

Evolvability Tradeoffs in Emergent Digital Replicators

Thomas LaBar^{**}
Arend Hintze[†]
Christoph Adami^{**‡}
Michigan State University

Abstract The role of historical contingency in the origin of life is one of the great unknowns in modern science. Only one example of life exists—one that proceeded from a single self-replicating organism (or a set of replicating hypercycles) to the vast complexity we see today in Earth’s biosphere. We know that emergent life has the potential to evolve great increases in complexity, but it is unknown if evolvability is automatic given any self-replicating organism. At the same time, it is difficult to test such questions in biochemical systems. Laboratory studies with RNA replicators have had some success with exploring the capacities of simple self-replicators, but these experiments are still limited in both capabilities and scope. Here, we use the digital evolution system Avida to explore the interplay between emergent replicators (rare randomly assembled self-replicators) and evolvability. We find that we can classify fixed-length emergent replicators in Avida into two classes based on functional analysis. One class is more evolvable in the sense of optimizing the replicators’ replication abilities. However, the other class is more evolvable in the sense of acquiring evolutionary innovations. We tie this tradeoff in evolvability to the structure of the respective classes’ replication machinery, and speculate on the relevance of these results to biochemical replicators.

Keywords

Evolvability, digital life, origin of life

I Introduction

The science surrounding the origin of life presents an obvious difficulty to scientists: Our data consist of a single example. It is unknown whether the events that followed the origin of life are deterministic consequences of life or unique to this specific origin. One question of interest is whether all life, or specifically all self-replicators, are evolvable. Evolvability (the ability to evolve) has many similar, but differing definitions [59, 29, 58]; here, we define it as the ability to increase in fitness. We also distinguish between two possibilities for such fitness increases: *optimization*—defined

* Contact author.

** Department of Microbiology & Molecular Genetics, BEACON Center for the Study of Evolution in Action, and Program in Ecology, Evolutionary Biology, and Behavior, Michigan State University, East Lansing, MI 48824. Email: labartha@msu.edu

† BEACON Center for the Study of Evolution in Action, Department of Integrative Biology, and Department of Computer Science and Engineering, Michigan State University, East Lansing, MI 48824. E-mail: hintze@msu.edu

‡ Department of Microbiology & Molecular Genetics, BEACON Center for the Study of Evolution in Action, Program in Ecology, Evolutionary Biology, and Behavior, and Department of Physics and Astronomy, Michigan State University, East Lansing, MI 48824. E-mail: adami@msu.edu

as the ability to improve an already present phenotypic trait, and *innovation*—the ability to evolve novel phenotypic traits.

A common assumption is that life originated with self-replicating RNA molecules [21, 28]. Thus, most empirical studies have focused on RNA replicators, either emergent replicators [53, 8, 9] or those created by experimenters [33, 48]. Experiments (both computational and biochemical) have also explored the evolvability of RNA replicators, usually involving extensive mapping of their fitness landscapes [25, 26, 6, 46]. While RNA is an enticing candidate for a pre-biotic molecule, the so-called “RNA world hypothesis” has its own problems—foremost the “asphaltization” that tends to befall formose carbohydrates under the expected conditions of an early Earth [14, 40]—which has led researchers to explore other origins of life not necessarily RNA-based [42, 47], or even to move the origin of life from Earth to Mars [37].

In recent years, more and more theoretical models concerning the origin of life have focused on exploring the abstract concepts that could possibly be involved in any potential origination, independent of a particular biochemical system [60, 19, 2]. The field of artificial life is ideally suited to study various possibilities for the origin of life as it imagines life as it could be, not just as it is. The question of the random emergence of replicators has been addressed in various digital systems before [43, 44, 3, 35], while theoretical models have explored the factors that lead to differing evolvability [5, 16, 15]. Artificial life tools have also been used to explore the potential of evolvability in different systems [54, 51, 7, 17, 18, 20].

Recently, Adami used information theory to calculate the likelihood of the random emergence of a self-replicator—in a sense, the progenitor to life—without regard to a specific biochemistry [2]. Adami and LaBar tested this theory with Avida by generating billions of random Avidian sequences and checking for their ability to replicate themselves [3]. Such an investigation is akin to studying the chance emergence of self-replicating RNA molecules [8]. In previous work we explored the relationship between evolvability and self-replication using these emergent replicators [30], and found that almost all emergent replicators were evolvable, both in terms of optimization for replication and in terms of evolving beneficial phenotypic innovations. We also discovered that these replicators came in two forms: proto-replicators and self-replicators. Proto-replicators deterministically copy themselves inaccurately, but eventually evolve into self-replicators; self-replicators, on the other hand, produce an exact copy of themselves in the absence of stochastic mutation. We also noted the possibility of an optimization-innovation tradeoff in some of these replicators, especially the default Avida ancestor (a hand-written self-replicator specifically designed for evolution experiments).

Here, we extend our previous study and test a fundamental question concerning life’s origins: How does the genetic composition of the first replicator determine the future evolution of life? One extreme possibility is that all emergent replicators have similar genetic compositions and thus the future evolution of life will occur in a similar manner, no matter what the progenitor. On the opposite end of the spectrum is the possibility that every emergent replicator is sufficiently different from every other replicator in genetic composition. In this case, the future outcome of life may be entirely dependent on which replicator emerges first. In experiments with the artificial life system Avida, we find that the interplay between the genetic composition of the first emergent replicator and future evolutionary outcomes is between these two extremes. Emergent replicators can be classified into two distinct classes based on a functional analysis of their replication machinery. These classes differ in their ability to optimize their replication ability. However, those replication classes that display high evolvability towards optimizing replication also demonstrate low evolvability towards evolving novel phenotypic traits, and vice versa. Finally, we show that this difference in evolvability is due to differences in the replication machinery between the different replicator classes.

2 Methods

2.1 Avida

In order to study the interplay between emergent replicators and evolvability, we used the Avida digital evolution platform [1, 41]. In Avida, simple computer programs (*Avidians*) in a population

compete for the memory space and CPU time needed to replicate themselves. Each Avidian consists of a genome of computer instructions, where each locus in the genome can be one of 26 possible instructions. Contained within each genome are the instructions necessary for the Avidian to allocate a daughter genome, copy itself into this new daughter genome, and divide off the daughter genome. As the replication process is mechanistic (i.e., Avidians execute their genome's instructions, including those to replicate, sequentially), the speed at which replication occurs is also genome-dependent. Therefore, fitness is genome-dependent, as an organism's ultimate success is determined by its replication speed in these simple environments.

As Avidians directly copy and pass their genomes to their daughters, Avida populations also have heredity. During the copying process, errors may be introduced, resulting in mutations and population variation. Therefore, Avida populations possess heredity, variation, and differential fitness: They are an *instantiation* (as opposed to a simulation) of Darwinian evolution [45]. This simpler instantiation of evolution by natural selection has allowed for the exploration of many topics hard to test in biological systems; see for example [31, 4, 39, 12, 13, 22, 62].

Avida has been explained in greater detail elsewhere (see [41] for a full description); here, we will cover the details relevant for this study. An Avidian consists of a variety of elements: a genome of instructions, a *read head* that marks which instruction should be copied, a *write head* that denotes where the read-head-marked instruction should be copied to, and three registers to hold integer numbers (AX, BX, and CX), among other elements.

In order to undergo reproduction, an Avidian needs to perform four operations in the following order. First, it must execute an *b-allocate* instruction, which allocates a fixed number of *nop-A* instructions to the end of its genome (here 15, as we work exclusively with length-15 genomes). These *nop-A* instructions are inert by themselves, and serve as placeholders to be replaced by actual information. Allocation of this memory prepares the daughter genome space to receive the information from the parent. Next, the read head and write head must be set 15 instructions apart, allowing for the instructions in the parent genome to be copied into the daughter genome. These operations are algorithmically similar to creating a DNA replication fork at the origin of replication, preparing for the assembly of a copied sequence. Following this step, the genome must find a way of looping over the *b-copy* instruction in the genome to actually copy instructions from the parent into the daughter genome, similar to the action of a DNA polymerase fusing the new (daughter) nucleotides on the former parent strand. This copying can be done either by looping through the entire genome (using the circular nature of the genome to reuse a single *b-copy* command as many times as necessary to copy all instructions) or else by continuously looping over a smaller set of instructions in the genome (called the *copy loop*). The latter algorithm requires marking the set of instructions (the "replication gene") so as to control the forking of execution flow. Finally, an Avidian must execute an *b-divide* instruction to divide the duplicated genome into two Avidians, and thus successfully reproduce.

In the experimental design used here, we used two different mechanisms to produce mutations during replication. The first mechanism produces at most one *point mutation* at a random locus in the genome at a rate of 0.15 mutations per genome per generation upon successful division. The second mechanism is a deterministic (or *incipient*) mutation, which occurs when the instructions in an Avidian cause it to copy instructions into the daughter genome in the wrong place. For example, if the copy process skips the first one or two instructions of the genome, those instructions will remain in the default state (here *nop-A*) instead of in the state of the parent genome, that is, they appear to be mutated. In many cases, such a faulty copy algorithm results in the offspring being nonviable; however, in some cases viable offspring are produced. Sometimes, this incipient mutation mechanism results in the phenomenon of *proto-replicators*, defined as those replicators that deterministically (i.e., reproducibly because genetically controlled) make an offspring different from themselves [30]. Replicators that instead deterministically make an identical offspring are called self-replicators.

In the standard experimental evolution scenario, Avidians are placed into a landscape that rewards a variety of phenotypic traits. These traits commonly refer to the ability to perform certain

Boolean logic calculations on binary numbers that the environment provides. During an Avidian's life span, they can input (read from) and output (write to) binary numbers from/to the environment. Whenever a number is written, the Avida world checks if a Boolean logic calculation was performed. Successful performance results in an increase in the replication speed of that individual's descendants. In the standard Avida design (also the one used here in the experiments to explore evolvability in the sense of innovation), that environment rewards the performance of nine different calculations. This environment is usually referred to as the "Logic-9" environment and rewards calculations such as NOT and EQUALS [32]. The more complex the calculation performed, the greater the replication speed increase. The performance of such logic calculations can be viewed as an algorithmic analogue of performing different metabolizing reactions using different sugar resources (the binary numbers provided by the environment).

2.2 Experimental Design

To study the evolvability of emergent replicators, we first had to generate a collection of emergent replicators. We re-analyzed a list of 10^9 random Avidian length-15 genomes generated previously [3]. In order to discover replicators of fixed genome size, we set the OFFSPRING SIZE RANGE parameter in the main Avida configuration file to 1.0, which guarantees that the *allocate* command allocates exactly as much space as needed by the daughter genome. For our focus on replicators of length 15, this guaranteed that exactly 15 instructions are allocated for the daughter to receive.¹ Next, we looked at the relative abundance of proto-replicators compared to true self-replicators. We re-analyzed the above set of replicators, but made one additional parameter change: We set REQUIRE EXACT COPY to 1. Any replicator that registered as having zero fitness in this analysis would deterministically copy its genome inaccurately, and would be classified as a proto-replicator. Any replicators that could reproduce themselves accurately under this treatment were classified as self-replicators.

In order to test the evolvability of these replicators, we performed three different experimental tests. For all experiments, we used a population size of 3600. Individuals' offspring were placed in any random cell in the 60×60 grid, thus mimicking a well-mixed environment. Point mutations were applied upon division, and the genomic mutation rate was 0.15 mutations per generation. The genome size was fixed at 15 instructions. The first experiments tested the replicators' evolvability in the sense of replication optimization (optimization experiments). In this experiment, each replicator was used to seed 10 populations; these populations evolved for 10^3 generations. Phenotypic traits were not under selection in these experiments, that is, evolution only optimized the replication machinery. Next, we performed evolution experiments (referred to as the *innovation* experiments) where we rewarded individuals for evolved phenotypic traits (i.e., logic tasks) besides any increase in replicatory prowess.

We ran the innovation experiments for 10^4 generations (an order of magnitude longer than the optimization experiments). These experiments were designed to decrease the likelihood that our data resulted from one of the replicator classes taking longer to evolve complex traits. Finally, we repeated the innovation experiments, but used each ancestral replicator's fittest descendant from the optimization experiments to initialize the populations (tradeoff experiments). These experiments tested the presence of an optimization–innovation tradeoff.

2.3 Data Analysis

Once we determined the set of replicators, we first tried to determine if there were similarities in the genomes of the replicators. We realized that many of the genomes had similar instruction motifs within the sequence. To test whether different replication algorithms existed in different replicators,

¹ In previous studies [3, 30], OFFSPRING SIZE RANGE was set to 2.0, which allocated an extra number of loci that typically would not be filled by the replication process, and could interfere with accurate self-replication).

we used Avida's TRACE function to analyze the step-by-step execution of the replicators' genomes. This allowed us to cluster the emergent replicators into two main distinct functional replicator classes.

To examine the results from the evolution treatments, we analyzed the most abundant genotype at the end of each experiment. The data collected included the final evolved fitness and the number of evolved phenotypic traits. All Avida analysis was performed using Avida's Analyze mode with settings matching those under which the relevant experiments were performed. Statistical analyses were performed using the NumPy, SciPy, and Pandas Python modules [55, 27, 36]. Figures were generated using the Matplotlib Python module [24].

3 Results

3.1 Emergence of Replicators

Among the one billion randomly generated genomes, we found 75 genomes that could replicate themselves when genome size was fixed. Of these replicators, 22 were true self-replicators in the sense that they could perfectly copy their genomes even when mutations were turned off. The remaining 53 replicators were "proto-replicators" in the sense that they could *not* produce a perfect copy of their genome when the mutation rate was set to zero. This deterministic miscopying is due to the specific nature of a proto-replicator's genome (see Section 2). However, these replicators still produced viable offspring that would eventually lead to a self-replicator. The discovery of these proto-replicators extends the definition of proto-replicators from [30] to include fixed-length proto-replicators.

3.2 Replication Mechanisms of Emergent Replicators

Upon examination of these replicators, we detected the presence of distinct instruction motifs in their genomes. These motifs consist of instructions involved in genome replication and suggested that different replicators used different replication mechanisms. To explore this possibility, we performed a step-by-step analysis of each replicator's lifestyle by looking at the execution of their genome's instructions. Avidians must perform four steps to successfully reproduce: allocate a blank daughter genome, separate their read and write heads, copy their genome into the daughter genome through some looping process, and divide off the daughter genome (see Section 2 for a fuller description). We were able to cluster the replicators into two replication classes based on a difference in two traits: (1) how they separated their read and write heads to copy their genome and (2) how they looped through their genome in order to copy their genome.

We named the first class of replicators *bc* replicators because of the *hc* instruction motif they all share (see Table 1 for the Avida instructions and their descriptions). This class contains 62 replicators (9 self-replicators). Only 8 of these replicators have a standard copy loop (see Section 2 for a definition of a copy loop), which appears at the end of their genomes; these copy loops were marked by the presence of a *mov-head* (*g*) instruction.²

Not having a copy loop, the remaining *hc* replicators loop through their entire genome in order to replicate, that is, they use the circular nature of their genome to achieve self-replication without a dedicated copy loop. In order to separate their read and write heads, replicators of this class move the entry of the AX register placed by the *b-alloc* (*w*) execution into the CX register through the execution of one or more instructions (in all but one case, one or two *swap* (*r*) executions). Then, they would execute *jmp-head* (*h*) to move the write head 15 positions ahead to allow for genome copying (the read-head location is usually between the first instruction and the third instruction in the genome).

We named the second type of replicators the *fg* replicator class (given the *fg* motif they share). The class consists of 12 replicators (11 self-replicators). All of these replicators have a copy loop at

² Although Avidian genomes are circular, they have a defined beginning instruction where genome execution starts.

Table 1. Instruction set of the Avidian programming language used in this study.

Instruction	Description	Symbol
nop-A	no operation (type A)	a
nop-B	no operation (type B)	b
nop-C	no operation (type C)	c
if-n-eq	Execute next instruction only-if ?BX? does not equal complement	d
if-less	Execute next instruction only if ?BX? is less than its complement	e
if-label	Execute next instruction only if template complement was just copied	f
mov-head	Move instruction pointer to same position as flow head	g
jmp-head	Move instruction pointer by fixed amount found in register CX	h
get-head	Write position of instruction pointer into register CX	i
set-flow	Move the flow head to the memory position specified by ?CX?	j
shift-r	Shift all the bits in ?BX? one to the right	k
shift-l	Shift all the bits in ?BX? one to the left	l
inc	Increment ?BX?	m
dec	Decrement ?BX?	n
push	Copy value of ?BX? onto top of current stack	o
pop	Remove number from current stack and place in ?BX?	p
swap-stk	Toggle the active stack	q
swap	Swap the contents of ?BX? with its complement	r
add	Calculate sum of BX and CX; put result in ?BX?	s
sub	Calculate BX minus CX; put result in ?BX?	t
nand	Perform bitwise NAND on BX and CX; put result in ?BX?	u
h-copy	Copy instruction from read head to write head and advance both	v
h-alloc	Allocate memory for offspring	w
h-divide	Divide off an offspring located between read head and write head	x
IO	Output value ?BX? and replace with new input	y
h-search	Find complement template and place flow head after it	z

Notes. The notation ?BX? implies that the command operates on a register specified by the subsequent nop instruction (for example, nop-A specifies the AX register, and so forth). If no nop instruction follows, use the register BX as a default. Table from [30]. More details about this instruction set can be found in [41].

the beginning of their genomes; the copy loop was marked by the presence of an *if-label* and *mov-bead* (fg) instruction motif. In order to separate their read and write heads, replicators of this class do not move the sequence length (here, the number 15) from the AX register to the CX register, as hc replicators do. Instead, they loop through their copy loop 14 or 15 times to set the location of both the read head and write head to 15. Next, hc replicators send the read head to the origin of replication (the instruction with address 0), usually via the *mov-bead* (g) instruction, although two replicators used *jmp-bead* (h). Then, the read and write heads are a sequence length number of units apart, and copying of the genome commences through the copy loop.

We also found one replicator that was a hybrid of these two classes. This replicator had a copy loop at the beginning of its genome that ended with an *if-label* and *mov-bead* (fg) instruction motif. However, that replicator also used *swap* (r) and *jmp-bead* (h) to move its write head to 15 in order to begin genome replication, so it uses elements of both classes. This sequence turns out to be a true self-replicator, not a proto-replicator. All replicator sequences from the three classes are listed in Table 2 in the Appendix.

3.3 Evolvability of Emergent Replicators

We performed experiments to explore the evolvability of these replicators. First, we tested how evolvable these emergent replicators are in terms of replication optimization (i.e., the optimization of the replication algorithm). The majority of replicators in any class had a fitness less than 0.1 offspring per unit of time, while, by comparison, the default Avida ancestor genotype has a fitness of approximately 0.25 offspring per unit of time (Figure 1). The replicators with the largest fitnesses were all from the hc class.

To test for optimization evolvability, we evolved ten populations seeded with each replicator for 1,000 generations in an environment where they could not evolve new phenotypic traits (i.e., Boolean logic calculations), that is, they could only improve the speed at which they replicated. The genotypes with the highest mean evolved fitness all were of the hc type (Figure 2). Meanwhile, many of the least evolvable replicators came from the fg-replicator class. These results suggest that the replication machinery of an Avidian determines how well it can optimize its replication speed, regardless of its starting fitness.




Figure 1. Ancestral fitness of each emergent replicator. Blue bars are hc replicators, while red bars are fg replicators. The one yellow bar represents the single hybrid replicator.

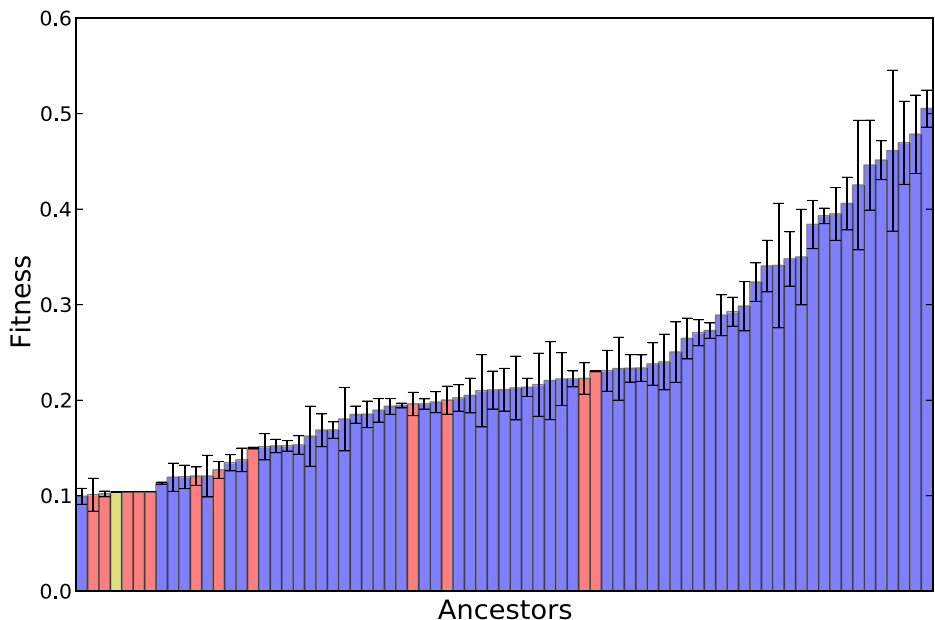


Figure 2. Mean evolved fitness of emergent replicators in the optimization experiments. Error bars are twice standard error. Blue bars are hc replicators, while red bars are fg replicators. The one yellow bar represents the single hybrid replicator.

Next, we performed experiments to test how an Avidian’s replication class affected its ability to evolve new phenotypic traits. We again evolved ten populations of each replicator, but here we selected for the ability of these Avidians to perform Boolean logic calculations (*traits*). We evolved these populations for 10^4 generations (as opposed to the 10^3 generations for replication optimization) to compensate for the possibility that trait evolution takes longer than optimization. Although most populations did not evolve any phenotypic traits within that time, the replicators that led to trait evolution were not evenly distributed across replication classes (Figure 3). We found that

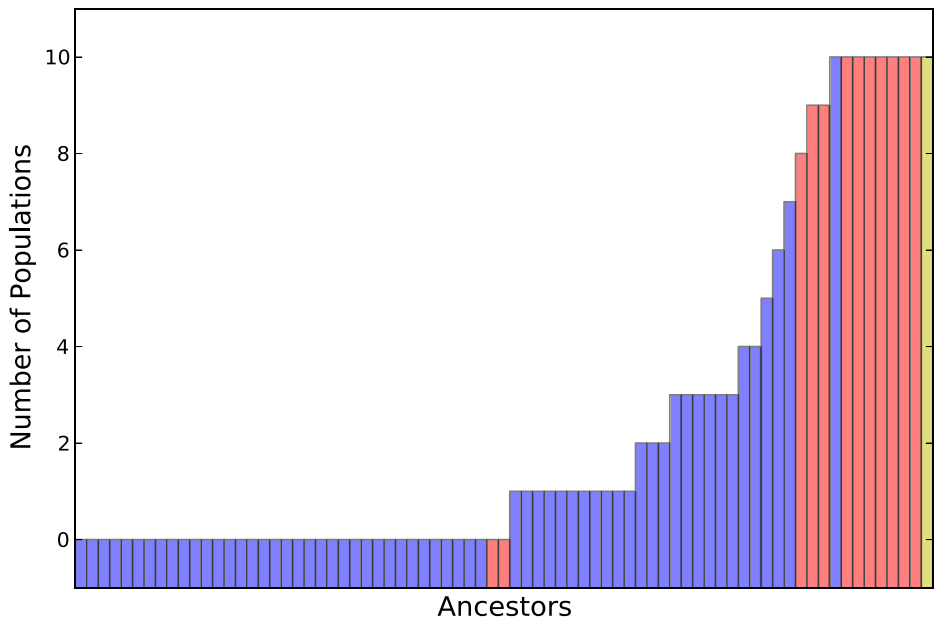


Figure 3. Number of populations (out of 10) that evolved any novel phenotypic traits in the innovation experiments for each emergent replicator. Color scheme as in Figure 2.

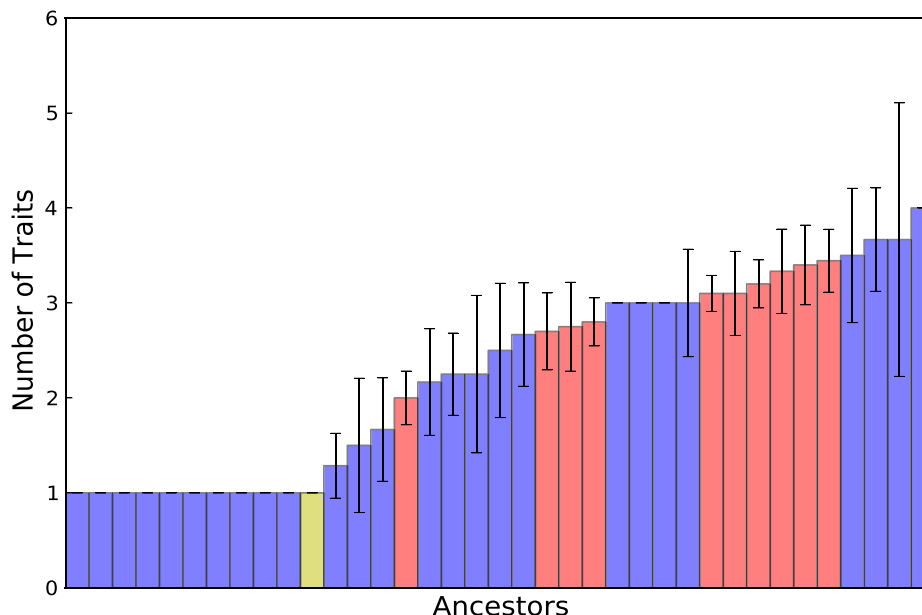


Figure 4. Mean number of evolved phenotypic traits in the innovation experiments. Only populations that evolved at least one trait were included in the data. Error bars are twice standard error. Color scheme as in Figure 2.

7/12 of the fg replicators evolved at least one trait in every replicate; only 2/12 fg replicators did not evolve traits in any experiment. In comparison, only one hc replicator evolved traits in every replicate, and 36 hc replicators never evolved any traits. However, the trend was less distinct when considering the number of traits evolved in replicates that evolved at least one trait (Figure 4). We found that some hc replicators could evolve as many traits as fg replicators.

3.4 Mechanism of Evolvability Tradeoff

The above results indicate that one class of replicators is better at optimizing replication ability (the hc-replication class) and the other is better at evolving innovations (the fg-replication class). This observation suggests a tradeoff between evolvability in the sense of phenotypic optimization, and evolvability in the sense of phenotypic innovation. To explore this further, we studied how the different replication mechanisms affect and determine evolvability. Specifically, we examined the respective algorithms the two classes use, and estimated the minimal number of instructions needed to replicate (as genotypes that execute fewer instructions have a greater fitness).

First, we explored the replication algorithm used by the fittest hc replicators. These replicators execute their entire genome (15 instructions). During this execution step, it will allocate its daughter genome, set the write head to 15 (the read head is set at 0 by default), and copy V instructions into the newly allocated daughter genome, where V is the number of *b-copy* (v) instructions in the genome. So far, 15 instructions have been executed and V instructions have been copied. The Avidian needs to copy a further $15 - V$ instructions and divide off the daughter genome to reproduce. During this first genome execution, the last instruction executed is a *mov-head* (g), which moves the instruction pointer back to the beginning of that Avidian's copy loop. The Avidian will then loop through the copy loop until its entire genome is copied into its daughter genome. As it needs to copy a further $15 - V$ instructions (and it can copy V instructions per traversal), it will loop through $\frac{15-V}{V}$ times. During each traversal through the copy loop, L instructions are executed, where L is the size of the copy loop. Thus, to replicate the rest of genome, $\frac{(15-V)L}{V}$ instructions are executed. Therefore, hc replicators execute $15 + \frac{(15-V)L}{V} = 15 - L + 15\frac{L}{V}$ instructions to reproduce.

Next, we will calculate the minimal number of instruction executions needed to replicate an fg replicator. Unlike the hc type with a copy loop at the end of their genomes, fg replicators have a

copy loop at the beginning of their genomes. Denote this copy loop size again by L , and the number of *b-copy* (v) instructions by V . First, fg replicators will loop through their copy loop 15 times, executing $15L$ instructions during this step. During this looping, the fittest replicators will have allocated their daughter genomes. This repeated looping process sets the location of both the read and the write heads to 15. After the last copy execution is complete, the rest of the genome is executed, requiring $15 - L$ instruction executions. During this step, the read-head location is set to 0. Next, the instruction pointer is moved back to the beginning of the genome to begin genome copying. As there are V *b-copy* (v) instructions in the copy loop, the copy loop will need to be traversed $\frac{15}{V}$ times. Because the copy loop is L instructions long, this step requires $\frac{15L}{V}$ instruction executions. Therefore, in total, fg replicators require $15L + 15 - L + 15\frac{L}{V}$ instruction executions to reproduce.

By comparing the numbers of instruction executions required for replication by the two classes, we can see that fg replicators require $15L$ more instruction executions when the copy loop size and the number of *b-copy* instructions are fixed. We note here that this does not mean that fg replicators have a greater mutation rate than hc replicators, as mutations are applied during division, not genome copying. One other factor contributes to the optimization advantage of hc replicators. The copy loop of fg replicators is located at the beginning of their genomes. These copy loops cannot easily expand, as they are bounded in the genome by the fg instruction motif; expansion requires at least two mutations. However, hc replicators have a copy loop at the end of their genome; this copy loop can more easily expand. While a copy loop expansion with a random instruction would increase the number of executions required for replication, a copy loop expansion with an additional *b-copy* instruction would decrease the number of times the copy loop needs to be traversed. This would decrease the number of instruction executions needed for reproduction and increase the fitness of an hc replicator.

Finally, we explored what prevented hc replicators from evolving any novel traits in many populations, for example whether there is an optimization–innovation tradeoff. Replicators from the hc class tend to fix beneficial mutations that enhance their replication ability. These mutations accumulate in regions of the genome that subsequently cannot be mutated into the type of instructions that discover novel traits. In turn, fg replicators do not optimize their replication ability as well as hc replicators, and therefore they do not (in a sense prematurely) fix as many beneficial mutations to

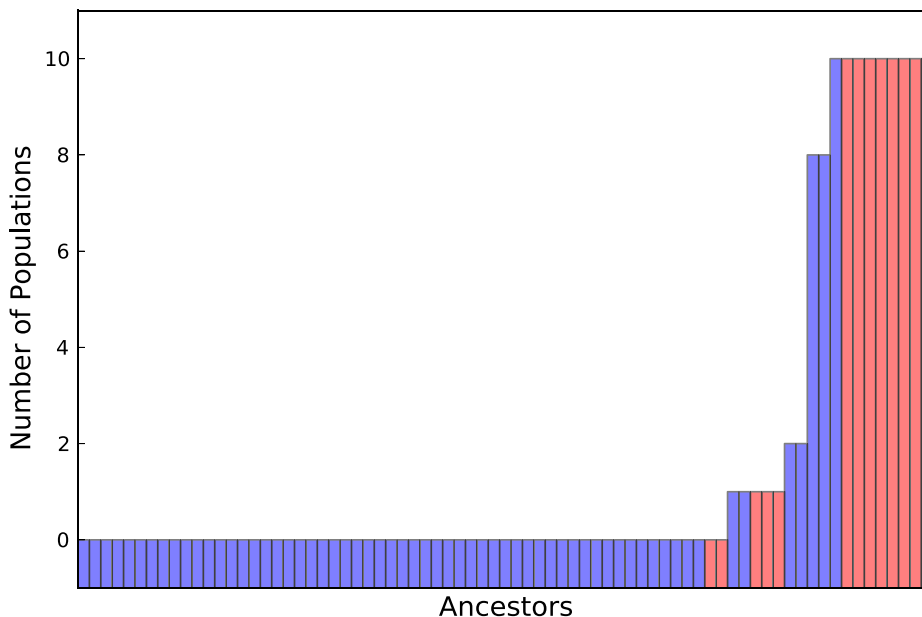


Figure 5. Number of populations (out of 10) that evolved any novel phenotypic traits in the tradeoff experiments for each ancestral replicator. Color scheme as in Figure 2.

optimize the replication machinery. The failure to do so appears to have opened the opportunity for those sequences to instead neutrally accumulate mutations that would later lead to trait evolution.

To test this idea, we took the fittest descendant of each ancestral replicator from the optimization experiments and evolved them in an environment where phenotypic traits were under selection for 10^4 generations (Figure 5). These experiments are the same as the innovation experiments, except we used evolved ancestors instead of the emergent replicators. Ten fg replicators and the hybrid replicator evolved novel traits in at least one population, and many of these replicators evolved traits in all ten populations. There is little difference in the evolution of traits in the fg replicators between these experiments and the original innovation experiments (Figure 3). However, the rate of trait evolution decreased in the hc replicators between these experiments and the original innovation experiments. For the hc replicators, 26/62 evolved traits in at least one population in the original innovation experiments; only 7/62 replicators evolved traits in these new innovation-after-optimization experiments. The decrease of trait evolution between the original innovation experiments and these experiments in the hc replicators (but not the fg replicators) indicates that their ability to optimize their replication prevents the hc replicators from later evolving novel phenotypic traits.

4 Discussion

We explored the evolvability of a set of emergent (i.e., found through random search), fixed-length, self-replicating sequences in the Avida digital life world. We found two distinct classes of replicators that differ in their general (algorithmic) strategy for self-replication. We found that one of these classes displayed an enhanced evolvability in terms of replication optimization, while the other was more evolvable in terms of phenotypic innovation. The experiments performed here complement our previous work on the interplay between evolvability and self-replication, where we found that most *variable-length* replicators were evolvable both in optimization and in innovation [30]. Using the present (distinct) set of replicators, we were able to provide a functional explanation for the differences of evolvability in emergent Avida replicators.

What consequences do these results have for the origin of life and its subsequent evolution? There is some evidence to suggest that an optimization–innovation tradeoff may have existed in the earliest RNA-based replicators. In experiments with RNA bacteriophage Q β , Mills, Peterson, and Spiegelman showed that this virus would shrink to a minimal genome and remove the genes essential for a viral lifestyle in order to optimize replication speed [38]. This suggests—although it certainly does not prove—that a selection pressure to minimize genome size may have been part of the RNA self-replication fitness landscape. When genome space is removed by selection, the ability for phenotypic innovation to occur is impeded as the remaining genomic space may become too small to contain the genes encoding novel traits.

We believe that the *potential* for an optimization–innovation tradeoff that hinders the evolution of phenotypic novelty may be universal. However, we also know that the earliest replicators on Earth, whether RNA-based replicators or other biochemical organisms, could eventually evolve phenotypic innovations. Furthermore, Voytek and Joyce showed that two RNA enzymes could adapt to distinct ecological niches in continuous evolution experiments, supporting the idea that novel phenotypic traits could evolve in RNA replicators [56]. In fact, a similar tradeoff was also observed in experiments with Avidian replicators (using a single hand-written ancestor) [61], which showed that evolution could take two different paths: one that favored fast replication, and another that traded replication speed for the exploitation of complex phenotypes.

The optimization–innovation tradeoff we have argued for here can also be interpreted as a tradeoff for evolvability in two separate fitness components. In a previous article, this tradeoff between replication rate and phenotypic complexity was compared to the well-known tradeoff between r and K reproduction strategies [61]. However, the tradeoff between fitness components in these replicators is not a fitness tradeoff, as classically discussed [52, 49, 23, 50]. Instead, the tradeoff is one in terms of evolvability between two fitness components. Each replicator class is biased in terms

of the generation of genetic variation, and this bias determines which fitness component is improved by each replicator class. This result complements the previous observation in *Avida*, which argued that historical contingency influenced which fitness component a population improved [61]. The present results identify another factor, namely that genetic background alone may determine a population's evolutionary trajectory towards a specific ecological strategy.

One factor thought to influence evolvability is mutational robustness and the presence of neutral mutations [58, 15, 34]. The presence of neutral mutations results in neutral networks throughout the fitness landscape, allowing populations access to novel phenotypes [57]. We found a similar result here. Replicators from the fg class fixed fewer beneficial mutations that optimize replication ability. As a consequence, more of their genome's instructions were neutral, and this in turn allowed them to accumulate many mutations that eventually led to the evolution of complex traits.

What role did the presence or absence of neutral networks play in the origin of life? Early work on (computational) RNA secondary structure predicted the existence of extensive neutral networks in RNA landscapes [25]. However, more recent experimental work suggests instead that the empirical fitness landscape of short RNAs contains isolated fitness peaks [26, 46]. It is possible that different emergent replicators could either be stuck on an isolated fitness peak (such as the hc replicators) or else have access to a neutral network that leads to phenotypic innovation (such as the fg replicators). Such a fitness landscape architecture would emphasize a strong role for historical contingency in the origin and subsequent evolution of life.

Finally, in the work described here, Avidians obtain the "energy" necessary to replicate by performing logic calculations on random binary numbers (the "tasks" metabolism), but the randomly generated replicators do not have any such metabolic instructions. This is possible because an early design decision gave each Avidian a certain amount of "base" energy for free, which they can augment by performing logic calculations. We can ask whether requiring that replicators possess some basic metabolic process (such as, for example, NOT, the simplest of the logic operations) would change the conclusions we arrived at here.

Requiring a minimal metabolism would certainly increase the minimal amount of information a replicator must have from the $I = -\log_{26}(75/10^9) \approx 5.23$ instructions [2] by a few instructions, making the emergence of a replicator by chance that much more rare. We also would expect that the metabolic instructions would be integrated with the replicator genes, which could affect the ability of a replicator to optimize its replication machinery without affecting the metabolism. On the other hand, having a seed metabolism present could also enhance the evolution of other logic operations, for example by code duplication. While a detailed investigation of these questions must await further work, an origin-of-life scenario without metabolic genes is not altogether far-fetched. For example, the Lost City hydrothermal vents located near the summit of the Atlantis Massif [11, 10] on the mid-Atlantic seafloor produces hydrogen and simple carbohydrates "for free," which could be utilized by the simplest primordial replicators without the need for explicit genes. Such replicators could then evolve the genes for autonomous metabolism later on, just as the Avidians do in the experiments described here.

Acknowledgments

Computational work in support of this research was performed at Michigan State University's High Performance Computing Facility. This material is based in part upon work supported by the National Science Foundation under Cooperative Agreement No. DBI-0939454. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

1. Adami, C. (1998). *Introduction to artificial life*. New York: TELOS Springer-Verlag.
2. Adami, C. (2015). Information-theoretic considerations concerning the origin of life. *Origins of Life and Evolution of Biospheres*, 45, 9439.

3. Adami, C., & LaBar, T. (2017). From entropy to information: Biased typewriters and the origin of life. In S. I. Walker, P. Davies, & G. Ellis (Eds.), *Information and causality: From matter to life* (pp. 95–105). Cambridge, MA: Cambridge University Press.
4. Adami, C., Ofria, C., & Collier, T. C. (2000). Evolution of biological complexity. *Proceedings of the National Academy of Sciences of the U.S.A.*, *97*(9), 4463–4468.
5. AnceL, L. W., & Fontana, W. (2000). Plasticity, evolvability, and modularity in RNA. *Journal of Experimental Zoology*, *288*(3), 242–283.
6. Athavale, S. S., Spicer, B., & Chen, I. A. (2014). Experimental fitness landscapes to understand the molecular evolution of RNA-based life. *Current Opinion in Chemical Biology*, *22*, 35–39.
7. Bedau, M. A., & Packard, N. H. (2003). Evolution of evolvability via adaptation of mutation rates. *Biosystems*, *69*(2), 143–162.
8. Biebricher, C. K., Eigen, M., & Luce, R. (1986). Template-free RNA synthesis by Q β replicase. *Nature*, *321*(6065), 89–91.
9. Biebricher, C. K., & Luce, R. (1993). Sequence analysis of RNA species synthesized by Q β replicase without template. *Biochemistry*, *32*(18), 4848–4854.
10. Brazelton, W. J., Nelson, B., & Schrenk, M. O. (2011). Metagenomic evidence for H₂ oxidation and H₂ production by serpentinite-hosted subsurface microbial communities. *Frontiers in Microbiology*, *2*, 268.
11. Brazelton, W. J., Schrenk, M. O., Kelley, D. S., & Baross, J. A. (2006). Methane- and sulfur-metabolizing microbial communities dominate the Lost City hydrothermal field ecosystem. *Applied and Environmental Microbiology*, *72*(9), 6257–6270.
12. Clune, J., Misevic, D., Ofria, C., Lenski, R. E., Elena, S. F., & Sanjuán, R. (2008). Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. *PLoS Computational Biology*, *4*(9), e1000187.
13. Covert, A. W., Lenski, R. E., Wilke, C. O., & Ofria, C. (2013). Experiments on the role of deleterious mutations as stepping stones in adaptive evolution. *Proceedings of the National Academy of Sciences of the U.S.A.*, *110*(34), E3171–E3178.
14. Decker, P., Schweer, H., & Pohlmann, R. (1982). Bioids: X. Identification of formose sugars, presumable prebiotic metabolites, using capillary gas-chromatography/gas chromatography–mass spectrometry of *n*-butoxime trifluoroacetates on OV-225. *Journal of Chromatography*, *244*, 281–291.
15. Draghi, J. A., Parsons, T. L., Wagner, G. P., & Plotkin, J. B. (2010). Mutational robustness can facilitate adaptation. *Nature*, *463*(7279), 353–355.
16. Earl, D. J., & Deem, M. W. (2004). Evolvability is a selectable trait. *Proceedings of the National Academy of Sciences of the U.S.A.*, *101*(32), 11531–11536.
17. Egri-Nagy, A., & Nehaniv, C. L. (2003). Evolvability of the genotype-phenotype relation in populations of self-replicating digital organisms in a Tierra-like system. In *Advances in Artificial Life* (pp. 238–247). Berlin Heidelberg: Springer.
18. Elena, S. F., & Sanjuán, R. (2008). The effect of genetic robustness on evolvability in digital organisms. *BMC Evolutionary Biology*, *8*(1), 284.
19. England, J. L. (2013). Statistical physics of self-replication. *Journal of Chemical Physics*, *139*(12), 121923.
20. Frénoy, A., Taddei, F., & Misevic, D. (2012). Robustness and evolvability of cooperation. In C. Adami, D. M. Bryson, C. Ofria, & R. T. Pennock (Eds.), *Proceedings of Artificial Life 13* (pp. 53–58). Cambridge, MA: MIT Press.
21. Gilbert, W. (1986). Origin of life: The RNA world. *Nature*, *319*(6055), 618.
22. Goldsby, H. J., Knoester, D. B., Ofria, C., & Kerr, B. (2014). The evolutionary origin of somatic cells under the dirty work hypothesis. *PLoS Biology*, *12*(5), e1001858.
23. Hereford, J. (2009). A quantitative survey of local adaptation and fitness trade-offs. *The American Naturalist*, *173*(5), 579–588.
24. Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science and Engineering*, *9*(3), 90–95.
25. Huynen, M. A., Stadler, P. F., & Fontana, W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation. *Proceedings of the National Academy of Sciences of the U.S.A.*, *93*(1), 397–401.

26. Jiménez, J. I., Xulvi-Brunet, R., Campbell, G. W., Turk-MacLeod, R., & Chen, I. A. (2013). Comprehensive experimental fitness landscape & evolutionary network for small RNA. *Proceedings of the National Academy of Sciences of the U.S.A.*, *110*(37), 14984–14989.
27. Jones, E., Oliphant, T., Peterson, P., et al. (2001). *SciPy: Open source scientific tools for Python*. <http://www.scipy.org>.
28. Joyce, G. F. (2002). The antiquity of RNA-based evolution. *Nature*, *418*(6894), 214–221.
29. Kirschner, M., & Gerhart, J. (1998). Evolvability. *Proceedings of the National Academy of Sciences of the U.S.A.*, *95*(15), 8420–8427.
30. LaBar, T., Adami, C., & Hintze, A. (2015). Does self-replication imply evolvability? In *Proceedings of the European Conference on Artificial Life 2015* (pp. 595–602). Cambridge, MA: MIT Press.
31. Lenski, R. E., Ofria, C., Collier, T. C., & Adami, C. (1999). Genome complexity, robustness and genetic interactions in digital organisms. *Nature*, *400*(6745), 661–664.
32. Lenski, R. E., Ofria, C., Pennock, R. T., & Adami, C. (2003). The evolutionary origin of complex features. *Nature*, *423*(6936), 139–144.
33. Lincoln, T. A., & Joyce, G. F. (2009). Self-sustained replication of an RNA enzyme. *Science*, *323*(5918), 1229–1232.
34. Masel, J., & Trotter, M. V. (2010). Robustness and evolvability. *Trends in Genetics*, *26*(9), 406–414.
35. Mathis, C., Bhattacharya, T., & Walker, S. I. (2015). The emergence of life as a first order phase transition. Preprint arXiv:1503.02776.
36. McKinney, W. (2010). Data structures for statistical computing in Python. In *Proceedings of the 9th Python in Science Conference* (pp. 51–56).
37. Mileikowsky, C., Cucinotta, F. A., Wilson, J. W., Gladman, B., Horneck, G., Lindegren, L., Melosh, J., Rickman, H., Valtonen, M., & Zheng, J. Q. (2000). Natural transfer of viable microbes in space: 1. From Mars to Earth and Earth to Mars. *Icarus*, *145*(2), 391–427.
38. Mills, D. R., Peterson, R. L., & Spiegelman, S. (1967). An extracellular Darwinian experiment with a self-duplicating nucleic acid molecule. *Proceedings of the National Academy of Sciences of the U.S.A.*, *58*(1), 217.
39. Misevic, D., Ofria, C., & Lenski, R. E. (2006). Sexual reproduction reshapes the genetic architecture of digital organisms. *Proceedings of the Royal Society of London B*, *273*(1585), 457–464.
40. Neveu, M., Kim, H.-J., & Benner, S. A. (2013). The “strong” RNA world hypothesis: Fifty years old. *Astrobiology*, *13*(4), 391–403.
41. Ofria, C., Bryson, D. M., & Wilke, C. O. (2009). Avida: A software platform for research in computational evolutionary biology. In A. Adamatzky & M. Komosinski (Eds.), *Artificial life models in software* (pp. 3–35). London: Springer.
42. Orgel, L. E. (2004). Prebiotic chemistry and the origin of the RNA world. *Critical Reviews in Biochemistry and Molecular Biology*, *39*(2), 99–123.
43. Pargellis, A. N. (1996). The spontaneous generation of digital life. *Physica D: Nonlinear Phenomena*, *91*(1), 86–96.
44. Pargellis, A. N. (2001). Digital life behavior in the amoeba world. *Artificial Life*, *7*(1), 63–75.
45. Pennock, R. T. (2007). Models, simulations, instantiations, and evidence: The case of digital evolution. *Journal of Experimental & Theoretical Artificial Intelligence*, *19*(1), 29–42.
46. Petrie, K. L., & Joyce, G. F. (2014). Limits of neutral drift: Lessons from the in vitro evolution of two ribozymes. *Journal of Molecular Evolution*, *79*(3–4), 75–90.
47. Robertson, M. P., & Joyce, G. F. (2012). The origins of the RNA world. *Cold Spring Harbor Perspectives in Biology*, *4*(5), a003608.
48. Robertson, M. P., & Joyce, G. F. (2014). Highly efficient self-replicating RNA enzymes. *Chemistry & Biology*, *21*(2), 238–245.
49. Roff, D. A., & Fairbairn, D. J. (2007). The evolution of trade-offs: Where are we? *Journal of Evolutionary Biology*, *20*(2), 433–447.
50. Shoval, O., Sheftel, H., Shinar, G., Hart, Y., Ramote, O., Mayo, A., Dekel, E., Kavanagh, K., & Alon, U. (2012). Evolutionary trade-offs, Pareto optimality, and the geometry of phenotype space. *Science*, *336*(6085), 1157–1160.

51. Smith, T., Husbands, P., & O'Shea, M. (2002). Fitness landscapes and evolvability. *Evolutionary Computation*, 10(1), 1–34.
52. Stearns, S. C. (1989). Trade-offs in life-history evolution. *Functional Ecology*, 3(3), 259–268.
53. Sumper, M., & Luce, R. (1975). Evidence for de novo production of self-replicating and environmentally adapted RNA structures by bacteriophage Q β replicase. *Proceedings of the National Academy of Sciences of the U.S.A.*, 72(1), 162–166.
54. Taylor, T. (1999). On self-reproduction and evolvability. In *Advances in Artificial Life* (pp. 94–103). Berlin Heidelberg: Springer.
55. Van Der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The NumPy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2), 22–30.
56. Voytek, S. B., & Joyce, G. F. (2009). Niche partitioning in the coevolution of 2 distinct RNA enzymes. *Proceedings of the National Academy of Sciences of the U.S.A.*, 106(19), 7780–7785.
57. Wagner, A. (2011). *The origins of evolutionary innovations: A theory of transformative change in living systems*. Oxford: Oxford University Press.
58. Wagner, A. (2013). *Robustness and evolvability in living systems*. Princeton, NJ: Princeton University Press.
59. Wagner, G. P., & Altenberg, L. (1996). Perspective: Complex adaptations and the evolution of evolvability. *Evolution*, 50(3), 967–976.
60. Walker, S. I., & Davies, P. C. W. (2013). The algorithmic origins of life. *Journal of the Royal Society Interface*, 10(79), 20120869.
61. White, J. S., & Adami, C. (2004). Bifurcation into functional niches in adaptation. *Artificial Life*, 10(2), 135–144.
62. Zaman, L., Meyer, J. R., Devangam, S., Bryson, D. M., Lenski, R. E., & Ofria, C. (2014). Coevolution drives the emergence of complex traits and promotes evolvability. *PLoS Biology*, 12(12), e1002023.

Appendix: Sequences of Replicators

See Table 2.

Table 2. Sequences of fixed-length replicators organized by replication class.

hc Genomes			
nuvhcwzscrchrwx	wdphcrxswvprcff	cxrphcwravgmspm	dwtwrvhcbxivzrr
wmjrcshcikcivxg	ahcqdqvwzjrcxet	smvqztrchcwxxwj	vvtxwfiwvzrchch
exwakshcvjjrckp	zezvkrwcfsshclx	dzzkzpvrcbwoxhc	kpcchcpavbcrxwo
zwrchceekovxyzi	rwvwtzubrcxkthc	mzvjwmwrcuxhcng	jvwpwhulxdarchc
pwzrvxqaarochcl	upathcvuwwxkrc	lxvwowxhcrcqnml	lbrcwavhcwaahwx
yhcxxclwveblzrc	whcwemcxrchoxp	qhcCowvrptjlbr	nxzjvwzxsrchcjl
axzrchctvwijsca	swvrhcxbllrpdir	lvlkwovtrcxphcp	xshcvxwrcdetbck
ynprcvxhcdfwlcj	ezvxwprchcdkqc	bwhcpvqzrcfhdtx	awexwcbvrvbxhcr
qsozvdretcrwxhc	xlczararhchvwqx	izvrcqhczzbawdw	oovwxnhrcqhclmy
ojrchcwxfwgczi	fjyrrcbwhcixvbg	kpavwrchcnbxupb	wrchcsndlzfrvxg

Table 2. (continued).

udptwrcxvnwvhca	ozpvoabkwxcrhc	unhcwklcvzrcxjg	wwrcvohckqcxhvp
xhwfkmrcdjhcixv	yzpftzvrcdthcw	frmvwzkrcv xvjhc	slcrvhcwsslxzra
aarchcekvwxtpeg	wifhorvthcxppwr	awmmqlvxhmmrhc	lbvxttawrcvlahc
szhcwrcwvxhcfwa	lvxhtwjmwwrchr	mphcwvadiqslxrc	lkexjvozxrhcw
cpvzrcqxjhcwtdj	bfqacswwrcvxhcz	peenvzwrchcxjgw	pshcwrcpvfxqxld
eazhwrcxhcvodsh	iklkvwozljrchcx		
fg Genomes			
vfgxqhmwmfjphgb	wvshxfgkfooxugb	kvxfgwfyxukujgb	dywqvphfguxqdgb
vmfifgwvpowxgbt	aytvvmwxkfgohbi	xnkwovyfgtxwqgb	vnfgudsftwxwhgb
vlfgvmhwxwlgbq	cxvdfgepkhbtshi	wvfghtzoxjirgb	wvfgxdmoprllwgb
Hybrid Genome			
vmfgqwrmdkyxuhc			
