

Potentiating Mutations Facilitate the Evolution of Associative Learning in Digital Organisms

Austin J. Ferguson^{1,2,3,*}, Charles Ofria^{1,2,3}

¹Department of Computer Science and Engineering ²Program in Ecology, Evolution, and Biology

³BEACON Center for the Study of Evolution in Action

Michigan State University, East Lansing, MI 48824

*fergu358@msu.edu

Abstract

Scientists have long tried to predict evolutionary outcomes in order to design vaccines for next year's diseases, stabilize endangered ecosystems, or make better choices in designing evolutionary algorithms. To predict, however, we must first be able to retroactively identify the key steps that determined the evolved state. Researchers have long examined the role of historical contingency in evolution; when do small, seemingly insignificant mutations substantially shift the probabilities of what traits or behaviors ultimately evolve? Practitioners of experimental evolution have recently begun to investigate this question using a new technique: analytic replay experiments. We can found many populations with a given genotype in order to measure the probability of a particular trait evolving from that starting point; we call this the "potentiation" of that genotype. Moving along a lineage, we can identify which mutations altered potentiation. Here we used digital organisms to conduct a high-resolution analysis of how individual mutations affected the potentiation of associative learning. We find that the probability of evolving associative learning can increase suddenly – even with a single mutation that appeared innocuous when it occurred. While there was no obvious signal to identify potentiating mutations as they arose, we were able to retrospectively identify mechanisms by which these mutations influenced subsequent evolution. Many of the most interesting and complex evolutionary adaptations that occur in nature are exceptionally rare. Here, we extend techniques for understanding these rare evolutionary events and the patterns and processes that produce them.

Introduction

How likely is the evolution of a particular trait? Researchers have long been interested in predicting evolutionary outcomes, but the inherent stochasticity in the process makes this goal exceptionally challenging. In order to make more accurate predictions, we would need to better understand how and why the underlying probabilities of potential outcomes change over time. Looking purely retrospectively at evolution in nature, this type of analysis is not possible (at least not without a time machine). Leveraging the flexibility and controls available in experimental evolution, however, allows us to empirically test questions that were previously only hypothetical (Kawecki et al., 2012). Here, we focus on Stephen Jay Gould's idea of "replaying the tape of life"

(Gould, 1990). The idea is simple: If we were to start life over again from the same initial conditions, would evolution follow the same pathway? Alas, Gould remarked that this experiment is unfortunately impossible.

While it may be impossible to replay the *entire* tape of life, practitioners of experimental evolution have conducted this experiment on a smaller scale. Travisano et al. (1995), Wagenaar and Adami (2004), and Blount et al. (2008) introduced and refined methods of investigating the role of historical contingency in evolving populations: parallel and analytic replay experiments. By evolving multiple populations from the same starting organisms, researchers can identify the range and distribution of outcomes. These populations can be evolved simultaneously (parallel replays), however many microbial and digital populations allow us to preserve a "fossil record", opening up another possibility. Analytic replay experiments systematically revive historical populations to re-evolve them, allowing researchers to identify alternative possibilities after the fact (Blount et al., 2018). When one strain of *E. coli* in Dr. Richard Lenski's long-term evolution experiment (Lenski et al., 1991) unexpectedly evolved the ability to digest citrate, Blount et al. (2008) used analytic replay techniques on previously frozen samples (spaced across the lineage) to identify the potentiation of this unlikely evolutionary outcome. In their replay experiments, restarts from earlier time points never re-evolved citrate utilization, but successful re-evolution of the behavior in restarts from later time points indicated that the population had become potentiated. In later work, Blount et al. (2012) used genetic sequencing and manipulation to identify the specific potentiating mutations associated with this increased probability.

Analytic replay experiments provide a powerful new tool for understanding the role of history in evolution. In addition to studying the evolution of *E. coli* citrate metabolism, analytic replay experiments have also been used to study the evolution of novel receptor usage of Phage λ into *E. coli* (Meyer et al., 2012), and colistin resistance in *Pseudomonas aeruginosa* (Jochumsen et al., 2016). For a review of these experiments and other uses of analytic replay experiments, see (Blount et al., 2018).

In this work we use digital evolution, specifically the evolution of self-replicating computer programs in the Avida Digital Evolution Platform (Ofria and Wilke, 2004), which has previously been used to conduct replay experiments. Yedid et al. (2008) employed this technique to investigate the re-evolution of traits following an extinction episode, while Covert III et al. (2013) used analytic replay experiments to study the importance of individual deleterious mutations in the evolution of complex traits.

We selected associative learning as a complex behavior to study potentiation. Associative learning is a non-trivial capability exhibited by most complex organisms. In digital evolution systems like Avida, it serves as a rare yet evolvable trait (Pontes et al., 2020). For an Avida organism to exhibit associative learning, it must be capable of sensing its environment, taking action, and storing information in memory. The evolution of associative learning has been studied via experimental evolution in both digital (Pontes et al., 2020; McGregor et al., 2012) and natural systems (Dunlap and Stephens, 2014; Mery and Kawecki, 2002), yet many questions remain about how it evolves. While more complex forms of learning are found in nature, associative learning remains an important building block for most others and insights about how it arises may be informative for understanding the broader evolution of intelligence, especially within digital contexts.

In this work, we begin to analyze how the likelihood of evolving a complex trait changes along a successful lineage. Using analytic replay experiments, we identified individual mutations that cause drastic increases in the potentiation of associative learning. We then analyzed those mutations and their mutational neighborhoods to begin characterizing how a mutation is potentiating. While these replay experiments are informative and useful for exploring counterfactual evolutionary possibilities, they are also computationally intensive. As such, we start by focusing on a set of case-study lineages to develop an initial framework for understanding how potentiation can occur.

Analyzing four successful lineages, we find that potentiation can increase suddenly, even due to a single mutation. Since these lineages were selected because they successfully evolved associative learning, potentiation generally increases in each, though some decreases do occur. Potentiating mutations vary in initial effect, making them challenging to detect. Retrospective analysis allows us to identify them, however, and begin hypothesizing about the dynamics that allow these mutations to potentiate associative learning. This work demonstrates using analytic replay experiments for quantifying potentiation along a lineage and establishes baselines and techniques for future studies.

Methods

Here we describe the digital evolution system and experiment setup used to conduct this work.

The Avida Digital Evolution Platform

This work uses an early build of version 5.0 of the Avida Digital Evolution Platform (Ofria and Wilke, 2004), currently under development as part of the Modular Agent Based Evolver 2 (MABE2) framework (<https://github.com/mercere99/MABE2>). In Avida, populations of self-replicating computer programs perform tasks to compete for CPU cycles, creating an evolution testbed that can support a wide array of experimental controls. Avida has been used for numerous studies on the evolution of complexity (Lenski et al., 2003; Zaman et al., 2014), associative learning (Pontes et al., 2020; Grabowski et al., 2010), and historical contingency via replay experiments (Yedid et al., 2008; Covert III et al., 2013). Fundamentally, Avida is designed to have tools necessary to conduct work at the scale required for replay experiments. While Avida is more complex, and thus slower, than other digital evolution models, we argue that this is appropriate for initial measurements of potentiation dynamics. Future work can isolate which of these dynamics are explained with simpler systems and which require more complicated interactions.

Avida genomes consist of assembly-like instructions that transfer data between registers, make basic comparisons, perform mathematical operations, *etc.* We use an extended instruction set that includes extra flow control and environment-specific instructions (Ferguson, 2023).

We used Avida populations on a 60x60 toroidal grid, resulting in a population cap of 3,600 organisms. Offspring are placed in a grid cell next to their parent, overwriting any existing organism in that cell; the parent organism is also reset. During reproduction, point mutations occur in offspring at a rate of 0.0075 per instruction, while single-instruction insertion and deletion mutations occur at a rate of 0.05 per reproduction. Organisms reproduce by executing the `Repro` instruction. To prevent organisms from immediately replicating, organisms must execute 1,500 instructions before the `Repro` instruction can be activated.

Associative learning To test the evolution of associative learning, we created a simplified version of the Avida path following environment (Pontes et al., 2020). Instead of navigating 2D space, an organism exists in one of four states: *forward*, *left*, *right*, or the error state, *backward*. Organisms are given a `Sense` instruction, which will give them the nutrient cue of their current state. Using this nutrient cue, organisms need to execute the appropriate instruction (one per state) to progress along the path. The *forward* and *backward* states have fixed cues (0 and -1, respectively), while, at birth, each organism is assigned random cue values for *left* and *right* in the range of $[1, 10^6]$. Organisms can genetically encode *forward* and *backward*, but must learn *left* or *right* in their lifetime to perfectly solve the task. Each path begins with one of four preset starting state sequences, chosen randomly for each organism at birth, followed by additional

random states. The four preset paths are the “one-fixed turn” paths from (Pontes et al., 2020), where organisms are guaranteed to encounter a *left* state before a *right*.

If the organism is not in an error state and executes the appropriate instruction, they are rewarded and move to the next state. If an organism executes the wrong instruction (e.g., the *Left* instruction in the *right* state), it is penalized and placed in the error state. While in the error state, the organism must execute the *Backward* instruction to return to the previous state and be allowed to try again; it will be penalized for any other action. A cooldown is applied, however, such that executing the *Backward* instruction causes the organism to wait for the equivalent of 10 additional instruction executions. Organisms are scored based on the number of valid states they successfully traversed minus the number of incorrect moves made, with a maximum score of 300. Fitness is calculated as 1.25^{score} , so each additional correct movement grants a 25% boost in fitness regardless of the total number of correct movements.

In this environment, optimal behavior requires associative learning in the form of imprinting. Since the paths are guaranteed to have a *left* state before a *right* state, the optimal behavior is to find and store the first positive cue value as the *left* cue. Combined with genetically-encoded *forward* and *backward* logic, storing and using the *left* cue is enough for organisms to identify the *right* cue through a process of elimination. Other possible behaviors involve error correction (assuming all turns are one direction, then correcting when wrong), bet-hedged learning (assuming more about the paths, e.g., that there are no instances of two lefts in a row), and various mixed strategies.

To categorize the behavior of a genotype, we evaluate it in 100 trials to ensure we observe how it performs in all four environments with different random cues. We then classify each of the 100 trials. Trials are classified as learning if the organism correctly handles greater than 90% of the states they were in, error correction if they always successfully navigated one turn state but not the other, and “low activity” if they failed to successfully navigate at least 25 states. To be categorized as learning or error correction, *all* 100 trials of that genotype must be of that class. If one or more trials were low activity, the genotype was categorized as “bet-hedged learning” or “bet-hedged error correction”. If a genotype displayed at least one learning trial and at least one error correction trial, they were classified as “mixed bet hedging”. Finally, all remaining genotypes were categorized as “low activity”. This categorization system was used across all three phases of this work.

Experiment framework

To identify mutations that substantially increased the likelihood that learning evolved, we split the work into three phases (see Figure 1 for an overview). First, we seeded 200 initial parallel replicates in the associative learning environ-

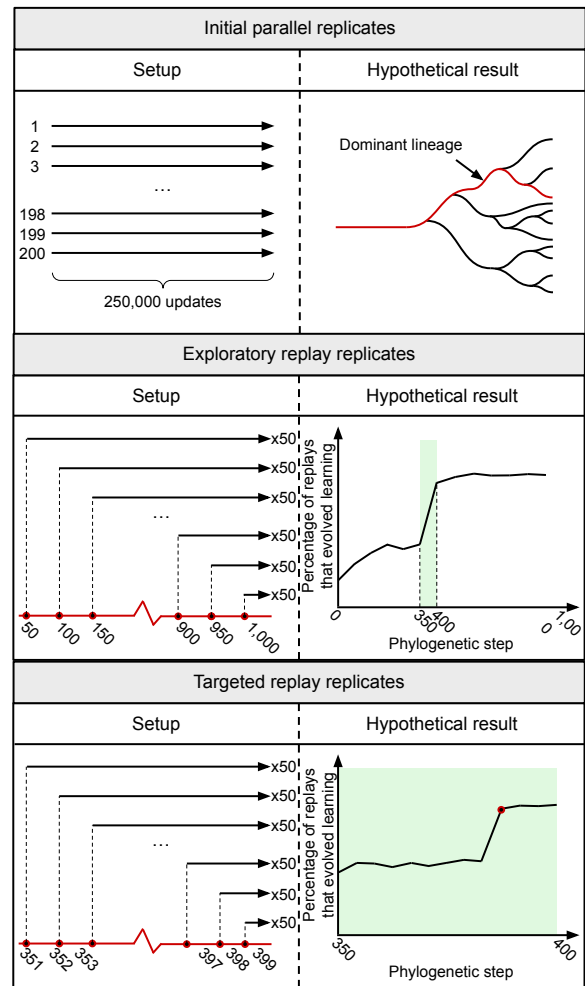


Figure 1: Illustration of the experimental design and hypothetical results. The top panes show the 200 initial parallel replicates seeded with the ancestral genotype and evolved for 250,000 updates. We extracted the lineage of the most abundant genotype in the evolved population (the dominant lineage), shown in red. Next, we conducted exploratory replays (middle panes) by launching replay replicates at regular intervals along focal lineages. These exploratory replays give a coarse-grained view of how potentiation changed over a lineage. We identified the window with the largest potentiation gain, shown as the shaded region. Finally, we ran targeted replay replicates for every step in this potentiation window. These fine-grained replay replicates show mutations that resulted in large potentiation increases (shown here with a red dot).

ment with a default ancestor only capable of reproduction. Each replicate was given 250,000 updates, where one update is the time it takes for all organisms to execute 30 instructions, on average. We identified the most abundant genotype in each final population to represent the replicate and classi-

fied its behavior. We then extracted the “dominant lineage”, stretching from the ancestor to the representative genotype. Each step in the lineage corresponds to a change in genotype between parent and offspring, with clonal offspring occupying the same step as their parent. While a step is often a single mutation, it is possible that one step is composed of multiple co-occurring mutations.

To begin analyzing changes in potentiation, we ran exploratory replays replicates on four lineages capable of learning. For each, we seeded independent replicates for every 50th step in the lineage, up to step 1,000. All replay replicates evolved in the same associative learning environment, and replays were given the same number of updates as had occurred after that genotype first appeared (e.g., replays for a genotype that appeared at update 150,000 would be evolved for the remaining 100,000 updates). Potentiation was measured as the portion of replay replicates that evolved learning. Because replays were seeded with a single organism, some replay populations went extinct before ever reproducing and thus were not factored in (the minimum number of finished replay replicates from a given lineage step was 38, while three case study lineages had a minimum of 48).

While the exploratory replays provide an overview of how potentiation changed, we dug deeper by running targeted replays to further explore windows of increasing potentiation. Specifically, we found the 50- or 100-step “potentiation window” that sees the largest increase in potentiation in the exploratory results, and seeded additional replays for every step in that range. These targeted replays were conducted identically to the exploratory replays, only they did not skip steps. Though computationally expensive, these replays illuminated the impact every genotypic change had on potentiation. Running 50 replay replicates per step still results in considerable noise, but we were able to identify mutations that clearly and substantially increased potentiation using these targeted replays.

We hand-analyzed algorithms in all potentiating mutations, here defined as single lineage steps that result in a potentiation increase of 25 percentage points or more. Further, we assessed genotype fitness in context of their lineage to identify if potentiation mutations were beneficial, deleterious, or neutral. Finally, we characterized the local fitness landscape of each genotype (one and two mutations out), measuring the presence and fitness of nearby genotypes capable of learning.

Data and software availability

Both the data and the software used to conduct this work are available in the supplemental material (Ferguson, 2023). Analyses were conducted in the R statistical computing language (R Core Team, 2021) using the *dplyr* package to summarize data (Wickham et al., 2022). Data was visualized using the *ggplot2* and *cowplot* packages (Wickham et al., 2020; Wilke, 2020).

Results

Here we discuss the generation and analysis of the initial replicate runs, followed by the more detailed results from each of the four case study lineages.




Figure 2: Behavior classification of the final dominant genotypes from the 200 initial parallel replicates.

Evolution of learning in the initial replicates

In the first phase of this study, we evolved 200 independent replicates for 250,000 updates (about 3600 generations) in the associative learning domain, each starting with a default ancestor. The distribution of evolved behaviors is shown in Figure 2. Only 16 of 200 replicates exhibited associative learning. An additional 15 replicates evolved forms of bet-hedged learning, with two of those replicates gaining and then losing associative learning along their lineage. The majority of replicates relied on some form of error correction, either as a sole strategy (140), a bet-hedged variant (2), or as a fallback due to limited learning (3). Finally, 24 replicates failed to navigate enough states to categorize them, leading us to label them as “low activity”.

We analyzed all 16 replicates that evolved and maintained learning, identifying the length of their lineages from onset of evolution until learning stabilized, no longer showing further improvement. Given the substantial computational power required for each time point studied, we performed replay experiments only on the shortest three such lineages (lineages A-C), plus the shortest lineage that exhibited error correction at some point in its evolutionary history (lineage D). Selecting these particular replicates to replay has the potential to bias the results, as discussed below.

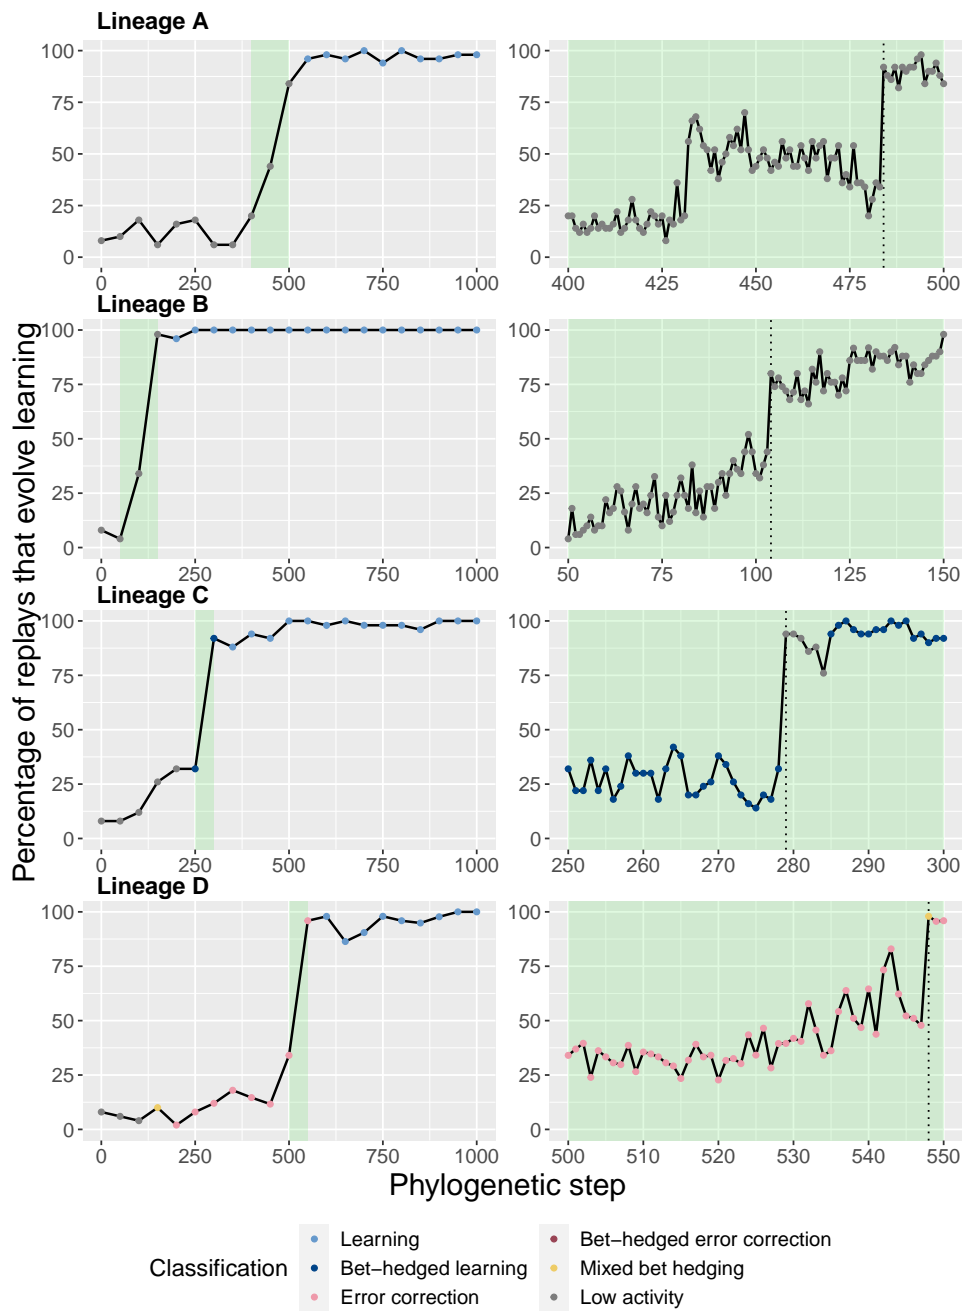


Figure 3: Potentiation of associative learning for all case studies, shown as the percentage of replicates that evolve associative learning when replayed from a given genotype. For each case study, the left plot shows the results of the exploratory replays. We identified a window of potentiation gain in each lineage, indicated by the shaded region. Within that window, we conducted targeted replays for every step along the lineage, shown on the right. The color of the points corresponds to the behavior exhibited at that step of the lineage. A dotted line in the targeted replays indicates the step that conferred the most potentiation.

Case studies of individual lineages

Below we present the results of the replay experiments performed on the four focal lineages and provide a step-by-step analysis of how key mutations altered both immediate fitness and evolutionary potential (potentiation). Where pos-

sible, we explained how these mutations altered the underlying algorithms. For each lineage, potentiation across both exploratory and targeted replays can be found in Figure 3.

Lineage A Our first case study is one of the shortest lineages and it contains a quick jump to learning (at step 537) from low activity, with only a brief time spent in bet-hedged learning. Exploratory replays for this lineage revealed a stark jump in potentiation between 400 and 500 steps from the ancestor, which we call the potentiation window. From step 400 to step 450, potentiation increased from 20% of replicates to 44%, and increased again to 84% of replicates by step 500. Before this window, potentiation fluctuated around the original 8% found starting from the default ancestor. After the selected window, potentiation increased once more and then fluctuated between 94% and 100%.

With this in mind, we seeded 50 replays for each genotype in the potentiation window. Even with the noise due to a small sample size, we identified two mutations that conferred sizable increases in potentiation: steps 432 and 484. The mutation at step 432 brought potentiation above 50% for the first observed time in the lineage. Surprisingly, potentiation then decreased (on average) back to a local minimum of 20% of replicates at step 480. Finally, the mutation at step 484 substantially increased potentiation to 92%, where it stayed for all subsequent replays.

Even though the largest jump in potentiation occurred at step 484, learning did not appear in the lineage until step 537. That said, only steps 516 and 525 caused any change in behavior; all other interim mutations occurred in unexecuted regions of the genome. The potentiating mutation at step 484 made a key instruction in the main loop of the genome redundant. It had no immediate effect on fitness, but later (in intermediate step 516) allowed the redundant instruction to be replaced by a right turn that granted a small fitness increase as organisms could now navigate until they reached the second left turn. Step 525 further improved navigation, but used a comparison that made an unfounded assumption on whether the left or right random cue is larger. When the assumption was correct organisms were capable of learning the cues, however the assumption is only correct 50% of the time, so this genotype is categorized as bet-hedged learning. Finally, step 537 swapped that comparison with one that makes no assumptions about cue values, enabling the genotype to learn in all environments.

Looking at the local mutational neighborhood, the potentiating mutation at step 484 increased the number of two-step mutations that conferred learning from 2 to 9 (of approximately 56 million). Additionally, the fitness of the learning mutations in the local neighborhood increased by three or four orders of magnitude.

What about the earlier potentiation that was gained and then lost? The mutation that substantially increased potentiation at step 432 introduced a comparison that had no immediate fitness effect. This comparison remained unimportant until step 525 when it became integral in introducing bet-hedged learning. Neither step 432 nor its predecessor had access to learning within a two-step mutational range. Thus,

it is likely that the potentiation comes from that comparison given that we observed it being utilized for learning later on.

Why then, did potentiation decrease between steps 432 and 484? At step 432 (and indeed before it), the algorithm had a section where if register B was non-zero, then B stored the cue associated with a left turn. While this information was likely to make the evolution of learning easier, it was unused at that time. As such, the mutations between steps 432 and 484 dismantled that machinery, requiring a replacement to be built before learning could evolve.

Lineage B Similar to Lineage A, this lineage transitioned from low activity to learning through a brief period of bet-hedged learning. Exploratory replays on this lineage reveal that learning was potentiated almost immediately; by step 150 potentiation had climbed above 95%, where it stayed for the rest of the lineage. As such, the potentiation window included steps 50 through 150.

Unlike Lineage A, the targeted replays reveal a general trend of increasing potentiation, with step 104 as a notable outlier. Mutations from steps 50 through 103 slowly increased potentiation from 4% to 44%, but the mutation at step 104 jumped to 80%. From there, another slow increase continued to raise potentiation to a peak of 98% at step 150.

Learning did not appear until step 195, over 90 steps beyond the largest potentiating mutation. Given that 34 intermediate mutations altered the encoded algorithm, the mechanistic pathway to achieve learning is more complicated than can be broken down in this work. However, the potentiating mutation at step 104 modified the execution flow of the genome, which appears to have been essential for the later evolution of learning.

While two mutations occurred at step 104, only one caused a functional change: an instruction to swap data between registers was mutated to a left turn. Prior to this mutation, the genome encoded a left turn later on, after which the execution became trapped in an endless loop. The potentiating mutation was immediately beneficial; it allowed organisms to take the left turn earlier, which, in turn, allowed them to avoid the loop. As a side effect, a large portion of the genome that was previously executed was now skipped, and these instructions remained skipped when learning evolved 91 steps later. Looking at the local fitness landscape, learning was neither present in the potentiating step's landscape nor in the step before. We hypothesize that the potentiation came from the change in execution flow, and that skipping over those instructions avoided a pitfall and freed up execution time that may have been useful in evolving learning.

Lineage C Lineage C has the biggest single-mutation potentiation increase (64 percentage points), and that mutation was deleterious when it occurred at step 279 along the lineage. Learning later appeared at step 305.

The potentiating step mutated a no-operation instruction

into a conditional flow control instruction. At step 278 the genotype was capable of bet-hedged learning, but the mutation at 279 knocked out all instances of learning, reclassifying the behavior as “low activity.” The next step restored some fitness, and then step 285 interacted with the mutation at 279 to not only restore fitness, but to dramatically improve it. Ultimately, the potentiating mutation at step 279 allowed the algorithm to more precisely discriminate between the left and right cues. Prior to the mutation, a less-than comparison was used, which only functioned correctly in 50% of instances, specifically those where the *left* cue was less than the *right* cue. The potentiating mutation switched to an equality comparison, which alleviated the assumption and initiated the transition from bet-hedging to learning.

Interestingly, the potentiating mutation lowered both the number and average fitness of learning mutations available in the local fitness landscape,

Lineage D Of the four replicates we analyzed, Lineage D is the only one that evolved error correction before learning. Like the earlier lineages, the exploratory replays show that almost all potentiation comes from a single window. In this case, potentiation grew from 34% of replicates to 96% between steps 500 and 550. Targeted replays are especially noisy for this lineage, but generally show an increase in potentiation, especially in the latter half of the window. The largest jump in potentiation occurred at step 548, near the end of the window. Several prior mutations also showed notable potentiation increases, but in each case later mutations appeared to counteract them. Specifically, steps 542 and 543 appear to have higher potentiation than the points around them, but potentiation dipped back below 50% before the largest jump at 548.

Out of all four lineages, D has the fewest steps between the largest potentiating step (548) and the first appearance of learning (556). At the time, the potentiating mutations at 548 caused no discernible change in fitness even though they increased potentiation by 50 percentage points. Two mutations occurred at step 548: a point mutation swapped a flow control instruction for a math instruction and an insertion mutation added a comparative conditional instruction into the main execution loop. At step 548 the genotype encoded a naive error correction algorithm: after setup, organisms could always handle *right* states, but always failed *left* states, recovered, and then continued. A mutation at step 556 swapped a sensing instruction with a math instruction, and this combined with the prior comparison instruction from step 548 to allow the organism to move left when needed and shifted the behavior to learning. The local fitness landscape supports the idea that the comparison instruction was useful to the evolution of learning, as the potentiating mutation increased the number of learning genotypes in the two-step mutational neighborhood from under 800 to over 100,000.

Looking back at the apparent false start at steps 542 and

543, it is not clear what algorithmic changes these mutations conferred. These steps did, however, alter the set of learning behaviors that fell within the local mutational neighborhood. At step 541, there were only 324 two-step mutations that conferred learning, and only three of those resulted in a substantial fitness increase (a merit $> 10^{15}$). After steps 542 and 543, that number rose such that over 900 two-step mutations could confer learning, with over 500 resulting in a substantial fitness increase (including over 200 that reached a merit $> 10^{25}$). We may be unsure of the exact effects of these mutations on the mechanics of the algorithm, but the changes in the local fitness landscape are profound.

Discussion and Conclusion

Potentiation can rise suddenly

We have documented several cases where single mutations dramatically increased the probability of associative learning later evolving. Of the four lineages analyzed, each had a single step in the lineage that resulted in a substantial increase in potentiation (ranging from 36 to 64 percentage points). Indeed, two of the lineages had an additional potentiating mutation that resulted in an increase of over 30 percentage points. Looking only at exploratory replays, each lineage has a 50-step window that resulted in a potentiation increase of at least 40 percentage points.

While four lineages are insufficient to make any strong claims, these results demonstrate that it is *possible* for single mutations to drastically increase potentiation, and provide compelling evidence that they may, in fact, be common. In Lineages B and D, however, we do also observe regions with smaller, incremental increases in potentiation. Further studies are clearly necessary to more fully understand the general patterns and processes by which potentiation rises across different representations and environments.

Potentiation can decrease along a successful lineage

In two of the lineages we analyzed (A and D), we see evidence of potentiation decreasing over spans of the lineage. With only 50 replay populations per lineage step, our results are noisy and it is difficult to isolate what is occurring during these periods of potentiation decline. While we were unable to identify any “anti-potentiating” mutations with effects as large as the positive potentiation mutations, it is possible for a single step in a lineage to greatly decrease potentiation.

Since we limited our analyses to runs where associative learning arose in the original replicate, we did not expect a preponderance of anti-potentiating mutations, but were intrigued to see evidence of them, even if at low effect. These same analytic replay experiments could be applied to lineages that failed to evolve the target behavior, to see if potentiation of that behavior experiences sudden drops. Similarly, our replays targeted windows with substantial increases in potentiation; other windows would be more likely to include decreases. Finally, failed replays from starting points with

otherwise high potentiation must have failed for a reason; they too could be used as a likely source (albeit more artificial) of anti-potentiating mutations.

Potentiating mutations can appear innocuous when they first occur

We analyzed the mutational step in each of the four lineages that conferred the greatest increase in potentiation. Of those four mutational events, two were neutral, one was deleterious, and one was beneficial. Even among these few replicates, there is no obvious pattern in the properties of potentiating mutations. Of the two neutral mutations, one made an instruction redundant while the other added a conditional instruction that had no effect when it was initially introduced. The deleterious and beneficial mutations both caused the execution flow to loop back earlier than it did before. Additionally, the number of mutations between the potentiating mutation and the appearance of learning varied wildly between lineages, ranging from 8 steps up to 91. The potentiating mutations in these four lineages are unique, and at the current time there is no pattern emerging among them. So, how are these mutations any different from other mutations? Untangling this mystery could be critical for predicting evolutionary outcomes or accelerating adaptive evolution.

We can identify *how* a mutation is potentiating

There are many mechanisms by which a mutation could facilitate the evolution of associative learning. For example, the mutation could provide a building block that is helpful to perform the task. But for a mutation to be potentiating it must notably increase the probability of associative learning appearing in the future. Any change, no matter how helpful, that was already likely to occur would not be considered potentiating. Indeed, it is the earlier mutations that made that change so likely that would be potentiating. Of course, those mutations are also more challenging to identify.

We have three different hypotheses for how a mutation could be potentiating: (1) It is the initial move into a genetic neighborhood with associative learning, (2) It is a shift into a genetic neighborhood with a more valuable version of learning, or (3) It is a “gateway” mutation that unlocks a beneficial pathway to learning, even though learning is not in the immediate genetic neighborhood.

Across the potentiating mutations we analyzed, we have found evidence for each of these hypotheses. The largest potentiating mutation in Lineage D supports Hypothesis 1, as it is the first time in the lineage that learning is only one mutation away. The main potentiating mutation in Lineage A and the earlier potentiation mutations in Lineage D support Hypothesis 2 as both cause drastic increases in the fitness benefit of learning mutations in the two-step neighborhood. Finally, Lineages B, C and the early potentiating mutation from Lineage A all provide support for Hypothesis 3. The mutations from Lineages A and B both have *zero* learning

mutations in their two-step neighborhoods. Interestingly, Lineage C sees a *decrease* in the number and fitness of learning mutations in the local neighborhood.

Hypothesis 3 has many possible mechanisms by which it may work. For example, new traits may produce a single, clear, beneficial pathway of improvements to follow. Alternatively, a new building block may open a larger region with many different ways of evolving associative learning. Finally, the mutation may actually damage existing functionality or remove existing interactions that were impeding further evolution. While all three hypotheses have some support, future work can begin to uncover if a certain hypothesis is seen more often, what conditions might result in each scenario, or if additional analyses are needed to truly characterize these potentiating mutations.

Outlook

This work is only an early step, focused on developing techniques and expectations for performing fine-grained analyses of replay experiments. Next, we must expand beyond four lineages, to collect broader, more systematic replay data, automating as much of the process as possible. We conducted this study on associative learning in Avida, but the underlying techniques must be examined broadly in other environments and substrates to ensure that our results are not unique to Avida or the evolution of associative learning. Within the current study system, there are many questions that remain unanswered: We focused on large *increases* in potentiation, but are there more obvious signals associated with *decreases*? How much of the noise that we see in our data is due to limiting ourselves to 50 replicates, and how much of it is due to actual shifts in potentiation with each mutation? What does potentiation look like in replicates that fail to evolve learning? Finally, it would be valuable to compare the specific evolutionary pathways the different replays take. Do they follow the same trend or do they differ? This would allow us to understand if, for example, a potentiating mutation funnels evolution in a fixed direction.

Ultimately, these analytic replay techniques provide us with a tool for examining evolution in a prospective fashion, not just the retrospective approach that we are traditionally limited to. They will allow for the development of new evolutionary theory and predictive capacity that will be invaluable, both for understanding how meaningful complexity is produced in the natural world and for improving evolutionary applications.

Acknowledgements

We thank Anselmo Pontes, Cliff Bohm, the reviewers, and the MSU Digital Evolution and ECODE labs. This work was supported by the BEACON Center for the Study of Evolution in Action, U.S. National Science Foundation grant DBI-0939454. Computational resources were provided by the MSU Institute for Cyber-Enabled Research.

References

- Blount, Z. D., Barrick, J. E., Davidson, C. J., and Lenski, R. E. (2012). Genomic analysis of a key innovation in an experimental *Escherichia coli* population. *Nature*, 489(7417):513–518.
- Blount, Z. D., Borland, C. Z., and Lenski, R. E. (2008). Historical contingency and the evolution of a key innovation in an experimental population of *Escherichia coli*. *Proceedings of the National Academy of Sciences*, 105(23):7899–7906.
- Blount, Z. D., Lenski, R. E., and Losos, J. B. (2018). Contingency and determinism in evolution: Replaying life’s tape. *Science*, 362(6415):eaam5979.
- Covert III, A. W., Lenski, R. E., Wilke, C. O., and Ofria, C. (2013). Experiments on the role of deleterious mutations as stepping stones in adaptive evolution. *Proceedings of the National Academy of Sciences*, 110(34):E3171–E3178.
- Dunlap, A. S. and Stephens, D. W. (2014). Experimental evolution of prepared learning. *Proceedings of the National Academy of Sciences*, 111(32):11750–11755.
- Ferguson, A. J. (2023). Supplemental material. Zenodo and GitHub. <https://doi.org/10.5281/zenodo.7953692>.
- Gould, S. J. (1990). *Wonderful Life: The Burgess Shale and the Nature of History*. WW Norton & Company.
- Grabowski, L. M., Bryson, D. M., Dyer, F. C., Ofria, C., and Pennock, R. T. (2010). Early Evolution of Memory Usage in Digital Organisms. In *ALIFE*, pages 224–231.
- Jochumsen, N., Marvig, R. L., Damkjaer, S., Jensen, R. L., Paulander, W., Molin, S., Jelsbak, L., and Folkesson, A. (2016). The evolution of antimicrobial peptide resistance in *Pseudomonas aeruginosa* is shaped by strong epistatic interactions. *Nature Communications*, 7(1):13002.
- Kawecki, T. J., Lenski, R. E., Ebert, D., Hollis, B., Olivieri, I., and Whitlock, M. C. (2012). Experimental evolution. *Trends in Ecology & Evolution*, 27(10):547–560.
- Lenski, R. E., Ofria, C., Pennock, R. T., and Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936):139–144.
- Lenski, R. E., Rose, M. R., Simpson, S. C., and Tadler, S. C. (1991). Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations. *The American Naturalist*, 138(6):1315–1341.
- McGregor, S., Vasas, V., Husbands, P., and Fernando, C. (2012). Evolution of Associative Learning in Chemical Networks. *PLOS Computational Biology*, 8(11):e1002739.
- Mery, F. and Kawecki, T. J. (2002). Experimental evolution of learning ability in fruit flies. *Proceedings of the National Academy of Sciences*, 99(22):14274–14279.
- Meyer, J. R., Dobias, D. T., Weitz, J. S., Barrick, J. E., Quick, R. T., and Lenski, R. E. (2012). Repeatability and contingency in the evolution of a key innovation in phage lambda. *Science*, 335(6067):428–432.
- Ofria, C. and Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial life*, 10(2):191–229.
- Pontes, A. C., Mobley, R. B., Ofria, C., Adami, C., and Dyer, F. C. (2020). The evolutionary origin of associative learning. *The American Naturalist*, 195(1):E1–E19.
- R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Travisano, M., Mongold, J. A., Bennett, A. F., and Lenski, R. E. (1995). Experimental tests of the roles of adaptation, chance, and history in evolution. *Science*, 267(5194):87–90.
- Wagenaar, D. A. and Adami, C. (2004). Influence of Chance, History, and Adaptation on Digital Evolution. *Artificial Life*, 10(2):181–190.
- Wickham, H., Chang, W., Henry, L., Pedersen, T. L., Takahashi, K., Wilke, C., Woo, K., Yutani, H., and Dunnington, D. (2020). *Ggplot2: Create Elegant Data Visualisations Using the Grammar of Graphics*.
- Wickham, H., François, R., Henry, L., and Müller, K. (2022). *Dplyr: A Grammar of Data Manipulation*.
- Wilke, C. O. (2020). *Cowplot: Streamlined Plot Theme and Plot Annotations for Ggplot2*.
- Yedid, G., Ofria, C. A., and Lenski, R. E. (2008). Historical and contingent factors affect re-evolution of a complex feature lost during mass extinction in communities of digital organisms. *Journal of evolutionary biology*, 21(5):1335–1357.
- Zaman, L., Meyer, J. R., Devangam, S., Bryson, D. M., Lenski, R. E., and Ofria, C. (2014). Coevolution drives the emergence of complex traits and promotes evolvability. *PLoS Biology*, 12(12):e1002023.