DOI: http://dx.doi.org/10.7551/978-0-262-31050-5-ch004

The role of deleterious mutations in the adaptation to a novel environment

Arthur W. Covert III[1,2]*, Jared Carlson-Stevermer[1,2], Dakota Z. Derrberry[1,2], and Claus O. Wilke[1,2]

[1] Section of Integrative Biology, Center for Computational Biology and Bioinformatics, and Institute for Cellular and Molecular Biology, The University of Texas at Austin.
[2] BEACON, Center for the Study of Evolution in Action,
*Author to whom correspondence should be sent
Art.Covert@austin.utexas.edu

Abstract

Organisms adapt by accumulating beneficial mutations. Yet sometimes these beneficial mutations are not directly accessible, and organisms may have to cross a fitness valley before further adaptation is possible. A few recent works have shown that crossing of fitness valleys, as evidenced by fixation of deleterious mutations, may be surprisingly common in adaptation, and may be an important contributor to long-term fitness increase. Here we ask how important crossing of fitness valleys is for organisms that have reached a local fitness peak in one environment and are then placed into a new environment. We compare two treatments of evolving digital organisms, one in which organisms are exposed to deleterious mutations and thus can freely explore fitness valleys, and one in which they are prevented from experiencing deleterious (but not lethal) mutations and thus cannot. We find that organisms that are exposed to deleterious mutations always do at least as well as organisms that are not. Whether organisms exposed to deleterious mutations do better depends on the relative similarity and complexity of the old and new environment. We conclude that crossing of fitness valleys is important for successful adaptation to certain types of novel environments.

Introduction

Conventional wisdom holds that beneficial mutations are good and deleterious mutations are bad. Yet in a finite population, deleterious mutations can contribute to the long-term evolutionary success by allowing the population to traverse a fitness valley leading to a higher fitness peak (Weissman et al. 2010, van Nimwegen and Crutchfeild 2000). Therefore, deleterious mutations cannot be unconditionally bad. There must be well-defined scenarios under which a hypothetical population experiencing no deleterious mutations would fare worse than a population that does experience them.

Indeed, several recent works have highlighted that crossing fitness valleys may be important for long-term evolutionary success. Simulations of RNA folding have found transitions between fitness peaks in the form of "fitness reversal" of deleterious mutations (Cowperthwaite et al 2007). Experiments with self-replicating computer programs have shown that such fitness reversals may actually open up new areas of the fitness landscape, areas which were inaccessible except via a deleterious mutation (Covert 2010, Lenski et al 2003). Experiments in *Saccharomyces cerevisiae* have uncovered at least one instance of a fitness reversal in an organic system (Kvitek and Sherlock 2011). This finding suggests that the fitness landscape of yeast is rugged, and that it requires at least the occasional valley crossing via a fitness reversal.

However, some authors have conjectured that a change to a novel environment can eliminate the need for crossing fitness valleys (Whitlock 1997, Whitlock et al 1995). A sufficiently large environmental change may turn a high-fitness peak into low-lying region in the fitness landscape, from which there are many new ways to climb up. If an environmental change creates a large number of new adaptive opportunities, a population should be able to adapt without having to cross any fitness valleys. Fixation of beneficial mutations alone should drive the population towards new fitness peaks. On the other hand, if the environmental change results in a rugged fitness landscape, then high-fitness regions may still only be accessible by traversal of fitness valleys, and deleterious mutations may be required for successful adaptation.

Here, we test the importance of crossing fitness valleys in the adaptation of digital organism. Digital organisms are self-replicating computer programs that evolve to perform various logical functions (corresponding to phenotypic traits). We can manipulate environmental complexity by changing how many and which logical functions are rewarded (Lenski et al 1999). We also can monitor all mutations as they appear in the population, and prevent mutations with certain characteristics (such as deleterious mutations) from ever entering the population. This setup allows us to directly compare the evolution of populations experiencing and not experiencing deleterious mutations (Covert 2010). We find that deleterious mutations are most important for long-term evolutionary success if the new environment rewards a small number of new traits that are complex and difficult to evolve. By contrast, if the new environment rewards either a large number of new traits or a small number of new traits that are less complex, deleterious mutations provide less benefit.

Methods

Experimental system. We used the digital life system Avida, version 2.12.2, for all experiments (Ofria and Wilke 2004). In Avida, digital organisms evolve and adapt to perform various one and two input logical functions (Table 1). Populations are

normally seeded with organisms that can do nothing but self-replicate. The code that makes up a digital organism is composed of simple CPU instructions, which are colloquially referred to as the genome. Mutation acts on these genomes by changing one instruction to a new randomly chosen instruction. Over time, these organisms evolve to do logical tasks that reward them with more CPU cycles they can use to execute their genomes more rapidly. Thus, they metabolize inputs from the environment to perform logical functions that give them additional energy to self-replicate more quickly.

All experiments were done with a mutation rate of 25% on divide and all populations were allowed to grow to a maximum size of 10,000 organisms.

Adaptations in one- and two-trait environments. Seed organisms were generated by adapting populations to one logic function, but not others. We called these initial replicates "priming populations" and began them with 50 replicates of populations with a standard seed organism that could do nothing but self-replicate. The replicate populations were first evolved for 100,000 updates in an environment that rewarded all one and two input logical functions except EQU. Starting at update 100,000, the NOT function was turned into a deleterious trait (i.e., organisms which performed NOT received a 75% fitness reduction). Making a trait deleterious creates selective pressure for the digital organisms to evolve away from that trait and towards other traits. Every next 20,000 updates, another logical operation was turned into a deleterious trait, until only XOR remained. The populations were then evolved for another 100,000 updates in an environment in which only XOR was beneficial and the other 7 functions (7 lower order functions, except EQU) were deleterious. We identified all priming populations in which the final dominant genotype was exclusively performing XOR at the end of the experiment, and randomly selected 25 of the final dominant genotypes to seed populations which would adapt to novel environments.

We used the 25 priming-population final dominant genotypes to seed 2 experiments, each with two paired treatments. Each treatment had 8 populations founded from each priming genotype, for a total of 200 replicate populations per treatment. In the first experiment, populations were adapted to environments with one of two new functions, NOR or EQU, but no other functions. In the second experiment populations were adapted to environments that rewarded XOR and one additional task, NOR or EQU. The new functions were respectively equally complex and more complex than XOR. Environments which rewarded one function where considered to generate a single-peaked fitness landscape and those rewarding multiple functions where considered to generate a multi-peaked landscape. None of the environments punished organisms for performing other logical functions.

All replicates were subjected to two treatments, *Control* and *Replace Deleterious* (RpD). RpD monitored every mutation that arose in the evolving organisms and replaced every deleterious mutation that occurred with a new, randomly chosen neutral, beneficial, or lethal mutation. The RpD protocol is identical to the one used in Covert (2010).

The Control treatment and the RpD treatment differ in that organisms can enter fitness valleys under the Control treatment but not under the RpD treatment. To enter a fitness valley, an organism has to suffer a deleterious mutation. These mutations are eliminated under RpD. To maintain a comparable overall mutation rate under this treatment, deleterious mutations are replaced by either a neutral or beneficial mutation, which cannot lead into a fitness valley by definition, or by a lethal mutation, which simply kills the offspring organism and prevents any further adaptation along this lineage.

Adaptations in multi-trait environments. We started 50 replicates with a standard seed organism, as before. The priming environment rewarded four traits (NOT, ORN, OR, and NOR). Populations evolved in this environment for 50,000 updates. All evolved populations were then transferred into a novel environment that rewarded all 9 possible one- and two-input logic functions (Table 1). Populations evolved a further 200,000 updates under the novel environment, exposed to two separate treatments: Control and RpD, as before.

Statistical analyses. We carried out all statistical analyses with SciPy (version 0.9). Fitnesses of evolved populations were measured on the dominant (most abundant) genotype in the final population. Fitness comparisons were performed using paired *t*-tests on log-transformed fitness values. When multiple replicates were derived from identical priming populations, we averaged log-transformed fitness values of those replicates before performing paired *t*-tests, to avoid pseudo-replication.

Function Name	Logic Operation	Energy Bonus
NOT	~A; ~B	x2
NAND	~(A AND B)	x2
AND	A AND B	x4
OR_N	(A OR ~B) (~A OR B)	x4
OR	A OR B	x8
AND_N	(A AND ~B) (~A AND B)	x8
NOR	~A AND ~B	x16
XOR	(A AND ~B) OR (~A AND B)	x16
EQU	(A AND B) OR (~A AND ~B)	x32

Table 1: The standard nine logical functions in the Avida environment and their energy bonus. Digital organisms have only the NAND operation available to them and must construct other logical functions out of NAND operations. The energy bonus for each for each function is equivalent to 2^n, where n is the minimum number of NAND operations needed to complete it. Each logical function corresponds to a phenotypic trait in the environment.

Results

Adaptation to environments that reward one or two traits. How important are deleterious mutations for adaptation to a new environment? To address this question, we carried out experiments with the following general design: We first let populations of digital organisms adapt to a chosen environment. (We call this initial environment the *priming environment* and populations evolved in this environment the *priming populations.*) From the priming populations, we selected a subset that had acquired the optimal phenotype in the priming environment. We then further adapted these populations to novel environments, with two separate treatments. The *Control* treatment was standard adaptation as used for the priming populations. The *Replace Deleterious* (RpD) treatment prevented organisms from experiencing deleterious mutations and hence from exploring fitness valleys. After adaptation in the novel environment, we compared which treatment, if any, led to higher fitness values in the final dominant (most abundant) organisms in the evolved populations.

We first studied the case where the priming environment rewarded one phenotypic trait and the novel environment rewarded either one other or one additional phenotypic trait. For digital organisms in the Avida world, phenotypic traits are defined via two-input logical functions. In our first set of priming adaptations, we evolved organisms to carry out only the XOR function. (We refer to this environment as the *XOR* environment.) We subjected organisms primed in *XOR* to three novel environments, which rewarded either the NOR function (*NOR* environment), both the NOR and the XOR functions (*NOR/XOR* environment), or both the XOR and EQU functions (*EQU/XOR* environment).

We found that the Control and RpD treatments performed similarly when switching from the *XOR* environment to either the *NOR* or the *NOR/XOR* environment. In *NOR*, the final dominant fitness was not significantly different among the two treatments (mean pairwise difference of log fitness $\bar{d}$=0.83, p=0.41, paired t-test), and neither was the number of times each organism evolved NOR (198 for the control and 196 for RpD). Likewise, in *NOR/XOR*, the final dominant fitness was not significantly different among the two treatments ($\bar{d}$=0.98, p=0.34, paired t-test). The number of times that XOR and NOR evolved in each treatment also did not differ signiricantly (NOR 181/185, XOR 181/180 between C/RpD). However, note that in all cases the final fitness under the Control treatment was higher than the final fitness under the RpD treatment, as evidenced by $\bar{d}>0$.

We consider XOR and NOR to be equally difficult to evolve or perform because they require the same number of NAND operations to complete. Thus, the switch from the priming *XOR* environment to either the *NOR* or the *NOR/XOR* environment was a switch to a novel environment of comparable complexity. To assess whether environment complexity played a role in experimental outcome, we also adapted Control and RpD treatments to the *EQU/XOR* environment. The EQU function is more complex than XOR or NOR because even the most parsimonious solution to EQU requires at least one more NAND operation than either XOR or NOR. In *EQU*, we found that the Control and RpD groups showed significant differences in both final dominant fitness and the number of replicates that evolved the EQU function (Table 2). The final dominant fitness of the Control group (with deleterious mutations) significantly exceeded that of the RpD group (without deleterious mutations) ($\bar{d}$=2.70, p=0.012, paired t-test). The Control group also evolved the EQU function and retained the XOR function significantly more often than the RpD group did (15 more evolved EQU, 20 more evolved both, p=0.015 odds ratio 2.57, and p=0.029 odds ratio 1.7 Fisher's exact test, respectively) .

In summary, these results show that deleterious mutations may be of benefit in long-term adaptation when organisms adapt to a novel environment of increased complexity (*EQU/XOR*) but not necessarily when they adapt to a novel environment of comparable complexity (*NOR*, *NOR/XOR*).

Treatments	XOR	EQU	Both
C	178	189*	159*
RpD	165	170	139

Table 2: Number of replicates which evolved to the tasks present in the *XOR/EQU* environment. The control evolved EQU and both significantly more than the RpD treatment did (p=0.015 and p=0.029 respectively, Fishers exact test)

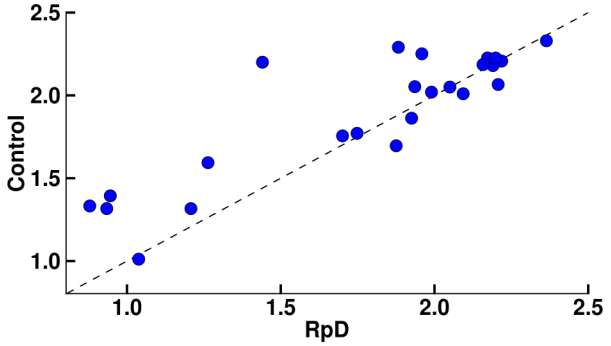

Figure 1: Average $\log_{10}$ final dominant fitness of all replicates initialized with one of 25 seed organism from a priming population, in the XOR/EQU environment. The control treatment tended to have significantly higher fitness than the RpD treatment ($\bar{d}$=2.65, p=0.014, paired t-test).

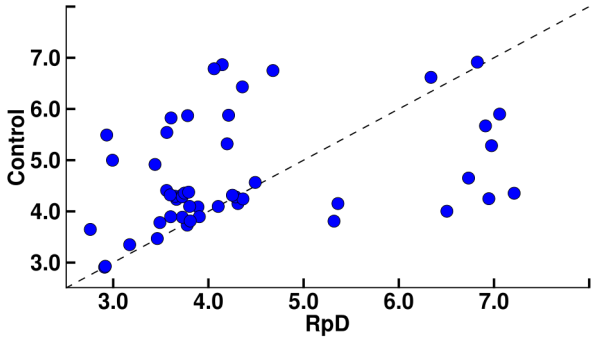

Figure 2: Average $\log_{10}$ final dominant fitness of all 50 replicates that underwent an environmental change from 4 tasks to 9 tasks after 50,000 updates of evolution. No significant differences where detected,($\bar{d}$=1.62, p=0.11, paired t-test).

Adaptation to environments that reward multiple traits. The previous experiments added only a single novel trait. One could hypothesize that if the novel environment rewards many new traits, then beneficial mutations should be plentiful and deleterious mutations should not be required for successful long-term adaptation. On the flip side, a complex novel environment may also provide many fitness barriers which would require deleterious mutations for efficient traversal. To test which of these two scenarios applied in our system, we carried out additional experiments in environments rewarding multiple traits. As priming environment, we used an environment rewarding four traits. As novel environment, we used an environment rewarding all 9 possible one- and two-input logic functions available in Avida. Unlike the previous experiment, we did not select a subset of priming populations that had evolved specific phenotypic traits. We simply evolved 50 replicates for 50,000 updates in the priming environment. We then took all evolved populations and subjected them for another 200,000 updates to the novel environment, under both the Control and RpD treatments as before.

We found that adaptation to the new environment lead to significantly higher fitness values under the control treatment ($\bar{d}$=2.52, p=0.0124, paired t-test). Deleterious mutations seemed to play a key role in the evolution of these populations, despite the introduction of a large number of new beneficial mutations. These new beneficial mutations where presumably on rugged parts of the fitness landscape and inaccessible from the fitness peaks in the earlier environment. The complex logical functions AND and EQU evolved more often under the Control treatment, but not significantly so (see Table 3 for full analysis). These findings suggest that the role of deleterious mutations in changing environments depends not just on the influx of new beneficial adaptations but also on the complexity of the environmental change. Adaptation to more complex environments may require more deleterious mutations.

	NOT	**NAND**	**AND**	**OR_N**	**OR**
C	196	187	136	198	197
RpD	194	181	113	200	196

	AND_N	**NOR**	**XOR**	**EQU**
C	185	197	51	72
RpD	180	185	46	54

Table 3: Number of replicates that evolved the logical tasks in the multiple traits environment. No significant differences were detected between the control and RpD treatments, but differences were noticeable between the AND and EQU tasks (p=0.023 and p=0.067 respectively, fishers exact test). However, neither p-value is significant when we consider that the analysis consists of 9 repeated tests, significance would have required a p-value of less than 0.0055.

Discussion

We have shown that deleterious mutations can play a significant role in adaptive evolution, depending on both the initial environment in which a population has evolved and the novel environment to which the organism becomes exposed. We used Avida to evolve digital organisms with and without deleterious mutations in a variety of environments. We found that in a switch from a one-trait environment to a one- or two-trait environment with similar complexity, deleterious mutations did not provide a significant long-term adaptive advantage. However, when we switched to a two-trait environment of greater complexity, we found a significant effect. Likewise, when we switched from a 4 trait environment to a nine trait environment we found a significant role for deleterious mutations. Thus, the effect of deleterious mutations does not seem to be universal; what works in one environment may not work in another. The question then becomes which specific elements of the environment make deleterious mutations advantageous or not. Whitlock (1997) suggested that the type of change may be less important than the frequency of change. Here, we did not assess the impact of frequency of change, since all population experienced exactly one switch from priming environment to novel environment. Thus, our experiments do not speak directly to Whitlock's conjecture. However, our experiments clearly show that frequency of change is not the only relevant variable; changes in environmental complexity and in similarity between the priming and novel environments are sufficient to significantly alter the importance of deleterious mutations to long-term adaptation.

When a population is exposed to a new environments, two parameters should determine how important the traversing of fitness valleys is for successful long-term adaptation: the ruggedness of the fitness landscape in the new environment and the number of novel adaptive opportunities (i.e., number of paths that lead uphill in a fitness landscape). When we switched the environment from *XOR* to *NOR, XOR* to EQU or to *XOR/NOR*, we likely did not add a large number of novel adaptive opportunities, but we also did not increase the ruggedness of the landscape by much. Thus, traversals of fitness valleys were not particularly important. Switching from *XOR* to *EQU/XOR* similarly did not provide a large number of novel adaptive opportunities, but it likely did increase the ruggedness of the landscape, owing to EQU's increased computational complexity relative to XOR. Thus, traversals of fitness valleys were important in this scenario. Finally, by switching from the four-trait environment to the nine-trait environment, we introduced additional ruggedness, but we certainly also added a large number of additional adaptive opportunities. Many of the new adaptive opportunities seem to require a deleterious mutation in order to exploit them, as evidenced from the higher fitness values in the control treatment.

Fitness interactions between genes can sometimes create fitness effects that deviate from our expectations. This effect is called epistasis, and can occur when two or more mutations have a greater or smaller sum fitness together than they do individually. Sign-epistasis is the most extreme form of this, in which individually deleterious mutations mutations may provide a net benefit (or vic-versa). Recent theoretical works

have suggested that while escape from fitness peaks via sign-epistatic mutations is possible (Weinreich and Chao 2005, Weissman et al. 2010), the main role of sign-epistasis is to constrain adaptive paths. Our work suggests that sign-epistasis may be a driving force in evolution.

One limitation of our work is that we did not identify the key adaptations that caused the net fitness benefit. Previous works with deleterious mutations in digital organisms identified and isolate key sign-epistatic interactions on the line of descent that opened up new areas of the adaptive landscape to explore (Covert 2010). A second limitation is that the seed organisms for all experiments were evolved in populations that allowed deleterious mutations. It is possible that the initial exposure to deleterious mutations made it easier for adaptation in later experiments to proceed without additional deleterious mutations.

Evolution is a combination of three factors: variation, inheritance, and selection. Most works on evolution focus on just one of these factors: selection. Variation and inheritance may be thought of as synonyms for chance and history, without which evolution cannot proceed (Blount et al 2010). There is no better example of a chance evolutionary event becoming important than a sign-epistatic fitness reversal, a previously deleterious mutation which becomes critically important to future evolution. With modern technology it is now possible to observe and measure the importance of history and chance in simple evolving systems. While selection is the ultimate arbiter of which variation and which history will be successful, selection requires raw material to work with, and this raw material may be the result of highly unlikely events.

Acknowledgments

This material is based in part upon work supported by the National Science Foundation under Cooperative Agreement No. DBI-0939454 and by the Freshman Research Initiative at the Univeristy of Texas at Austin. We thank Charles Ofria and Richard Lenski for stimulating discussions.

References

Blount, ZD., Borland CZ., and Lenski, RE (2008), Historical Contingency and the Evolution of a Key Innovation in an Experimental Population of Escherichia coli. *PNAS* 105:7899-7906.

Covert III, A. W. (2010), *On the Beneficial Effects of Deleterious Mutations.* Ph.D thesis, Department of Computer Science Michigan State University.

Cowperthwaite MC, Bull JJ, Meyers LA (2006), From bad to good: Fitness reversals and the ascent of deleterious mutations. *PLoS Comp Biol* 2:1292-1300.

Kvitek DJ, Sherlock G (2011), Reciprocal Sign Epistasis between Frequently Experimentally Evolved Adaptive Mutations Causes a Rugged Fitness Landscape. *PLoS Genet* 7:e1002056.

Lenski RE, Ofria C, Collier TC, Adami C (1999), Genome complexity, robustness and genetic interactions in digital organisms. *Nature* 400:661-664.

Lenski RE, Ofria C, Pennock RT, Adami C (2003), The evolutionary origin of complex features. *Nature* 423:139-144.

Ofria C, Wilke CO (2004), Avida: a software platform for research in computational evolutionary biology. *J Artif Life* 10:191-229.

van Nimwegen E, Crutchfield JP (2000), Metastable evolutionary dynamics: Crossing fitness barriers or escaping via neutral paths? *Bull Math Biol* 62:799-848.

Weinreich DM, Chao L (2005), Rapid evolutionary escape by large populations from local fitness peaks is likely in nature. *Evolution* 59:1175-1182.

Weissman DB, Desai MM, Fisher DS, Feldman MW (2009) *J Theor Biol* 75:286-300.

Whitlock, MC. (1997), Founder Effects and Peak Shifts Without Genetic Drift: Adaptive Peak Shifts Occur Easily When Environments Fluctuate Slightly. *Evolution.* 51:1044-1048

Whitlock, MC., Phillips PC., Moore, FBG., and Tonsor, SJ. (1995), Multiple Fitness Peaks and Epistasis, *Annu. Rev. Ecol. Syst*. 26:601-629