

Evolution of Robustness in Digital Organisms

Jeffrey A. Edlund*

California Institute of
Technology
Pasadena, CA 91125
jedlund@ugcs.caltech.edu

Christoph Adami†

Digital Life Laboratory 136-93
California Institute of
Technology
Pasadena, CA 91125
adami@caltech.edu

Abstract We study the evolution of robustness in digital organisms adapting to a high mutation rate. As genomes adjust to the harsh mutational environment, the mean effect of single mutations decreases, up until the point where a sizable fraction (up to 30% in many cases) of the mutations are neutral. We correlate the changes in robustness along the line of descent to changes in directional epistasis, and find that increased robustness is achieved by moving from antagonistic epistasis between mutations towards codes where mutations are, on average, independent. We interpret this recoding as a breakup of linkage between vital sections of the genome, up to the point where instructions are maximally independent of each other. While such a recoding often requires sacrificing some replication speed, it is the best strategy for withstanding high rates of mutation.

1 Introduction

The standard picture of evolutionary adaptation as a steady upward march towards the highest peaks in the fitness landscape has received important additions recently. These amendments concern evolution at high mutation rates rather than the low rates of standard population genetics, and take into account the *quasi-species* nature of populations evolving at high mutation rates (see [22] for a review). A quasi-species is an assembly of closely related self-replicating molecules (the term “species” harks back to its use in chemistry, not in biology) that is highly structured, and that is *as a whole* the target of selection. This deviation from the standard picture in which single molecules (or organisms) are the unit of selection is made possible by the high rate of mutation that turns members of the *cloud* of sequences into each other at a high rate. Many of these mutations do not affect fitness (they are *neutral*), and therefore change only the genetic, not the functional, identity of the sequence. One consequence of evolution at high mutation rate is that it is possible that sequences with a strongly impaired fitness can appear on the line of descent if they are part of a quasi-species of adapting molecules [9]. Mathematically, this can be explained by noting that the probability of fixation of deleterious mutations is finite within quasi-species theory even for very large populations [20] (strongly deleterious mutations can go to fixation in standard population genetics only in very small populations). Also, while such a mutation may be deleterious to the organism carrying it, it is in fact advantageous to the quasi-species as a whole.

Another recent observation is that a novel selective pressure starts to emerge at high mutation rates, which favors mutationally robust genomes over fragile ones, even (if the

* Present address: Jet Propulsion Laboratory 169-506, California Institute of Technology, Pasadena, CA 91109.

† Present address: Keck Graduate Institute for Applied Life Sciences, 535 Watson Drive, Claremont, CA 91711.

mutation rate is high enough) if the robust genomes replicate less fast. This unexpected observation, now commonly known as the “survival of the flattest” phenomenon, was first postulated theoretically [10, 18, 11], and then observed experimentally within populations of digital organisms [17].

The phenomenon of robustness selection [2] can be understood in two equivalent but complementary ways. On the one hand, it can be described in terms of a pressure for sequences evolving on neutral networks to occupy the highly connected, rather than sparse, areas of the network. Or, in a picture emphasizing a fitness landscape of peaks and valleys, it can be viewed as a pressure to occupy broad rather than narrow fitness peaks, because organisms occupying high but narrow peaks are easily mutated to nonfunctional (zero-fitness) sequences. While this dynamics has been extensively documented [19], with clear evidence of populations of sequences *leaving* high but narrow peaks to occupy nearby wide but shallower peaks when exposed to a high mutation rate, it is unclear exactly how it is possible that such nearby peaks are even available. In other words, while it is clear that, faced with a high mutation rate, organisms shed replication speed for robustness, it is not clear how this robustness is achieved in the first place. Is robustness achieved just by increasing sequence length (i.e., adding so-called “junk” DNA)? Or, if sequence length is held constant, is it simply because a previously functional region turned neutral? Do organisms develop a type of redundancy that protects them from harmful mutations? What are the *rules of recoding* the existing information so that the average effect of a mutation becomes less lethal? In this contribution, we consider the hypothesis that the average deleterious effect of a single mutation is lowered simply by decreasing the amount of antagonistic epistasis between mutations.

2 Epistasis and Neutrality

The term “epistasis” is generally used in genetics to indicate an interdependence between genes, or more precisely, to indicate that the functionality of one gene depends on the functionality of one or more others [24]. Gene epistasis has been shown to play a major role in genetics, and seems to affect the genetic load [5] and the topology of fitness landscapes in general [16], as well as the evolution and maintenance of recombination [7, 13].

While it seems intuitively obvious that this must be the case for a number of genes, it turns out that very little such interdependence has been seen in mutation experiments that knocked out single genes, then pairs of genes, and so on in microorganisms [14, 3, 4]. In general, genes can interact in such a way that the loss of two genes is worse than the combined losses of each of them alone (called *synergistic*, or reinforcing epistasis), or they could interact in a *mitigating* manner, so that the combined loss of both genes simultaneously weighs less heavily than the cumulative loss of both individually (*antagonistic* epistasis). While it was found that many pairs of genes interact, they appear to do so in roughly equal proportions of reinforcing and mitigating pairs, leading to a net absence of epistatic effect (no net *directional* epistasis). A similar trend has been observed in digital organisms [8], where the effect of single or multiple mutations on the fitness of the whole organism (rather than single genes) has been tested, much as in [3, 4].

The average directionality of epistatic interactions between mutations can be modeled by a single parameter if we parameterize the mean effect of mutations on fitness in the following manner [8, 1]. If the fitness of the wild-type sequence is u_0 , we write the average fitness of all n -mutants (sequences with n mutations) as

$$\bar{w}(n) = u_0 \exp(-\alpha n^\beta), \quad (1)$$

with parameters α and β that are to be obtained from a fit to experimental data. The parameter α characterizes the rate of decay of fitness as mutations accumulate, and can be estimated by examining the fitness of all one-point mutations:

$$\alpha = -\ln(\bar{w}(1)/u_0). \quad (2)$$

The average rate of decay, on the other hand, just reflects the *robustness* with respect to deleterious mutations of the organism: If α is large, the organism is very fragile and fitness decays quickly, while a low α implies that many mutations are *neutral*, that is, they do not change the organism's fitness. The *neutrality* ν of a sequence, defined as the number of sequences that are one point mutation away from the wild type sequence and that have the same fitness as the wild-type, is related to α via

$$\nu = \ell(D - 1)e^{-\alpha}, \quad (3)$$

where ℓ is the sequence length, and D is the dimension of the alphabet the sequence is drawn from ($D = 4$ for RNA or DNA sequences, $D = 20$ for proteins, etc.).

The epistasis parameter β can be obtained from Equation 1 if we can extract at least $\bar{w}(2)$ by doing all 2-point mutations, because

$$\beta = \frac{1}{\ln 2} \ln \left[\frac{\ln(\bar{w}(2)/u_0)}{\ln(\bar{w}(1)/u_0)} \right]. \quad (4)$$

Ideally, though, we would like to obtain an estimate of $\bar{w}(3)$ and higher by sampling from the possible 3- and more-point mutations, to verify that the parameterization in Equation 1 is justified.¹ This parameterization has proven to be quite good out to $n = 7$ and higher for a number of different computational systems, including digital life [8], RNA folding [21, 23], and a computational model of the T7 bacteriophage [25].

If $\beta = 1$, then on average all mutations are *independent* (no directional epistasis on average), even though pairs of mutations can still interact synergistically or antagonistically; it's just that their contributions cancel [8]. If mutations are independent, then $\bar{w}(2) = \kappa \bar{w}(1) = \kappa^2 u_0$ with some constant $\kappa \leq 1$, so that, from Equation 4

$$\beta = \frac{1}{\ln 2} \ln \left[\frac{\ln \kappa^2}{\ln \kappa} \right] = 1. \quad (5)$$

A β value that exceeds 1 signals synergistic epistasis on average, while $\beta < 1$ indicates antagonistic epistasis (see Figure 1).

According to Kondrashov's theory [6], synergistic epistasis is conducive to the evolution of sexual recombination. One possible mechanism for synergistic interactions is genetic *redundancy* (so that a few mutations are not harmful, while a larger number may knock out the last remaining redundant pathway). And indeed, genetic redundancy is difficult to preserve in asexual organisms.² Within our asexual digitals, we thus expect antagonistic epistasis throughout, as is observed [8]. In antagonistic epistasis, a few mutations impair the fitness seriously, while additional ones do not carry the same punch. This phenomenon was recently shown to be due to compensatory mutations that restore some of the lost fitness [23]. Here, we follow a different (but

¹ The number of possible 3-point mutations for a sequence of length $\ell = 100$ and $D = 20$, for example, is already $n(3) \approx 1.1 \times 10^9$. In practice, we cannot exhaustively test all $n > 3$ -mutants even for digital genomes.

² We thank Charles Ofria for pointing this out to us.

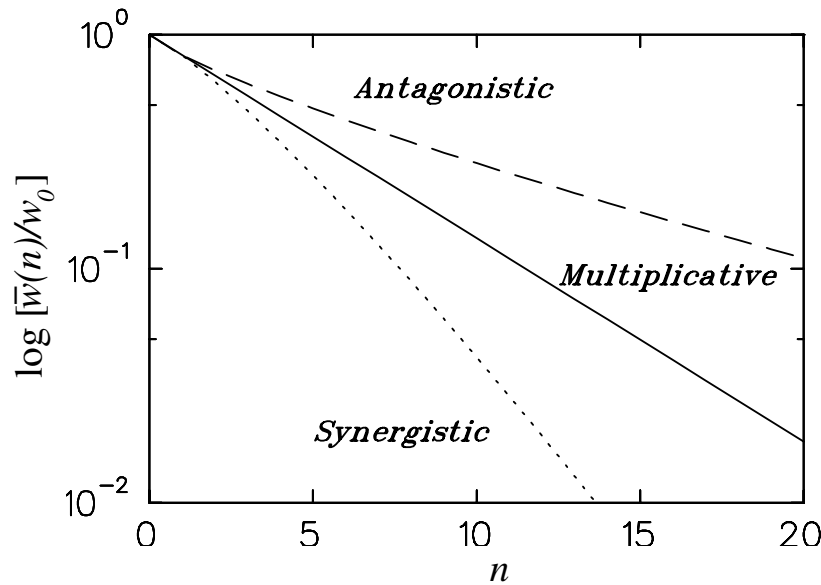


Figure 1. Model fitness decay with different epistasis parameters β (but with the same decay parameter $\alpha = 0.2$): multiplicative, $\beta = 1$ (solid line); antagonistic epistasis, $\beta = 0.8$ (dashed curve); and synergistic epistasis, $\beta = 1.2$ (dotted curve).

possibly complementary) reasoning. Is antagonistic epistasis due to efficient, but fragile coding?

That epistasis can reflect poor coding choices is perhaps best illustrated by a simple example. Suppose a code is designed in such a way that three blocks coding for different functions are inherently interdependent so that mutating any of the three blocks will reduce the fitness close to zero. In that case, any subsequent mutation will not have much of an effect, as most of the fitness was lost due to the first mutation. This is the typical sign of antagonistic epistasis. Imagine that it is possible to recode the same information in such a way that each block is independent of the other. In that case, we would see no epistasis, that is, $\beta \approx 1$. This is reminiscent of the *linkage* of genes in genetics, which can be broken up through sexual recombination. Linkage can be harmful to an organism if beneficial genes are linked (by location) with detrimental ones. In a sense, we hypothesize that robustness evolves in digital organisms by breaking up harmful linkages in the code so as to move a $\beta < 1$ towards $\beta = 1$.

To test the hypothesis, we evolved populations of digital organisms to high fitness at low mutation rate, only to transfer them to a high mutation rate environment, where we expect selection of robustness to take place. From the moment of transfer, we measure the change in robustness by monitoring α *along the line of descent*. In this manner, we avoid having to calculate α and β for each member of the population, while still capturing the essential evolution. At the same time, we measure β along the line of descent, and check for a tendency of β to move towards $\beta = 1$.

3 Materials and Methods

For these experiments we used Avida version 2.0b3, freely available under the GNU General Public License (GPL) from SourceForge.net.³ For setting up the experiment,

³ <http://sourceforge.net/projects/avida/>.

we used the default values in the `genesis` file (see [12]) unless otherwise specified. The population size was limited to 2,500 organisms by setting the world size to 50×50 . The organisms were also constrained to length 100 by setting `CHILD_SIZE_RANGE` to 1. Additionally, all insert and delete mutations were set to 0. By fixing the length to 100 we are able to guarantee a fixed genomic mutation rate and focus on how the organism recodes information in a fixed length.

We conducted 21 independent experiments (*replicas*), which involved first breeding an adapted organism at low mutation rate, and then transferring to a high mutation rate to observe the evolution of robustness. From the replicas, we collected 5,699 organisms on the line of descent, excluding six for which α and β measurements were unreliable. These measurements are performed in analyze mode in the Avida *Test-CPU* (see [12]), and involve running the organism on an isolated CPU to obtain fitness estimates. Because the test environment is different from the population the organism was extracted from, and because the length of the test is necessarily finite, sometimes this test returns the value zero for u_b , preventing a calculation of α and β . This is clearly an artifact of the Test-CPU, because in almost all cases an organism on the line of descent is viable.

We averaged the data from these runs by creating variable-size bins so that each bin contained the organisms appearing in the line of descent of 1,000 updates. This gave rise to 10 bins with between 459 and 624 organisms. The average and standard error of α , β , $\log w$, and P_v was calculated for each bin.

3.1 Evolution at Low Mutation Rate

We started by evolving the default length 100 ancestor for 100,000 updates in a low mutation environment. This environment had a per-site copy mutation rate of 0.001, which translates to a genomic mutation rate of 0.1.⁴ In order to save computational time, all statistical logging functions of Avida were turned off in the `genesis` file and the `events.cfg` file. A single event “u 100000 print_dom DOMINANT_ORG” was added to the `event.cfg` file to print the dominant organism at update 100,000 to the file `DOMINANT_ORG`. The `environment.cfg` was modified to include: the `ECHO` task with an *additive* reward of 0.5, all the two-input tasks (`ADD`, `SUB`, `NOT`, `NAND`, `AND`, `ORN`, `OR`, `ANDN`, `NOR`, `XOR`, `EQU`) with an additive reward of 2, and all the three-input tasks (`logic_3AA` through `logic_3CP`) with an additive reward of 5. For each generation the organism was only rewarded for a specific task once.

We used a different random seed for each run, and the population was allowed to evolve and adapt to its environment for 100,000 updates. At this point, we extracted the genotype with the greatest abundance (the “dominant”) from the population, and saved its genome along with the list of tasks that it was capable of doing, for the next phase of the experiment.

3.2 Introduction to a High Mutation Environment

We proceeded by seeding each of the dominant organisms from the low mutation runs into an environment with the following two changes. We set the per-site copy mutation rate to 0.02, which implies a genomic mutation rate of 2 (on average two mutations per genome per generation). Also, we made sure that the organisms could not evolve any new tasks after transfer, by changing the `environment.cfg` file accordingly, that is, limiting the available bonuses to those that each particular seed organism could already do at transfer). In this manner, we are able to focus on evolutionary changes that increase neutrality, as beneficial mutations are all but ruled out.

⁴ This mutation rate is low only compared to the high-rate environment we transfer the populations into later. Compared to biochemical organisms, a genomic mutation rate of 0.1 is fairly high, comparable to that of many RNA viruses.

We recorded fitness, robustness, and code neutrality, in particular the coefficients α and β defined in Equations 2 and 4, for all organisms on the line of descent of the most abundant organism at the end of the high mutation rate run. Naturally, this line of descent ends with the dominant organism used to seed the run. In this manner, we are able to track all important changes in the lineage leading from the seed to the end of the run, without having to analyze a large number of organisms at fixed time intervals. To facilitate this, we turned on the option `TRACK_MAIN_LINEAGE` in the `genesis` file. Also, we modified the `events.cfg` to execute the `detail_pop` and the `dump_historic_pop` events after 10,000 updates (see the online documentation to Avida v2.0⁵). These events output information about the population that are necessary to reconstruct the lineage of the dominant organism after 10,000 updates.

3.3 Analysis of the High Mutation Lineage

In order to use Avida in its analyze mode, we modified the `analyze.cfg` file to instruct Avida to load in the files `detail_pop` and `historic_dump`. We then used the command `FIND_LINEAGE num_cpus` to reduce the organisms in the analysis batch to those on the lineage of the most dominant organism after 10,000 updates in the high mutation environment. We obtained statistics about the fitness effect of all possible 1- and 2-point mutations on genomes on the line of descent by using two `LANDSCAPE` commands, generating $\bar{w}(1)$ and $\bar{w}(2)$. This data allows us to calculate α and β using Equations 2 and 4. The probability of a neutral mutation, P_v is calculated in the distance 1 landscaping calculation as the fraction of single-point mutations that do not change the fitness:

$$P_v = \frac{\text{\#neutral mutations}}{(D-1)\ell} . \quad (6)$$

This probability is related to the fraction of neutral mutations $e^{-\alpha}$ (see Equation 3), but only becomes equal to it in the limit that all mutations are either neutral or lethal.

3.4 Knockout Arrays

Knockout (functional-genomic) arrays were generated for the ancestral and final organisms in the high-mutation run depicted in Figures 2–4 in Avida analyze mode, by inserting the command `MAP_TASKS` in the file `analyze.cfg` with a list of tasks that the organism is proficient in.

4 Results

After 100,000 updates in a low mutation rate environment, the mean fitness across runs had increased to 15.92 ± 0.79 by acquiring 11.2 ± 0.3 tasks. Upon transfer into the high mutation environment, fitness was lost on average, as expected (mean fitness across runs after 10,000 updates was 9.92 ± 0.45).

4.1 Case Study

In Figure 2, we show the logarithm of the replication rate along the line of descent, $\log w$, as well as the probability for a mutation to be neutral, P_v , for a particular representative run. The tendency for fitness to be shed in exchange for increased neutrality is evident. There are 278 organisms on the line of descent of this replica. The increase in neutrality is mirrored by a decrease in α , as shown in Figure 3, where we also show the corresponding increase in the epistasis parameter β . That α and β

⁵ <http://www.dllab.caltech.edu/avida/v2.0/docs>.




Figure 2. (a) Logarithm of replication rate, $\log(w)$, and (b) probability P_v for a mutation to be neutral along the line of descent for a typical run, as functions of time measured in updates. Note the presence on the line of descent of a severely impaired organism around update 6,000.

are strongly correlated (as predicted by theory [21]) can be seen in Figure 4, where we plotted α versus β for all 278 organisms along the line of descent.

We can gain some insight into the code rearrangement that led to the increase in neutrality by studying the *functional-genomic arrays* of the first and last organism on the line of descent of this example, shown in Figure 5. These arrays are obtained by successively replacing each instruction in the genome by a nonfunctional (null) instruction and testing for the fitness of such a mutated organism and for how the mutation affects the tasks the program does (its functionality). Such arrays reveal those sections of code that affect particular tasks, and those that are neutral. The organism at the start of the line of descent could perform 12 tasks: the 1-input task ECHO, four 2-input tasks (ADD, SUB, NAND, ORN), and seven 3-input tasks. The last are worth the most (in terms of energy production) because they are rewarded with a bonus of 5 (see Section 3). A total of 326 mutations knocked out tasks for the fragile organism, compared to 149 mutations for the robust one. To achieve this, the robust organism at the end of the line of descent got rid of six tasks: all of the 2-input tasks and two 3-input tasks. But the gain in robustness was remarkable, as mutations on all the eliminated tasks also affected the remaining five 3-input tasks! Furthermore, 15 mutations on the robust organism actually recreated some of the lost tasks. By removing the lesser-rewarded 2-input tasks and two of the 3-input tasks, the final organism was able to reduce the number of mutations affecting the five retained 3-input tasks from 174 to 134, a reduction of 23%.




Figure 3. (a) Robustness parameter α , and (b) epistasis parameter β along the line of descent (same run as Figure 2), as functions of time.

We also note another remarkable observation from these knockout arrays. Almost all the changes in code between the organisms at the beginning and end of the line of descent took place in one particular stretch of code (from instructions 7 to 71), while the rest was virtually unchanged. This is understandable for the copy-loop coded in instructions 1–6 and 97–98, but it was unexpected to see a block of code (from 72 to 96) survive intact while almost every other instruction changed. We surmise that this block (coding for the five 3-input tasks that were retained) is particularly difficult to break up. Indeed, most mutations in this area affect at least three of the five tasks. It is conceivable that more robustness can evolve if these modules can be broken up further, but this may take considerably more evolutionary time.

4.2 Average across Replicas

The trend apparent in Figures 2–4 is consistent among all the runs, with variations. For example, while the fitness always drops during the first few thousand updates after transfer, it recovers in some runs to close to the original level. This happens despite the fact that new tasks cannot be acquired in this phase, and organisms cannot gain fitness by repeating a task that is rewarded. It is possible that this fitness gain is due to the evolution of a more efficient replication loop, but we suspect that this is not the case. Rather, we surmise that while organisms jettison some code in order to decrease epistasis between mutations (for example, by sacrificing this or that minor task, as seen in the case study), they sometimes succeed in recovering this fitness *without increased epistasis*. Thus, while organisms tend to leave high and narrow peaks in favor of




Figure 4. α versus β for all organisms along the line of descent (same run as Figures 2 and 3).

broader but lower peaks initially, sometimes they succeed in ultimately finding a peak that is just as high as the initial one, but offering a lot more neutrality to boot. In Figures 6–8 we show the statistics shown above averaged over all runs, by binning organisms from all replicas as described in Section 3.

5 Discussion

In this article, we have examined the evolution of mutational robustness in digital organisms, and found that the changes in robustness are directly correlated to changes in the epistasis parameter β , from largely antagonistic to multiplicative. Because no fitness increases were possible once the organisms were transferred to the high mutation environment, the only path towards higher realized fitness for the population lies in recoding the information at hand in a more robust manner. It appears that this can be achieved fairly easily, but we also note that the increase in robustness goes hand in hand with a loss of replicative ability. Indeed, it seems that recoding has its price, or conversely, that the benefit of high fitness has to be bought with fragile, interdependent coding.

The give-and-take between robustness to single mutations and directional epistasis that we observe in digitals has previously been seen in other computational systems [21] and even in quantitative trait loci (QTL) data from mice [15]. The correlation between α and β was explained in [21] by a simple geometric argument, and appears to be due to the finiteness of neutral networks. The same correlation allows us to answer the

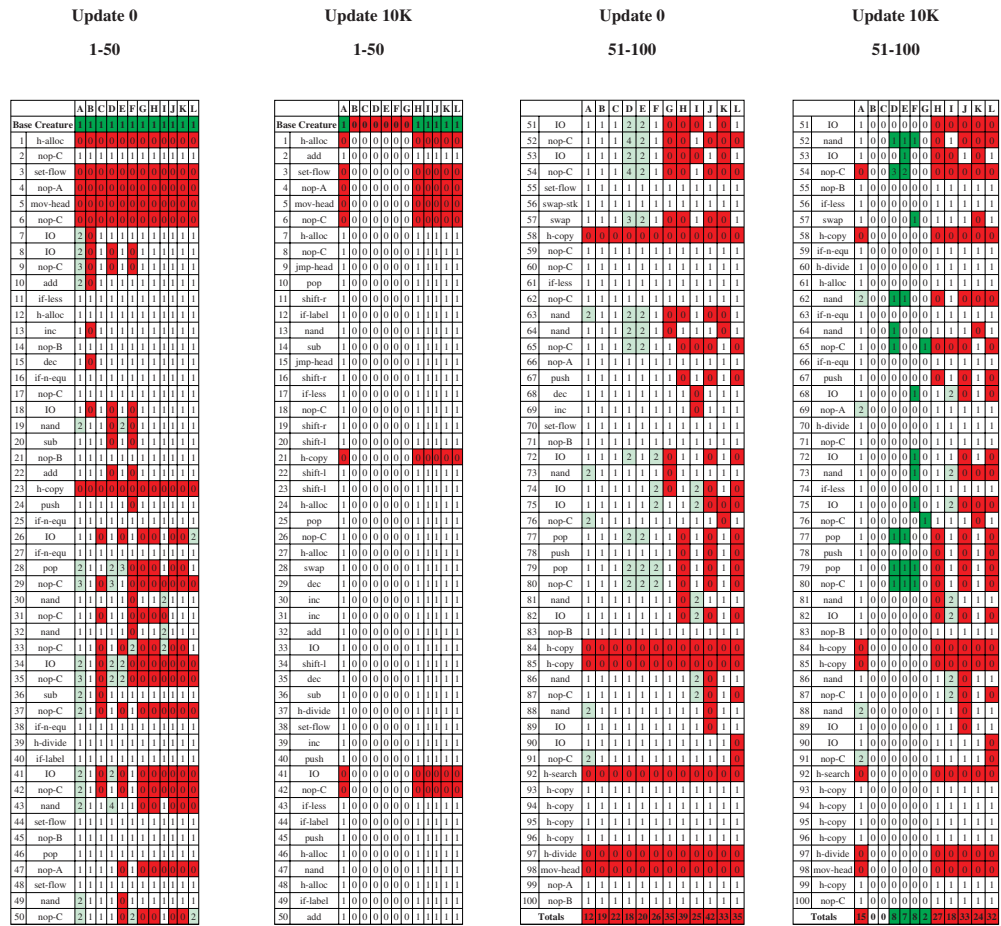


Figure 5. Comparison of the functional-genomic array for the two organisms at the beginning (update 0) and end (update 10,000) of the line of descent for the run depicted in Figures 2–4. The first 50 instructions are compared in the left half of the figure, and the last 50 in the right half. Each column represents a task, where A=ECHO, B=ADD, C=SUB, D=NAND, E=ORN, F=3AG, G=3AX, H=3BA, I=3BO, J=3BY, K=3BZ, and L=3CI. (The abbreviations 3XX stand for particular three-input logical tasks, which are described in the environment configuration file of Avida.) Each row represents an instruction, with the instruction number given in the first column and the mnemonic in the second. The effect of a mutation of this instruction on the task in the respective column is indicated by the color of the box: no color means the task is unaffected by the knockout of this instruction, red means the task is eliminated, and green indicates that the task is created or enhanced. The numeral inside the box represents the number of times this task is executed per generation by the organism. In this environment, only the first execution is rewarded. The first row, labeled “Base Creature,” designates whether the task in the columns above is being performed by the organism.

perennial question that crops up when studying the evolution of robust codes: how do they do it, how is that robustness achieved? It appears “they” do it by loosening the tight linkage between important sections of the code, up until they are largely independent.

Acknowledgments

We thank Charles Ofria and Claus Wilke for suggestions and discussions that led to the final version of this project. J.A.E. thanks D. Allan Drummond and Daniel Wagenaar for guidance. This work was supported by the National Science Foundation’s Biocomplexity Program under Contract No. DEB-9981397. All data were obtained on




Figure 6. (a) Logarithm of replication rate, $\log(w)$, and (b) probability P_v for a mutation to be neutral, along the line of descent averaged over all runs, as functions of time measured in updates.

a cluster of 160 processors at Michigan State University, supported by the NSF under the aforementioned contract.

References

1. Adami, C., Ofria, C., & Collier, T. (2000). Evolution of biological complexity. *Proceedings of the National Academy of Sciences of the U.S.A.*, *97*, 4463–4468.
2. de Visser, J. A. G. M., Hermisson, J., Wagner, G. P., Meyers, L. A., Bagheri-Chaichian, H., Blanchard, J. L., Chao, L., Cheverud, J. M., Elena, S. F., Fontana, W., Gibson, G., Hansen, T. F., Karakauer, D., Lewontin, R. C., Ofria, C., Rice, S. H., von Dassow, G., Wagner, A., & Whitlock, M. C. (2003). Evolution and detection of genetic robustness. *Evolution*, *57*, 1959–1972.
3. Elena, S., & Lenski, R. E. (1997). Test of synergistic interactions among deleterious mutations in bacteria. *Nature*, *390*, 395–398.
4. Elena, S. F. (1999). Little evidence for synergism among deleterious mutations in a nonsegmented RNA virus. *Journal of Molecular Evolution*, *49*, 703–707.
5. Kimura, M., & Maruyama, T. (1966). The mutational load with epistatic gene interactions in fitness. *Genetics*, *54*, 1337–1351.
6. Kondrashov, A. S. (1982). Selection against harmful mutations in large sexual and asexual populations. *Genetical Research*, *40*, 325–332.
7. Kondrashov, A. S. (1993). Classification of hypotheses on the advantage of amphimixis. *Journal of Heredity*, *84*, 372–387.




Figure 7. (a) Robustness parameter α , and (b) epistasis parameter β , along the line of descent, averaged over all runs as described in Section 3.

8. Lenski, R. E., Ofria, C., Collier, T. C., & Adami, C. (1999). Genome complexity, robustness and genetic interactions in digital organisms. *Nature*, *400*, 661–664.
9. Lenski, R. E., Ofria, C., Pennock, R. T., & Adami, C. (2003). The evolutionary origin of complex features. *Nature*, *423*, 139–144.
10. van Nimwegen, E., Crutchfield, J. P., & Huynen, M. (1999). Neutral evolution of mutational robustness. *Proceedings of the National Academy of Sciences of the U.S.A.*, *96*, 9716–9720.
11. Ofria, C., Adami, C., & Collier, T. C. (2003). Selective pressures on genomes in molecular evolution. *Journal of Theoretical Biology*, *222*, 477–483.
12. Ofria, C., & Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, *10*, 191–229.
13. Peters, A. D., & Lively, C. M. (2000). Epistasis and the maintenance of sex. In J. Wolf, E. D. Brodie, III, & M. J. Wade (Eds.), *Epistasis and the evolutionary process* (pp. 99–112). Oxford, UK: Oxford University Press.
14. de Visser, J. A. G. M., Hoekstra, R. F., & van den Ende, H. (1997). An experimental test for synergistic epistasis and its application to *Chlamydomonas*. *Genetics*, *145*, 815–819.
15. Wagner, G. P., Laubichler, M. D., & Bagheri-Chaichian, H. (1998). Genetic measurement theory of epistatic effects. *Genetica*, *102/103*, 569–580.
16. Whitlock, M. C., Philips, P. C., Moore, F. B.-G., & Tonsor, S. J. (1995). Multiple fitness peaks and epistasis. *Annual Review of Ecology and Systematics*, *26*, 601–629.
17. Wilke, C. O., Wang, J. L., Ofria, C., Lenski, R. E., & Adami, C. (2001). Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature*, *412*, 331–333.




Figure 8. α versus β for all organisms along the line of descent, all runs (see Section 3).

18. Wilke, C. O. (2001). Adaptive evolution on neutral networks. *Bulletin of Mathematical Biology*, 63, 715–730.
19. Wilke, C. O. (2001). Selection for fitness vs. selection for robustness in RNA secondary structure folding. *Evolution*, 55, 2412–2420.
20. Wilke, C. O. (2003). Probability of fixation of an advantageous mutant in a viral quasispecies. *Genetics*, 163, 467–474.
21. Wilke, C. O., & Adami, C. (2001). Interaction between directional epistasis and average mutational effects. *Proceedings of the Royal Society of London, B* 268, 1469–1474.
22. Wilke, C. O., & Adami, C. (2003). Evolution of mutational robustness. *Mutation Research*, 522, 3–11.
23. Wilke, C. O., Lenski, R. E., & Adami, C. (2003). Compensatory mutations cause excess of antagonistic epistasis in RNA secondary structure folding. *BMC Evolutionary Biology*, 3, 3.
24. Wolf, J., Brodie, E. D., III, & Wade, M. J. (Eds.) (2000). *Epistasis and the evolutionary process*. Oxford, UK: Oxford University Press.
25. You, L., & Yin, J. (2002). Dependence of epistasis on environment and mutation severity as revealed by in silico mutagenesis of phage T7. *Genetics*, 160, 1273–1281.