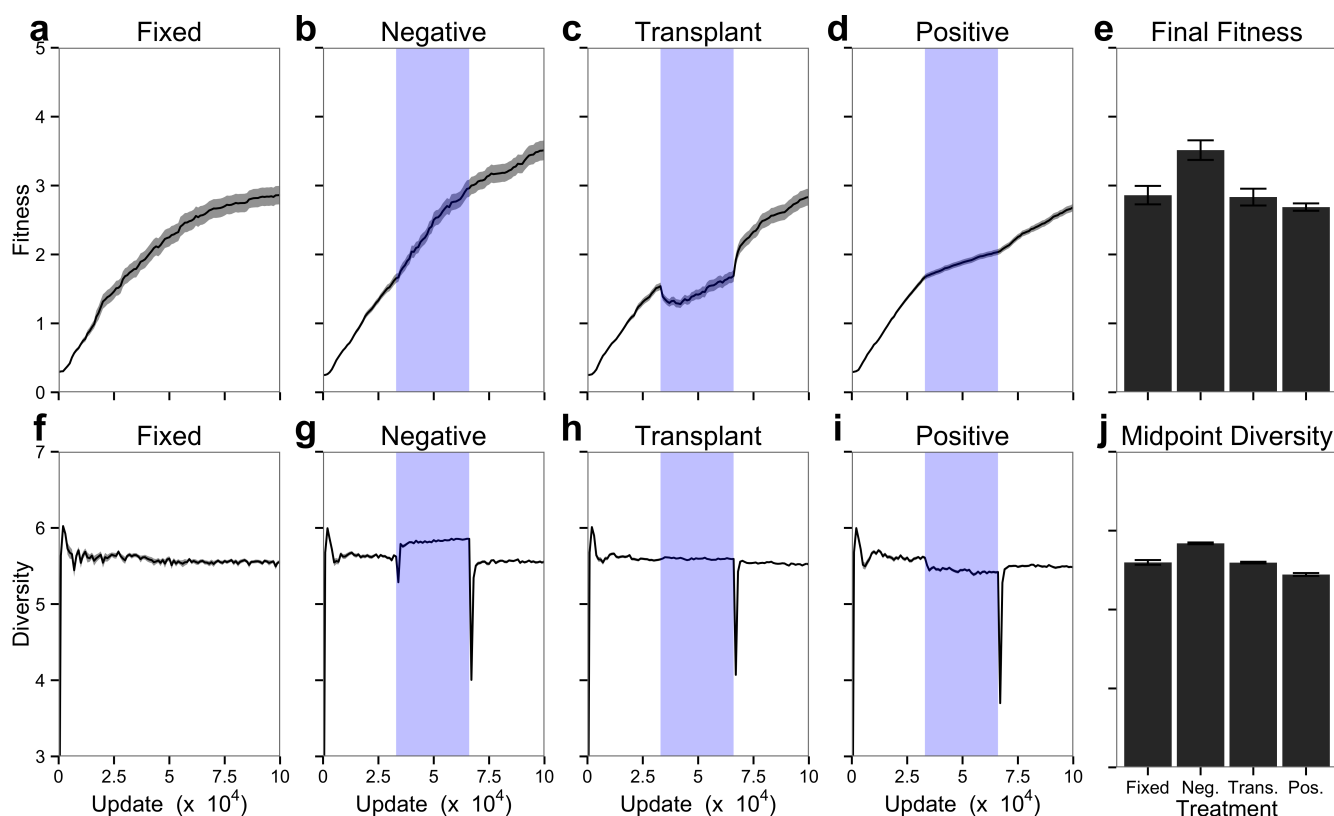


Figure 4: Fitness and Diversity in Endogenously Changed Environments Above, mean fitness of the line of descent from of 60 replicates of runs in the *Fixed* (panel a), *Negative Frequency Dependence* (b), *Paired Transplant* (c), and *Positive Frequency Dependence* (d) treatments, and their final mean fitness as a bar plot (e). Below, the mean genotypic diversity (measured by Shannon Entropy across genotypes) of the populations over time for the same treatments, and the mean diversity value at update 50,000 (the midpoint of the run) is shown. Blue shading denotes exposure to a different environment(s) and ribbons and error whiskers denote standard error of the mean.

tions in genotypes at this time conferred benefits (Fig. 3). Changing the environment for the middle third of the run serves to liberate the population trapped on a suboptimal peak. This result raises the question of why populations generally evolved to better positions in the reference landscape, as opposed to different positions. There is no *a priori* reason why such populations would not have moved into a section of the reference landscape with worse evolutionary prospects.

We postulate two reasons why populations spending time in alternative environments reach higher fitness genotypes after returning to their reference environment. If the fitness of a peak (its “height”) is proportional to the size of its basin of attraction (the number of genotypes that the peak can be reached by solely beneficial mutations), the additional movement in genotype space afforded by evolving in a different environment, should be beneficial. For example, if a population initially evolved to a low peak, movement in the alternative environment would be likely to shift the population to a different/better basin of attraction of a higher peak.

If, instead, the population initially evolved to a high peak, taking steps in the other environment will be less likely to cause the population to leave the larger basin, and thus it would likely return to the same high peak.

A second mechanism leading to better adaptation relies on association between fitnesses of the reference and alternative environments. If regions of higher fitness are shared between the environments (certain genotypes/traits may be favored in both), traversing the alternative environment may favor traits that are also beneficial in the reference environment. Prior work by Lenski et al. (Lenski et al., 2003), finds that many of the more complex traits favored in our reference environment (requiring a specific arrangement of instructions in the genome) could be co-opted by a few mutations to be beneficial in the alternative environment, and vice versa. However, the benefit of evolutionary exposure to a rigid, different environment will depend on the nature of the alternative landscape and the position of the population in genotype space. For example, Kashtan et al. (Kashtan et al., 2007) found that switching between environments that

shared common sub-problems yielded better solutions than unrelated environments. This dynamic is likely at play in the Avida system as well.

In light of these findings, we return to Wrights Shifting Balance Theory (SBT). Wrights SBT relies on drift and selection; however, these features are antagonistic, as circumstances that support drift hinder selection (e.g., small population sizes). Both Fisher and Wright understood that a change in the environment could move a population off a former peak. This dynamic requires environmental sign epistasis, where the fitness effect of a mutation is beneficial in one environment, but detrimental in another (Lindsey et al., 2013). Here we could imagine a recasting of SBT for a peak shift in a reference environment. A population is structured into demes, which need not be small. Suppose the whole population starts on some sub-optimal peak. If demes experience heterogeneous environments (temporally or spatially), they may be able to take different paths to different peaks. Migration between demes allows the higher fitness genotypes to spread and fix globally. With such a model, dynamic environments may allow rapid evolution across rugged landscapes without the requirement for small subpopulation size.

In contrast to the previous descriptions of exogenously changed environments, malleable environments are shaped by organisms and, in turn, can affect their evolution. In our *Negative Frequency Dependent* treatment, as phenotypes exploiting novel resources increased in frequency, they became devalued. This reduction in fitness "flattens" the landscape, as the rich (higher fitness genotypes) get poorer, which diversifies the population into different niches. This effect promotes the exploration of other regions of the genotype space, and indeed, in our *Negative Frequency Dependent* treatment, diversity was enhanced during the middle third. Whereas, in *Paired Transplant* runs, superior phenotypes found in this population remained superior and, in this situation, some rich stay rich leading to reduced diversity. In the *Paired Transplant* of each *NFD* run, diversity does not increase, as novel phenotypes, which did not occur in the *NFD* run, can easily spread and fix. In *Positive Frequency Dependent* runs, an initial solution feeds back (enriches its environment) to make itself even more superior. In this case the rich get richer and displace other potential phenotypes. Thus some malleable environments, such as *NFD*, can yield enhanced adaptation by favoring diversity and innovation.

Incorporating dynamic environments, especially malleable ones, to our models of evolutionary change may enhance our understanding of adaptation, as has also been seen with host-pathogen coevolution (Zaman et al., 2014). Other natural systems that demonstrate negative frequency dependence such as resource consumption (Ross Gillespie et al., 2007; Svanback and Bolnick, 2007) or host-pathogen coevolution (Carius et al., 2001; Koskella and Lively, 2009; Thrall et al., 2012) have stably diverse populations that may

be better able to adapt to their environments. By incorporating common features of the natural world, namely, the dynamic aspect of environments and the feedback between populations and their environments, we found that populations may be able to adapt faster. This result has implications for understanding evolution in nature, but may also suggest useful features to incorporate into evolutionary algorithms to solve engineering problems (Back and Schwefel, 1993; Sauer, 2001). For instance, if the fitness function of an evolutionary algorithm discounted current high fitness solutions, alternative solutions can be more thoroughly explored. This principle is central to novelty search, one popular evolutionary algorithm (Inden et al., 2012). Thus natural and artificial populations may yield adaptive benefits from exposure to exogenously or endogenously altered environments (Goings and Ofria, 2009).

References

- Alto, B. W., Wasik, B. R., Morales, N. M., and Turner, P. E. (2013). Stochastic Temperatures Impede Rna Virus Adaptation. *Evolution*, 67(4):969–979.
- Ancel Meyers, L., Ancel, F. D., and Lachmann, M. (2005). Evolution of Genetic Potential. *PLoS Computational Biology*, 1(3).
- Back, T. and Schwefel, H.-P. (1993). An Overview of Evolutionary Algorithms for Parameter Optimization. *Evol. Comput.*, 1(1):1–23.
- Carius, H. J., Little, T. J., and Ebert, D. (2001). Genetic Variation in a Host-Parasite Association: Potential for Coevolution and Frequency-Dependent Selection. *Evolution*, 55(6):1136–1145.
- Collins, S., de Meaux, J., and Acquisti, C. (2007). Adaptive walks toward a moving optimum. *Genetics*, 176(2):1089–1099.
- Fisher, R. A. (1930). *The Genetical Theory Of Natural Selection*. At The Clarendon Press.
- Goings, S. and Ofria, C. (2009). Ecological approaches to diversity maintenance in evolutionary algorithms. In *2009 IEEE Symposium on Artificial Life*, pages 124–130.
- Heino, M., Metz, J. A., and Kaitala, V. (1998). The enigma of frequency-dependent selection. *Trends in Ecology & Evolution*, 13(9):367–370.
- Inden, B., Jin, Y., Haschke, R., Ritter, H., and Sendhoff, B. (2012). An examination of different fitness and novelty based selection methods for the evolution of neural networks. *Soft Computing*, 17(5):753–767.
- Jones, C. G., Lawton, J. H., and Shachak, M. (1994). Organisms as Ecosystem Engineers. *Oikos*, 69(3):373–386.

- Kashtan, N., Noor, E., and Alon, U. (2007). Varying environments can speed up evolution. *Proceedings of the National Academy of Sciences*, 104(34):13711–13716.
- Kauffman, S. and Levin, S. (1987). Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45.
- Koskella, B. and Lively, C. M. (2009). Evidence for negative frequency-dependent selection during experimental coevolution of a freshwater snail and a sterilizing trematode. *Evolution; International Journal of Organic Evolution*, 63(9):2213–2221.
- Laland, K. N., Odling-Smee, F. J., and Feldman, M. W. (1996). The evolutionary consequences of niche construction: a theoretical investigation using two-locus theory. *Journal of Evolutionary Biology*, 9(3):293–316.
- Laland, K. N., Odling-Smee, F. J., and Feldman, M. W. (1999). Evolutionary consequences of niche construction and their implications for ecology. *Proceedings of the National Academy of Sciences*, 96(18):10242–10247.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.
- Lindsey, H. A., Gallie, J., Taylor, S., and Kerr, B. (2013). Evolutionary rescue from extinction is contingent on a lower rate of environmental change. *Nature*, 494(7438):463–467.
- Odling-Smee, F. J., Laland, K. N., and Feldman, M. W. (2003). *Niche Construction: The Neglected Process in Evolution*. Princeton University Press, Princeton.
- Ofria, C. and Wilke, C. O. (2004). Avida: a software platform for research in computational evolutionary biology. *Artificial Life*, 10(2):191–229.
- Parter, M., Kashtan, N., and Alon, U. (2007). Environmental variability and modularity of bacterial metabolic networks. *BMC Evolutionary Biology*, 7:169.
- Ross Gillespie, A., Gardner, A., West, S., Griffin, A., Sherratt, A. E. T. N., and Whitlock, E. M. C. (2007). Frequency Dependence and Cooperation: Theory and a Test with Bacteria. *The American Naturalist*, 170(3):331–342.
- Sauer, U. (2001). Evolutionary engineering of industrially important microbial phenotypes. *Advances in Biochemical Engineering/Biotechnology*, 73:129–169.
- Sterelny, K. (2005). Made By Each Other: Organisms and Their Environment. *Biology and Philosophy*, 20(1):21–36.
- Svanback, R. and Bolnick, D. I. (2007). Intraspecific competition drives increased resource use diversity within a natural population. *Proceedings of the Royal Society of London B: Biological Sciences*, 274(1611):839–844.
- Thrall, P. H., Laine, A.-L., Ravensdale, M., Nemri, A., Dodds, P. N., Barrett, L. G., and Burdon, J. J. (2012). Rapid genetic change underpins antagonistic coevolution in a natural host-pathogen metapopulation. *Ecology Letters*, 15(5):425–435.
- Waxman, D. and Peck, J. R. (1999). Sex and adaptation in a changing environment. *Genetics*, 153(2):1041–1053.
- Whitlock, M. C. (1997). Founder Effects and Peak Shifts Without Genetic Drift: Adaptive Peak Shifts Occur Easily When Environments Fluctuate Slightly. *Evolution*, 51(4):1044–1048.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proc of the 6th International Congress of Genetics*, volume 1, pages 356–366.
- Zaman, L., Meyer, J. R., Devangam, S., Bryson, D. M., Lenski, R. E., and Ofria, C. (2014). Coevolution Drives the Emergence of Complex Traits and Promotes Evolvability. *PLOS Biol*, 12(12):e1002023.

Acknowledgment

This material is based in part on work supported by the National Science Foundation under Cooperative Agreement Number DBI-0939454 (to C.O.), a National Science Foundation Graduate Research Fellowship (to J.R.N.), and a NSF CAREER Award Grant (DEB0952825) (to B.K.).