

HARNESSING THE COMPLEXITY OF NATURAL EVOLUTION TO ENHANCE
EVOLUTIONARY ALGORITHMS

By

Vincent Romeo Ragusa

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Computer Science—Doctor of Philosophy

2023

ABSTRACT

Evolutionary computation is a powerful optimization tool, and an invaluable test bed for population genetics. Evolutionary algorithms can become stuck on local optima, but can escape these traps by temporarily losing fitness in order to discover even higher fitness in a process called valley-crossing. Valley-crossing is fundamentally linked to the balance between the forces of selection and variation, and as such, controlling this balance is important for optimizing the efficiency of evolutionary algorithms. Nature, in contrast, is not actively optimized for performance, and yet nature seems to overcome many challenges that evolutionary algorithms do not. It is possible that nature benefits from a highly dynamic balance between selection and variation, and this constant flux helps natural populations avoid stagnation and overcome obstacles in the fitness landscape. Working with this hypothesis in mind, I investigate the nature of selection and how natural phenomena strengthen or weaken it.

I find that selection strength can be thought of as the degree to which an evolving system is dissimilar to neutral drift. This perspective opens the door to accept all phenomena that affect the strength of selection as part of a unified theory of selection that treats selection strength as an emergent property. I present a new evolutionary dynamic – the free-for-all effect – that is the reduction of selection strength on organisms with higher-than-average fitness. Free-for-all can result in rapid evolutionary adaption that would otherwise seem impossible, and provides an elegant explanation for punctuated equilibrium. The discovery of free-for-all highlights the importance of spatial structure in evolving populations, and has led to the design of a new evolutionary search method called super explorers. Super explorers mimic the free-for-all effect, and improve evolutionary search, while placing full control into the hands of the algorithm designer.

Copyright by
VINCENT ROMEO RAGUSA
2023

ACKNOWLEDGMENTS

I want to thank my advisor Chris Adami, who exemplifies what being a mentor is all about. Chris welcomed me into his lab and helped me flourish. I would be remiss not to thank Arend Hintze, who mentored me for the first half of my journey, and always encouraged me to follow my instincts. I want to thank my committee members Charles Ofria and Wolfgang Banzhaf, who have lent me their vast knowledge and perspective over the years. Of course, I must thank Cliff Bohm. Cliff and I have shared an office for the last 6 years, and I am grateful for every moment. Cliff is a great collaborator, and an even greater friend. There are a number of people who have had a disproportionately large impact on my life. Many of them helped steer me to where I am today, knowingly or not. I want to thank Michael Harper, Matt Flemming, Anthony Rodriguez, and David Mathias for fostering my love of computer science and inspiring me to keep learning. Last but not least, thank you to my family and friends who have supported me with their love and laughter.

TABLE OF CONTENTS

Chapter 1 Introduction	1
1.1 Background and Motivation	1
1.2 Literature Review	2
1.3 Research Objectives and Contributions	17
1.4 Research Methods	18
1.5 Outline of Dissertation	21
1.6 Summary	22
Chapter 2 Quantifying the impact of selection	23
2.1 Introduction	23
2.2 Properties of selection strength	25
2.3 The Selection Impact Metric	27
2.4 Modeling the Offspring Distributions of Neutral Drift	29
2.5 Counting Offspring and Descendants	34
2.6 Empirical data	36
2.7 Discussion	42
2.8 Methods	47
Chapter 3 Noisy fitness and selection strength	49
3.1 Introduction	49
3.2 Methods	52
3.3 Results	57
3.4 Discussion	63
Chapter 4 The role of disequilibrium in evolutionary discovery	69
4.1 Introduction	69
4.2 Results	73
4.3 Discussion	81
4.4 Conclusion	89
4.5 Methods	90
Chapter 5 Augmenting evolutionary algorithms with Super explorers	94
5.1 Introduction	94
5.2 Methods	97
5.3 Results	104
5.4 Discussion	108
5.5 Conclusion	115
Chapter 6 Conclusions	117
6.1 Measuring the impact of selection	117
6.2 Altering selection strength with fitness noise	120
6.3 The evolutionary free-for-all effect	122
6.4 Super Explorers: protected, sustained, neutral drift	128
6.5 Final Remarks	131

BIBLIOGRAPHY	134
-------------------------------	------------

Chapter 1

Introduction

1.1 Background and Motivation

Evolutionary computation (EC) is a field of artificial intelligence concerned with solving optimization problems with algorithms inspired by evolution. These bio-inspired optimization algorithms are called evolutionary algorithms (EAs). In order to solve an optimization problem with EC, a population of digital organisms is created, which each represent a candidate solution to the optimization problem. Some candidate solutions will be of higher quality than others (they have higher fitness) and those high-fitness candidates are selectively chosen to reproduce to form the next generation of solutions. The selection bias towards higher fitness organisms increases the average fitness of the population over time. Digital organisms tend to produce offspring identical to themselves; the information that defines the parent is inherited by the child. Occasionally, however, small errors are introduced to create variation among the members of the population. These small mutations change the child organism and the candidate solution they represent. These random mutations are sometimes beneficial and increase the organism's fitness. Since organisms with high fitness are preferred by selection, the population will slowly evolve towards ever higher fitness as these beneficial mutations are discovered and build on one another. In the end, when there are no more

improvements to discover, the highest quality solution to the optimization problem is now available, and the optimization problem is solved. Hopefully.

The scenario I just described is the ideal outcome of using EC to solve an optimization problem. In practice, there are no guarantees that an evolving population will discover the highest fitness solution. Sometimes the population does not even discover a solution that can work practically in the real world. There are many reasons why EAs may struggle to discover paths towards higher fitness. The optimization problem may simply be difficult to solve, or difficult for EAs specifically. Alternatively, the problem may be solvable, but the design of the EA may be ineffective for this specific problem. Perhaps instead, the EA *is* well suited to solving the problem, but the poor performance is due to its various settings and parameters being configured ineffectively. When an EA fails to find good solutions, it is never quite clear which of these problems may be the one plaguing it. Each of these problems are the subject of ongoing research by many scientists and engineers around the world.

My general interest is in the design and tuning of the evolutionary algorithms, and the theory behind how evolution solves problems. Specifically, I am interested in selection, and how parts of an evolutionary algorithm that are ostensibly unrelated to selection can have a large influence on selection. These complex relationships can create unforeseen obstacles, as well as offer hidden benefits.

In search of both obstacles and benefits, I look to the complexity of nature. The study of natural evolution and related theory provide a rich source of inspiration. As EAs were originally inspired by nature, I turn once again to nature for new ideas about how to improve EAs and overcome the obstacles EC faces.

1.2 Literature Review

1.2.1 Evolutionary Algorithms

Evolutionary algorithms (EAs) [Holland, 1992, Koza, 1994, Rechenberg, 1984, De Jong, 2012] are search and optimization algorithms. In other words, given a description of an optimization

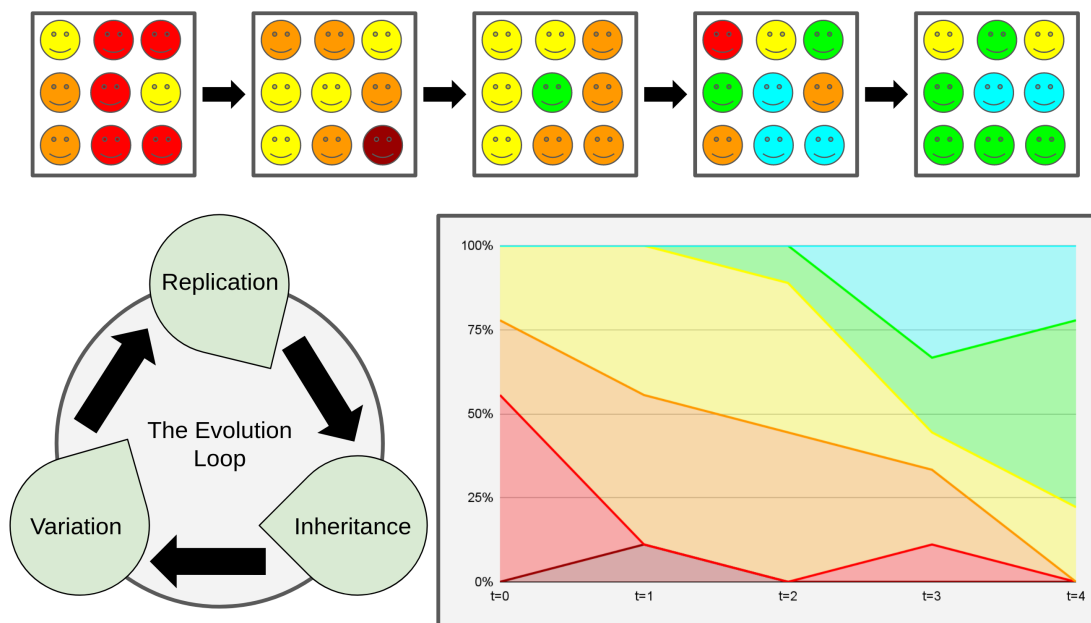


Figure 1.1: A toy evolutionary algorithm, evolving the color of 9 digital organisms. Green is the fittest color, with colors further from green in chromatic order becoming increasingly less fit. **Top:** five generations of organisms, with black arrows indicating the progression of time. **Bottom left:** the evolution loop: replication, inheritance, and variation. Each iteration of the loop produces the next generation of organisms. Selection occurs because the fittest organisms produce more offspring during the reproduction step. **Bottom right:** the population’s color composition over time. The average color of the population slowly changes to green.

problem, an EA will search for solutions to the problem, and attempt to maximize the quality of the solutions it finds. It does this by emulating the evolution of a population of digital organisms [Adami, 1998, Ofria and Wilke, 2004, Adami, 2006]. Each digital organism contains a blueprint for a solution to the optimization problem. We call this blueprint the organism’s genome¹. Each organism has a genotype — the unique set of genes the organism carries — and a phenotype, the set of all traits, functions, and behaviors the organism expresses as a result of its genotype. The organism’s phenotype is often the candidate solution it represents. Figure 1.1 provides a simple description of an evolutionary algorithm.

In any evolving system, the *true* measure of an organism’s fitness is its reproductive success². Therefore, in order to leverage evolution to solve optimization problems, we must

¹The genome of biological organisms are not ‘blueprints’ in this sense, but are more like the organism’s instruction manual or operating system, however even these analogies are imperfect.

²Some researchers talk about the long-term survival of the entire lineage that begins with the organism in

link the quality of an organism’s solution to the organism’s reproductive success. The quality of an organism is assessed according to its phenotype: its traits, functions, and behaviors. The specific way in which the organism’s phenotype is mapped to a metric of quality determines which solutions will evolve. A fitness function accomplishes the mapping from aspects of the organism’s phenotype to a quality score or scores. This quality score is then used by a selection algorithm to decide which organisms reproduce, and how many offspring they will have. It is commonplace (and sometimes confusing) to call the quality score assigned to the organism its fitness. In reality, an organism’s fitness can only be determined in retrospect, for example by counting its offspring. The organism’s quality score is merely *proportional to* the organism’s fitness. However, for ease of writing, and to conform to the conventional lexicon of EC, I will call the quality score ‘fitness’ whenever the distinction is irrelevant.

In EC, as opposed to nature, selection is performed by a selection algorithm. A selection algorithm looks at the quality scores assigned to each organism, and gives organisms with higher quality a chance to reproduce more often than those with lower quality. Selection algorithms are not unlike animal breeders in that they have full authority over which organisms reproduce and how many offspring they will have. Since selection favors organisms with higher quality scores, the population will evolve towards higher quality scores, and thus towards high quality solutions to the optimization problem.

When an organism is selected to reproduce, it passes its genotype to its offspring. The offspring’s genotype may also receive mutations which make it distinct from its parent. If sexual recombination is desired, two or more parents can be selected to produce offspring together. The newly born organisms must also be evaluated by the fitness function in order to participate in the next round of selection. Sometimes the mutation and recombination of genomes results in a solution whose quality is higher than any other seen previously, moving the population collectively towards higher fitness.

question. This ‘long-term fitness’ can only be determined in retrospect, however the organism’s reproductive success is (usually) a good predictor of its long-term fitness.

1.2.2 Fitness Landscapes and Valley-crossing

Fitness landscapes [Wright, 1932, Maynard Smith, 1970] are conceptual and visual aids for understanding the fitness relationships between genotypes. Like a topographical map with elevation markings, fitness landscapes show higher fitness genotypes at higher elevations than lower fitness genotypes. Adjacent locations in the fitness landscape correspond to genotypes that are similar to one another. Typically, the distance between two genotypes is the mutational (Hamming) distance between them, but other measures of distance can be used. Organisms at the same elevation in the fitness landscape have the same fitness, and therefore have no fitness advantage over one another. Organisms at higher elevations have a fitness advantage over organisms at lower elevations than them. Inheritance and variation create organisms at nearby locations, while selection favors those at higher elevations. The net result is a population of organisms that climb the hills and wander the plains of their fitness landscape.

An evolving population tends to climb the slopes of its fitness landscape. However, if climbing upwards towards higher fitness was *all* the population was capable of, it would quickly reach the top of some hill, more formally called a local optimum, and never retreat downward³. If a population reaches a local optimum, there are no *single* mutations that can increase the fitness of an organism further. This poses an obstacle for evolutionary computation algorithms, which are intended to climb the *highest* hill in the fitness landscape: the global optimum. The way around this obstacle is to allow the population to sometimes retreat down from a local optimum in hopes of discovering another higher hill to climb. The process of escaping one optimum to climb another is called “valley-crossing”, since the space in the fitness landscape between two hills is like a valley. Valley-crossing can happen gradually, for example through the sequential fixation of deleterious mutations due to Muller’s ratchet [Muller, 1964], or quickly, with mutations occurring quickly without intermediate

³This aspect of evolution is effectively the stochastic gradient descent algorithm, popular in machine learning.

genotypes fixing in a process called stochastic tunneling [Iwasa et al., 2004a].

Although a population does not always require a loss in fitness to cross a fitness valley [Iwasa et al., 2004a, Østman and Adami, 2014], deleterious mutations are often necessary to explain valley-crossing events. For example, [Covert et al., 2013] disallow deleterious mutations in an evolutionary algorithm and observe that the adaptability of the system is hindered due to the inability to cross valleys. Similarly, it has been shown that elitism, strictly selecting only the highest fitness organisms, hinders adaptation on fitness landscapes with many local optima [Oliveto et al., 2018], showing that deleterious mutations play an important role.

Besides crossing fitness valleys, there are other means to escape local optima. Extradimensional bypass theory [Conrad, 1990, Cariani, 2002] describes how a fitness landscape may change over time, resulting in new paths towards higher fitness that do not involve deleterious mutations. The neutral theory of molecular evolution [Kimura, 1983, Kimura, 1991] posits that fitness-neutral genetic variation accounts for nearly all variation seen in nature. If true, the theory would suggest that natural populations traverse fitness landscapes primarily along neutral ridges [Gavrilets, 1999]. In other words, the mutational distance to a beneficial mutation limits the rate of adaptation, not the difficulty crossing fitness valleys. Neutral theory has not gone unchallenged; evidence has accumulated to suggest that selection plays a more important role than neutral theory would suggest [Lewontin et al., 1974, Kreitman, 2000, Fay et al., 2002, Hahn, 2008, Kern and Hahn, 2018]. Modifications to the theory have been made by [Ohta, 1992, Ohta, 2002] to allow for nearly-neutral variation, and in particular deleterious mutations, which once again introduces valley-crossing into the equation. Regardless of how neutral the natural fitness landscape is, we can be sure that optimization problems will represent many kinds of fitness landscapes, and EAs must be equipped to navigate them all.

1.2.3 No free lunch: the balance between selection and variation

Viewing evolution as a search-and-optimization algorithm allows us to import a massive body of literature from the optimization community. This body of literature can help

predict how evolution will unfold and establish the limitations of evolution as a problem-solving device. All search and optimization algorithms, including both natural selection and computer models of evolution, are subject to the fundamental limitations of the no-free-lunch theorems [Wolpert and Macready, 1995, Wolpert and Macready, 1997, Ho and Pepyne, 2002, Wolpert and Macready, 2005], and particularly to the explore-exploit tradeoff [Millidge et al., 2021]. In order to explore the fitness landscape and discover better solutions, some genomes must mutate far from the mass of the population, becoming less like the solutions that birthed it. In order to exploit the fitness landscape for further optimizations, some genomes must mutate near the mass of the population to fine tune the best solutions. Some genomes must not mutate at all in order to preserve copies of the existing solutions for future generations. Managing this tradeoff is typically a main concern in the development of evolutionary algorithms. The manner in which an evolutionary algorithm balances these concerns will determine the algorithm’s performance on a given fitness landscape.

We can think of the explore-exploit dilemma in terms of selection, variation, and inheritance. Selection is a purifying force, by which I mean selection favors the high-fitness individuals, while slowly driving all others to extinction. Inheritance is a cohesion force, by which I mean members of the population accumulate in the same vicinity within the fitness landscape when reproduction occurs. If selection and inheritance were the only phenomena occurring in the population, every organism in the population would eventually become a clone of the most fit individual. Together, selection and inheritance are the main drivers of exploitation dynamics in the evolutionary optimization process. Variation is a diffusive force, by which I mean mutations randomly move organisms around in the fitness landscape, spreading them out in all directions. Variation is the primary driver of exploration dynamics in the evolutionary search process. Inheritance plays a critical role in exploration as well, since it allows for mutations to build on one another, increasing the range of exploration. Inheritance also focuses the exploration in the vicinity of the population where good solutions are known to exist. It is, therefore, the opposing forces of selection and variation which

result in the explore-exploit dilemma. Both selection and variation are necessary aspects of evolution, but shifting the emphasis towards one aspect can also shift emphasis away from the other.

If selection and variation are necessary but contrary forces, then perhaps there is an optimal balance between these forces that maximizes the rate of adaptation. The no-free-lunch theorems tell us that any algorithm will have some inputs for which the algorithm behaves optimally, while for other inputs the algorithm will perform terribly⁴. If we think of the evolutionary algorithm and a particular set of parameters (held constant) as a single unit, we may consider the fitness function as the input, and deduce that these particular parameters will lead to optimal evolutionary search on some fitness landscapes, and terrible search on others. If instead we think of the evolutionary algorithm and a particular fitness function as a single unit, we may consider the algorithm's parameters as the input, and deduce that some parameters will lead to optimal evolutionary search on this particular fitness function, while other parameters will lead to terrible search. Furthermore, consider that the forces of selection and variation are a direct consequence of the evolutionary algorithm and the values of its parameters.

Therefore, as a consequence of the no-free-lunch theorems, for a given location on a fitness landscape, there is an optimal balance between selection and variation that results in the fastest adaptation from that location. Equivalently, for every balance of selection and variation, there are locations in the fitness landscape that are most easily traversed by that population. However, in general, there is no single optimal balance between selection and variation for all areas of every fitness landscape, and no fitness landscape is optimally traversed by every balance of selection and variation. Put simply, the 'best' configuration of the population depends entirely on where the population resides in the fitness landscape.

⁴In truth, the theorems state that all algorithms perform the same when averaged across all inputs. We may deduce from this that unless all inputs result in exactly the same performance, some inputs must achieve better outcomes than others.

1.2.4 Selection regimes and rates of adaptation

It is common in the literature to talk about strong versus weak selection, or strong versus weak mutation. In particular, the literature considers two regimes: strong selection / weak mutation (SSWM) and weak selection / strong mutation (WSSM)⁵. The SSWM and WSSM regimes are also known as the sequential-fixation regime and the successive-mutations regime [Desai and Fisher, 2007, Desai et al., 2007]. These regimes help us understand the different ways that an evolutionary system can balance selection and variation.

In the sequential-fixation regime, the discovery of beneficial mutations is rare, and so, once discovered, beneficial mutations go to fixation before another is discovered. In this regime, the rate of adaptation is limited only by the rate at which beneficial mutations are discovered. The sequential-fixation regime may, in addition to correlating with strong selection and weak mutation, also correlate with strong selection and *strong* mutation. To illustrate this point, consider a population with a high mutation rate, but very few beneficial mutations to discover in the local fitness landscape. This population is experiencing a large inflow of neutral and deleterious variation that can work against selection and result in exploration and potentially valley-crossing. Such a population would be in the sequential-fixation regime if those valley-crossings rarely happen and the population goes to fixation between each discovery.

In the successive-mutations regime, many beneficial mutations are discovered faster than a single one can fix, resulting in competition between beneficial mutations, called clonal interference (CI) [Gerrish and Lenski, 1998, Campos et al., 2004, Desai et al., 2007, Fogel et al., 2008]. CI slows down the rate of adaptation, since the population makes inefficient use of the beneficial mutations it discovers, often simply losing those discoveries to competition and drift.

These two evolutionary regimes are the extremes. In reality, a population is likely

⁵Sometimes the order of selection and mutation in these acronyms are written in the reverse order: WMSS or SMWS.

somewhere between the two regimes and constantly changing. Thus, these two regimes serve more as simplified models of evolution, with properties that make them easy to analyze and model, than as complete descriptions of evolutionary dynamics.

1.2.5 Shifting the balance

Since no single configuration of an evolutionary algorithm’s selection and variation forces can solve every problem, evolution should be a more powerful tool when it changes its selection and variation over time. Indeed, this has been found to be true. There is a large body of literature that describes methods for self-adaptive evolutionary algorithms, algorithms that tune their own parameters in response to how the population of solutions has evolved so far [Bäck, 1994, Smit and Eiben, 2009, Kramer, 2010]. Self-tuning algorithms manipulate mutation rates, mutation sizes, and other aspects of variation, as well as parameters that affect the strength of selection. Typically, only one of either variation or selection is tuned, since their opposing nature makes tuning both unnecessary for many applications. Other algorithms use niching techniques (a.k.a. diversity maintenance or diversity preservation) [Shir, 2012] to maintain variation in the population by means other than mutations alone. Typically, niching algorithms co-opt selection to apply pressure against the collapse of diversity, effectively shifting the balance towards more variation.

Perhaps unsurprisingly, nature is full of phenomena that can also shift the balance of selection and variation. For instance, the size of the evolving population relates to both selection strength and standing variation [Jain et al., 2011, Ochs and Desai, 2015, Rozen et al., 2008, Ohta, 1972, Lynch and Conery, 2003, Takahata, 1987]. Larger populations suffer less sampling error and are more sensitive to small fluctuations in fitness, leading to strong selection. Small populations more easily experience genetic drift, a phenomenon where the frequency of a genotype in the population has a larger impact on its reproductive success than the quality of its corresponding phenotype. The decoupling of quality from fitness results in weak selection. It is important to note that adjusting the population size changes more than just the selection strength felt by the individuals in the population. Increasing

the population size also increases the mutation inflow, the number of individual mutations a *population* experiences in a unit of time. A larger inflow of mutations increases the standing variation of genotypes in the population [Feller, 1951, Kimura, 1964]. Whether increasing the population size results in a net shift of the balance towards selection or variation will depend on the circumstances⁶.

It has also been documented recently that noisy fitness, for example due to environmental noise, can reduce the strength of selection. For example, [Wang and Zhang, 2011] and [Melbinger and Vergassola, 2015] both see an increase in drift-like behavior when fitness is noisy and the selection algorithm has difficulty comparing organisms accurately. Both publications describe the change in dynamics as being similar to a reduction in the effective population size. A third study by [Van Egeren et al., 2018] demonstrates how noisy fitness can also benefit adaptation by weakening selection enough to enable valley-crossing.

When a population discovers new territory and subsequently moves or grows to inhabit it, the population is said to undergo a range expansion. Range expansion events have been documented to reduce selection temporarily on the leading edge of the expansion [Slatkin and Excoffier, 2012, Peischl et al., 2013, Peischl and Excoffier, 2015, Peischl et al., 2015, Gilbert et al., 2017, Burton and Travis, 2008]. Once the range expansion has ended, the reduction in selection strength ends as well.

Mass-extinction events are also known to increase the adaptive capability of the surviving members of the population [Lewis, 1962, Mathias and Ragusa, 2016, Engholm and Mathias, 2021, Lehman and Miikkulainen, 2015, Slater, 2013, Sahney et al., 2010], likely due to reduced competition on less-fit genotypes and the subsequent range expansion that occurs to replace the lost biomass.

Selection is not the only aspect of evolution that can change naturally. For example, mutation rates have also been documented to shift naturally [Lenski et al., 2003]. However, mutation rates, unlike selection strength, are a property of the organism and so subject to

⁶Increasing the population size “to infinity” results in both infinite selection strength and infinite variation, thus infinite populations are non-physical models that represent an impossible idealized population.

evolution [Lynch et al., 2016]. For this reason, shifts in mutation rate are not as capricious as shifts in selection strength over the same time scales.

The shifting balance between selection and variation was proposed long ago as a mechanism for valley-crossing in natural populations. The aptly named ‘shifting balance theory’ (SBT) was initially proposed by [Wright, 1932, Wright, 1982]. The theory described a particular mechanism for valley-crossing involving migration between small isolated populations, where selection is weak, and larger core populations, where selection is strong. However, a growing number of people have begun citing SBT in reference to a collection of phenomena that naturally affect selection and variation, such as range expansions [Johnson, 2008]. I support this extension of SBT, and further propose to include all the phenomena I describe above. For clarity, I will refer to the original narrow-sense SBT as ‘peripheral isolate theory’, and I will refer to the more general assertion, that dynamic selection strength drives adaptation, as shifting balance theory. SBT, defined this way, provides a robust and general mechanism for how natural populations can cross fitness valleys and keep the wheels of evolution turning.

1.2.6 Punctuated Evolution

The irregular rate of adaptation seen in nature is the subject of much discussion [Simpson, 1944, Fitch, 1995]. The fossil record shows a repeating, irregular pattern of long periods of stasis interrupted by rapid evolutionary change. This pattern was first called punctuated equilibrium (PE) by [Eldredge, 1972, Gould and Eldredge, 1993, Gould, 1977, Gould and Eldredge, 1977]. The same pattern has been seen in the observed variation of the speed of the molecular clock [Langley and Fitch, 1974, Gillespie, 1984a]. Episodic patterns of speciation have been inferred using different methods and on different taxonomic scales; from the fossils of the lower and middle Cambrian [Conway Morris, 1998], Phanerozoic marine fossils [Erwin et al., 1987, John Sepkoski Jr, 1998] and tracheophyte fossils [Niklas et al., 1983, Niklas, 1997], super-tree reconstructions of mammalian phylogenies [Bininda-Emonds et al., 2007], and from influenza sequence evolution [Wolf et al., 2006].

A number of mechanisms have been proposed to account for variable rates of evolution.

For instance, [Eldredge, 1972, Gould and Eldredge, 1993] originally proposed peripheral isolate theory [Gould, 1977, Wright, 1932, Wright, 1982] as a possible mechanism for PE. Sometimes new genetic mechanisms (or even mechanisms outside neo-Darwinian theory [Gould, 1987]) are invoked to explain this pattern, although the evidence for such mechanisms appears slim [Charlesworth et al., 1982]. PE has also been observed in populations that undergo periodic selection [Adami, 2006, Elena et al., 1996], and independently of whether adaptations lead to the exclusion or coexistence of species [Chow, 2004]. Changes in the environment leading to positive selection and the colonization of new niches have also been advanced as causes for accelerated evolution (the founder effect) [Mayr, 1954, Templeton, 1980, Carson, 1975]. Equally, small population sizes [Ohta, 1972, Lynch and Conery, 2003], changes in the population size [Takahata, 1987], catastrophic extinctions [Lewis, 1962], environmental stress [Hoffmann and Parsons, 1997, Nevo, 2001] changes in mutation rate [Lenski et al., 2003], and the introduction of new developmental or regulatory mechanisms [Erwin et al., 1987, Erwin and Valentine, 1984, Davidson and Erwin, 2006] as well as genetic instabilities [Fontdevila, 1992] and chromosomal evolution [Bush et al., 1977, Sites and Moritz, 1987] have all been implicated. All of the above points to a strong relationship between PE and SBT.

There are alternative hypotheses for what causes PE that does not involve SBT at all. One of the primary alternative hypotheses is neutral theory [Kimura, 1983, Kimura, 1991]. As discussed previously, neutral theory proposes that nearly all variation seen in nature is neutral variation. If true, a natural population may drift aimlessly for many generations before discovering a beneficial mutation. Thus, vast neutral landscapes can result in very long periods of stasis in between the discovery of beneficial mutations. This alone only accounts for the periods of stasis, and does not explain why beneficial mutations tend to cluster together to result in rapid evolutionary change. Further, evolutionary computation often displays punctuated equilibrium when evolving on fitness landscapes with far less neutrality, or none at all. Thus, while neutrality certainly can play a role in the punctuated equilibrium phenomenon, it cannot be the sole explanation.

Another hypothesis suggests that the evolutionary process exhibits ‘self-organized criticality’ (SOC) [Bak and Boettcher, 1997]. In brief, SOC applies when a system in disequilibrium moves towards equilibrium but stops at the ‘minimally stable state’ instead of progressing to the ‘maximally stable state’ [Bak et al., 1987]. This leads to the system perching on the edge of disequilibrium and thus being increasingly sensitive to further perturbations. The canonical example is that of a sand pile: the addition of a grain of sand represents may result in a perturbation from equilibrium for the entire pile, and the shifting of other grains restores the pile to a minimally stable state⁷. Sometimes more grains must shift to accomplish equilibrium, and if the system moves beyond the minimal stable state (to a more-than-minimally-stable state), there can be a prolonged state of stability where the pile is temporarily less sensitive to the addition of more grains.

The analogy with PE is easy to understand, since systems with SOC and evolution both exhibit unpredictable and sudden large changes. However, in order to map SOC to evolution, one must first explain what equilibrium the system is moving towards, what perturbations are moving the system away from equilibrium, how these perturbations accumulate a potential for change over time, what the minimally-stable state is, and why evolution stops there instead of proceeding to a more stable state. Until these questions are answered, SOC remains a less viable explanation for PE.

We can amend the SOC hypothesis slightly by first observing that environmental factors that may kick off evolutionary change can be self-organized. For example, earthquakes [Turcotte et al., 1985, Smalley Jr. et al., 1985] and forest fires [Malamud et al., 1998], among other things, are both thought to be self-organized. These kinds of events may create environmental instability, that then leads to evolutionary adaptation. Since these destabilizing events exhibit SOC, the adaptive response to them would result in the observed PE without evolution itself exhibiting SOC.

However, PE is observed in digital evolution all the time, and SOC does not appear in

⁷The maximally stable state is no pile at all: all grains of sand in a ground state with zero instability.

many of these systems. So, SOC cannot be the only explanation for PE, even if it can explain some spurts of evolutionary change. Shifting balance theory, in my opinion, remains the best contender to explain punctuated equilibrium.

1.2.7 Genetic Recombination

Hitherto now, I have avoided discussing processes by which two or more organisms can exchange genetic material with one another. Sexual reproduction and horizontal gene transfer are common examples of such an exchange. In evolutionary computation, it is common to use a simplified kind of sexual reproduction called crossover [Sastry et al., 2005]. Crossover produces an offspring whose genome is a combination of two or more parents' genomes. Although crossover is a crude analogy for sexual reproduction, I may nevertheless refer to 'sex' in an abstract sense when talking about crossover or genetic recombination generally.

When a population does not have the capability for genetic recombination, they are generally called asexual populations. Up to now, I have only discussed asexual populations. Adding sex to an asexual population creates several notable changes to the evolutionary dynamics.

First, sex allows the combination of two or more beneficial mutations that have arisen in separate lineages. This directly combats clonal interference by providing opportunities for competing benefits to come together and go to fixation instead of driving one another to extinction. Second, sex allows the combination of two or more deleterious mutations that have arisen in separate lineages. These doubly-deleterious individuals can be more easily purged by purifying selection, helping to slow or stop the spread of these mutations to future generations. Third, sex allows for the separation of beneficial mutations and deleterious mutations that have arisen in a single lineage. Offspring that receive only the beneficial mutation will outcompete organisms with the deleterious mutation, and offspring that receive only the deleterious mutation will be outcompeted by organisms that have the benefit. Like in the previous two cases, sex is assisting the beneficial mutations and hampering the deleterious mutations. All of these effects can be thought of as increasing the strength

of selection [Kondrashov, 1988, Rice and Chippindale, 2001], and we have already discussed the implications of increasing the strength of selection in previous sections.

Sex can also make exploration of the fitness landscape easier. Asexual populations traverse the fitness landscape through mutation alone, and valley-crossing is a major obstacle for asexual populations. Sex can join together mutations that might be difficult to join asexually. Consider the following example: Suppose we are evolving a population of organisms whose genomes are strings of four bits. The population begins with all organisms having genotype '0000'. The genotype '1111' is the global optimum, having a higher fitness than all other genotypes. However, all other genotypes suffer a fitness penalty for each '1' in their genome; there are four single-1 genotypes, six double-1 genotypes, and four tripple-1 genotypes, and each group suffers a larger fitness penalty than the last. If we assume that at most one site in an organism's genome can mutate in a single generation, then an asexual population must, at some point, have a tripple-1 genotype survive selection and produce a '1111' offspring to discover the global optimum. This may be an extraordinarily rare occurrence if the mutation rate is low and/or the fitness penalty for each '1' is severe. On the other hand, a sexual population could combine two genomes, each one a double-1 genotype, to produce a '1111'. Under the same circumstances, this could be a far more likely occurrence than the asexual path to the same genotype. Thus, sex can enhance the population's ability to explore the fitness landscape, and valley-cross.

Finally, although recombination produces novel combinations of existing genetic variants, it cannot create new genetic variants in the sense that mutation does. In other words, a population that evolves with recombination but without mutations can explore the fitness landscape, but exploration is limited to creating combinations of the existing genetic variation. As evolution progresses in such a system, genetic variation will be lost to selection, further limiting the combinations that can be produced. Without mutations to supply new variation, the population will eventually converge to a single genotype. The final result may not even be the optimal combination of the variation that was initially available in the population. Thus,

sexual recombination is neither necessary, nor sufficient for evolution to continue indefinitely.

Ultimately, sexual and asexual populations operate in largely the same way. Asexual populations experience weaker selection, and suffer from clonal interference, but more easily maintain genetic diversity. Sexual populations experience stronger selection and reduced diversity, but can mix and match existing mutations to create novel combinations and explore additional areas of the fitness landscape without waiting on the arrival of further mutations.

1.3 Research Objectives and Contributions

Many of the proposed causes of punctuated equilibrium are also phenomena I propose fit under the umbrella of a general shifting balance theory. This observation leads us to the core of my research: Can we leverage the natural phenomena, which are proposed to accelerate natural evolution, to augment digital evolution and evolutionary computation?

Since a general shifting balance theory predominantly involves phenomena that change the selection strength of the system, that is where I have focused my efforts. As such, it will be paramount to first understand selection in a more concrete way. How can we measure the strength of selection? I believe a good measure of the strength of selection must come from a good conceptualization of selection and its role in the evolutionary process. It also remains to be seen if we can understand why each of the phenomena that shift the strength of selection do so. Are these phenomena merely *like* changing the strength of selection or are they *actually* changing the strength of selection, and what (if any) distinction is there to be drawn between those two concepts? Answering these questions, again, comes down to first creating a better understanding of what selection strength is to begin with.

Many of the phenomenon that change the strength of selection are highly specific events or circumstances. For example, the original shifting balance theory dealt only with isolated populations and migration events. This mechanism is not expected to be a common occurrence, nor a primary driver of evolutionary change because of its rarity. Are there more universal phenomena that create a shifting strength of selection? How can we detect these

phenomena and rule out confounding factors? How plausible would these new hypothetical phenomena be as explanations of punctuated equilibrium?

Finally, how can we take answers to the questions above and turn them into solutions for evolutionary computation? Which phenomena are easily converted to new algorithms or best practices? What phenomena are difficult to leverage and why?

1.4 Research Methods

I use computational simulations to test both the development of new population genetics theory and any potential contributions to computational algorithms. Wherever applicable, existing population genetics theory is used to confirm or compare results generated by the computational simulations.

The MABE software package [Bohm et al., 2017] is one of the primary tools used to build the evolutionary algorithms used in my work. Whenever an evolutionary algorithm with little complexity was needed, I instead wrote a simple system in the python programming language. All data analysis and visualization was performed using python and assorted libraries. The details of each simulation are described in the corresponding chapters.

My collaborators and I have made a number of novel contributions to the research methods we use in the course of conducting the research. The contributions that appear in several chapters in this dissertation are described below. Other contributions are discussed in more detail in Chapter 6.

1.4.1 Fitness proportional selection without diminishing returns

We can take the standard fitness proportional selection algorithm (a.k.a. roulette wheel selection or roulette selection), with probabilities of reproduction given by

$$Pr(i) = \frac{w_i}{\sum_j^N w_j} , \tag{1.1}$$

and we can modify it to avoid diminishing returns at large fitness values. This is accomplished by ensuring that the same change in fitness results in the same fitness advantage. Mathematically, it is written as

$$Pr(i) = \frac{b^{w_i}}{\sum_j^N b^{w_j}} , \quad (1.2)$$

where b is a parameter to tune the strength of selection. This formulation of roulette selection no longer has diminishing returns, but it suffers from a computational problem, since b^w can easily reach the limit of what can be stored in most computing platforms. To fix this, we can offset the fitness values, working instead with values relative to the maximum fitness,

$$Pr(i) = \frac{b^{w_i - w_{\max}}}{\sum_j^N b^{w_j - w_{\max}}} . \quad (1.3)$$

In this form, the magnitude of $w_i - w_{\max}$ is bound by the range of fitness values $w_{\min} - w_{\max}$, helping to keep the magnitude of $b^{w_i - w_{\max}}$ small enough to hold in computer memory. Furthermore, the range of b^{w_i} is $(-\infty, \infty)$, while the range of $b^{w_i - w_{\max}}$ is $(-\infty, 1]$, ensuring that any loss of precision occurs at the smallest values, which has no great impact on the outcome of the simulations.

1.4.2 Saw-tooth fitness function

Measuring selection strength is hard (more about that in Chapter 2). In the absence of a more sophisticated methodology, selection strength can be inferred by the relative ease or difficulty of valley-crossing on a known fitness landscape. To this end, we have constructed a novel ‘saw-tooth’ fitness function that endlessly repeats the same fitness valley.

The saw-tooth fitness function requires organisms with a genome with a single locus that encodes a number k . When the organism reproduces, this allele mutates to $k + 1$ with probability $\frac{\mu}{2}$ or $k - 1$ with probability $\frac{\mu}{2}$, where μ is the probability to mutate. The saw-tooth function, shown in Figure 1.2, can be written in terms of the genotype k as

$$w(k) = Bn - Dm , \quad (1.4)$$




Figure 1.2: Diagram of the saw-tooth fitness function. The shape of the function is defined by the valley width v , crossing benefit B , and mutation disadvantage D . The score of a genotype k is calculated according to its position, given by n and m , multiplied by the benefit and disadvantage, respectively.

$$m = k \pmod{v} , \quad (1.5)$$

and

$$n = \frac{k - m}{v} , \quad (1.6)$$

where B is the additive fitness benefit of discovering each peak, n is the number of peaks to the right of the origin the genotype k has reached or passed, D is the additive fitness disadvantage of each mutation into a fitness valley, m is the number of mutations into the valley the genotype k has accumulated, and v is the width of the fitness valley, from peak to peak. The modified roulette selection probabilities (Equation 1.3) can then be used to ensure the selective forces are consistent everywhere in the saw-tooth fitness landscape.

The saw-tooth function allows us to relate the average rate of adaptation (fitness gain) to the average rate of valley-crossing, since each valley crossed awards the same fitness benefit. Further, since we know that fitness valleys are more easily crossed when selection is weaker, we can conceptually map the rate of valley-crossing to selection strength. These results are also easily compared with the theoretical expectations for an Ornstein–Uhlenbeck

process [Uhlenbeck and Ornstein, 1930, Arttime et al., 2018, Alili et al., 2005, Yi, 2010], by looking at the runtime of the simulation and the average number of valleys crossed, *i.e.*,

$$\bar{t}_{\text{cross}} = B \frac{t_{\text{final}}}{\bar{w}_{\text{final}}} . \quad (1.7)$$

1.5 Outline of Dissertation

The remaining chapters of this dissertation each explore a specific aspect of selection strength and how its changes affect evolutionary dynamics. Chapter 2 delves into the nature of selection itself, and explores the creation of a new metric of selection strength. The new metric of selection strength highlights important aspects of selection that will inform future attempts to improve evolutionary computation algorithms. Chapter 3 explores noisy fitness evaluation as a natural mechanism for reducing selection strength. Fitness noise is compared and contrasted with two other methods of controlling selection strength in an evolutionary system: population size and tournament selection. Chapter 4 introduces a previously undocumented evolutionary dynamic which affects the strength of selection of an evolving system. This newly described dynamic, called the free-for-all effect, reduces selection strength on populations during selective sweeps. The free-for-all effect also highlights that more than just weak selection is needed for efficient exploration of a fitness landscape. Chapter 5 builds on the discoveries of chapter 4 to develop a new method of augmenting evolutionary computation algorithms, further emphasizing the need for more than just weak selection. Chapter 6, the final chapter, summarizes the main results of each chapter and synthesizes a set of best practices for evolutionary computation based on these findings. The final chapter also includes a discussion of future research avenues and the benefits I believe are to be found by exploring them.

1.6 Summary

Evolutionary computation is a powerful optimization tool, and an invaluable test bed for population genetics. Although it is still debated how much evolutionary adaptation occurs via valley-crossing, it is undeniable that valley-crossing can provide faster avenues to higher fitness, or indeed the only avenue in some cases. Valley-crossing is fundamentally linked to the balance between the forces of selection and variation, and as such, controlling this balance is important for optimizing the efficiency of evolutionary algorithms. Nature, in contrast, is not actively optimized for performance, and yet nature seems to overcome many evolutionary obstacles that evolutionary algorithms do not. It is possible that nature benefits from a highly dynamic balance between selection and variation, and this constant flux helps natural populations avoid stagnation and overcome obstacles in the fitness landscape. Working with this hypothesis in mind, I investigate the nature of selection and how natural phenomena strengthen or weaken it. I will look at how to measure selection strength, explore how selection strength changes over evolutionary time, and construct new evolutionary algorithms and best practices based on my findings.

Chapter 2

Quantifying the impact of selection

2.1 Introduction

The effectiveness of an evolving population at navigating its fitness landscape is determined in large part by the strength of selection acting on it, and particularly so for deceptive fitness landscapes [Oliveto et al., 2018]. Evolving populations are subject to an explore-exploit trade-off [Wolpert and Macready, 1995, Wolpert and Macready, 1997, Ho and Pepyne, 2002, Millidge et al., 2021] that governs, among other things, how capable the population is at navigating the fitness landscape [Wright, 1932].

Strong selection restricts exploration [Feller, 1951, Kimura, 1964, Iwasa et al., 2004a, Covert et al., 2013, Oliveto et al., 2018, Lewis, 1962] and enhances exploitation by allowing only the most fit organisms to reproduce. This increases the probability that a new beneficial mutant will fix in the population, and also decreases the expected time to fixation, but reduces the probability to discover a new benefit in the first place.

Weak selection, on the other hand, allows many more less-fit organisms to reproduce, enhancing exploration at the cost of exploitation. This increases the probability to discover a new beneficial mutation, while reducing the probability to fix the mutation once it is discovered, and increasing the time to fixation. The most adaptive populations are those

that happen to balance exploration and exploitation in a way that aligns with the (local) topology of their fitness landscape [Ragusa and Bohm, 2021].

In spite of its important role in evolutionary theory, selection strength is not a fundamental (independent, or irreducible) property of evolving populations. Instead, it is an emergent property affected by many aspects of the evolving population, organism, and environment. For example, the following features all *simultaneously* affect the selection strength of an evolving system in their own way: population size [Jain et al., 2011, Ochs and Desai, 2015, Rozen et al., 2008, Ohta, 1972] and changing population sizes [Wright, 1982, Lynch and Conery, 2003], noisy genotype-to-fitness actualizations [Wang and Zhang, 2011, Melbinger and Vergassola, 2015, Van Egeren et al., 2018], temporospatial population structures such as range expansions [Slatkin and Excoffier, 2012, Peischl et al., 2013, Peischl and Excoffier, 2015, Peischl et al., 2015, Gilbert et al., 2017] and population topology [Mühlenbein, 1991, Lieberman et al., 2005, Mühlenbein, 2009, Askari and Samani, 2015, Möller et al., 2019, Tkadlec et al., 2019, Kuo et al., 2021], mass-extinction events and bottlenecks [Lewis, 1962, Mathias and Ragusa, 2016, Engholdt and Mathias, 2021], sexual selection [Bohm et al., 2019], and many more. While selection strength does not have a singular cause, quantifying the effective strength of selection can be extremely useful for designing experiments that aim to measure the effect a particular feature has on selection strength, understanding the present state of a population of interest, or for comparing populations with one another.

As one might expect, given the importance of selection in evolutionary theory, there is no lack of methods that already exist for quantifying different aspects of selection strength. These methods, as we will see, each have their drawbacks: some make too many assumptions and are either unnecessarily restrictive or simply not appropriate in all cases, some make too few assumptions and are consequently less meaningful, some have units that make comparing populations difficult, and other measures can only be computed relative to another population, which prohibits their use in some common cases.

In this work, I aim to build a theory of selection strength from a set of general assumptions.

I provide intuition and justification for each of the assumptions before proceeding to formalize them mathematically. I then present a new measure of selection strength, the *selection impact metric*, and demonstrate its use on a number of digital evolution experiments.

2.2 Properties of selection strength

Selection creates a bias in the reproductive success of genotypes in a population. Here, bias refers to a *non-uniform probability* of producing an offspring. Changes in the frequency of a trait are the result of a correlation between the trait and the reproductive success of the organism with that trait [Price et al., 1970, Price, 1972]. However, selection does not “intend” to change trait frequencies; a changing trait frequency emerges from the bias in reproductive success. Therefore, changes in trait frequencies, while correlated with selective forces, are not where we should be looking to measure selection. Instead, we will only consider offspring production when quantifying selection strength.

If selection is a bias in reproductive success, then no bias means no selection. This implies that a measure of selection strength should indicate zero in the case of neutral drift. In addition to satisfying our common sense, identifying neutral drift with the absence of selection has the additional benefit of defining an absolute zero for a metric of selection strength. A population consisting of a single organism also has no bias in reproductive success, since the probability of allocating offspring to among one individual is trivially uniform. This informs us that selection is inherently a property of a population, so our measure of selection strength will reflect this.

It is important to acknowledge that the selection strength experienced by a population can change over time. Indeed, the bias of reproductive success changes with each birth and death. Ideally, a measure of selection strength would be able to resolve changes in selection strength over time, as well as space.

A metric of selection strength should allow for comparison between different species. If, for instance, one wants to compare the selection pressure experienced by a population of beetles

to that experienced by a population of birds, the metric should facilitate such a comparison. However, some metrics of selection strength, such as those based on the correlation between an organism’s fitness and traits [Price et al., 1970, Price, 1972], are not capable of making such comparisons as they are dependent on the units used for the traits and fitness [Lande and Arnold, 1983, Houle, 1992, Hereford et al., 2004]¹. There are some ad-hoc solutions that attempt to standardize units, but there is a lack of consensus on which solution to use [Matsumura et al., 2012]. One way to avoid this problem is to use a unitless metric of selection strength, that would allow for cross-species comparison of the selection pressure.

The following list summarizes the desired properties of a selection strength metric.

- (F) Selection strength is calculated from observing only the number of offspring that organisms produce.
- (Z) Selection strength has an absolute zero: when the population experiences neutral drift.
- (P) Selection strength describes a force felt by a population, a collection of more than one organism.
- (C) Selection strength can be compared between two populations, regardless of their relationship to each other.
- (U) Selection strength is a unitless quantity.

I will reference these constraints in the following sections to both motivate the mathematical development of a new metric, and to call attention to when the new metric satisfies these requirements.

¹For example, in a population of beetles, a correlation may be found between the walking speed measured in inches per second and the fitness measured in number of offspring. Similarly, in a population of birds, a correlation may exist between wingspan measured in inches and number of offspring. However, it is uncertain how to compare the units of $\frac{[\text{inches}]}{[\text{second}] \times [\text{offspring}]}$ and $\frac{[\text{inches}]}{[\text{offspring}]}$.

2.3 The Selection Impact Metric

The core concept behind the selection impact metric is that selection creates a bias in the distribution of offspring among the available parents in a population. Conversely, there is an unbiased distribution of offspring that corresponds to neutral drift.

We can represent a population's *offspring distribution* as a probability mass function, $f(n)$ (defined in Section 2.5), that indicates what fraction of the population's parents produced n offspring during a particular period of observation.

The unbiased *drift distribution*, $f_D(n)$ (defined in Section 2.4), is a theoretical offspring distribution that represents the most likely distribution of offspring when the probability of reproduction is exactly $\frac{1}{N}$ for each organism in a population with N potential parents.

With both f and f_D in hand, we can measure how far from drift the population is and define the selection impact as

$$S(f) = \text{Dist}(f, f_D) , \tag{2.1}$$

where Dist is a suitable distance function between two probability mass functions.

Because the only data about the population required to compute the selection impact is the distribution of offspring f , the definition of selection impact as Equation (2.1) satisfies the constraint (F), that we only consider offspring counts. Further, by defining the selection impact as a distance from f_D , the constraint (Z), that the metric is zero in the case of neutral drift, is also satisfied, since $S(f) = 0$ implies $f(n) = f_D(n)$. The constraint (P), that S describes the impact of selection on a population, is satisfied implicitly by the use of probability mass functions (a population with $N = 1$ will always have $f(n) = f_D(n)$).

The remainder of Section 2.3 addresses the choice of distance function to be used in Equation (2.1). Section 2.4 describes the derivation of f_D from a minimal set of assumptions about the population, and Section 2.5 describes how to construct f from an observation of the population. Section 2.6 presents empirical data that verify the drift distributions, and showcase some applications of $S(f)$. In Section 2.7, I discuss where the selection impact fits

into the ecosystem of other selection strength measures and some future research directions.

2.3.1 Choosing a suitable distance measure

The ability to compare the selection impact of two observations, constraint (C), is satisfied if Dist satisfies the triangle inequality. Likewise, the constraint (U), that the selection impact is unitless, is satisfied if in turn Dist is unitless. Therefore, the choice of distance function is limited to a set of unitless metrics that operate on probability mass functions.

We have a choice between information theoretic measures of difference between probability masses, like the Kullback-Leibler (KL) divergence [Kullback and Leibler, 1951], or a true measure of distance between probability masses, like the Earth-Mover distance (EMD, a.k.a. the Wasserstein distance or the Kantorovich–Rubinstein metric) [Kantorovich, 1960] or the Cramér-von Mises distance (CMD) [Anderson, 1962]. We can rule out the KL-divergence and similar information theoretic measures because they treat the discrete elements of the distribution’s support as unordered symbols, instead of ordered values. For our application — measuring the difference of two offspring distributions — we care about the distance between offspring counts, not merely that two counts are different symbolically. A parent of two offspring should be considered quite different from a parent of thirty-two, while considered quite similar to a parent of three. Therefore, we will choose a proper distance measure over an information theoretic measure.

The choice between EMD and CMD is somewhat arbitrary, though the CMD has a diminishing sensitivity to large differences along the support of the probability mass function. Therefore, I have chosen to employ the EMD as the distance metric in Equation (2.1) for the remainder of this report.

2.3.2 Earth Mover Distance

For two probability distributions $P(x)$ and $Q(x)$ with a common one-dimensional interval support, $\chi = [x_0, \dots, x_n]$, the EMD is given by the following recursive definition:

$$r_{i+1} = P(x_i) + r_i - Q(x_i) , \quad (2.2)$$

$$r_0 = 0 ,$$

$$\text{EMD}(P, Q) = \sum_{i=1}^n |r_i| . \quad (2.3)$$

If $P(x)$ and $Q(x)$ have different supports, then the EMD can be calculated using the smallest interval that includes both, $\chi = [\min(\chi_P, \chi_Q), \max(\chi_P, \chi_Q)]$.

2.4 Modeling the Offspring Distributions of Neutral Drift

In order to calculate the selection impact metric, it is necessary to have both the actual offspring distribution of a population and a theoretical offspring distribution for a population of the same size that is evolving under neutral drift. These drift reference distributions model the most likely distribution of offspring among the parents in a neutrally drifting population.

The neutral drift expectation depends on several key factors: the generational model (*e.g.*, Wright-Fisher vs. Moran, Sections 2.4.1 & 2.4.2), population size (Sections 2.4.1.2 & 2.4.2.1), and sexual recombination (Section 2.4.1.1).

2.4.1 Non-Overlapping Generations

In a Wright-Fisher process, time advances in discrete intervals called generations. A generation consists of a reproduction step and a replacement step. During the reproduction step, all members of the population are given the opportunity to create an offspring. The probability that a member of the population is selected as a parent is proportional to their fitness. Parents are selected until a number of offspring equal to the population size are produced. The same organism may be selected as a parent more than once during the reproduction step. Next, in the replacement step, all members of the population are replaced by the offspring. Since no parent ever exists in the population at the same time as their children, the generations

are called 'non-overlapping'.

Let N denote the population size (e.g., $N \in \mathbb{Z}^+$). When the population experiences neutral drift, each organism has a probability of $\frac{1}{N}$ to be chosen to reproduce, and there are N opportunities to be selected. We can model the probability that any organism has exactly n offspring when the population drifts by a simple binomial probability distribution

$$f_D(n; N) = \binom{N}{n} \left(\frac{1}{N}\right)^n \left(1 - \frac{1}{N}\right)^{N-n}, \quad (2.4)$$

where $0 \leq n \leq N$ and $n \in \mathbb{Z}^+$. Because the population must be replenished with offspring each generation, the number of trials is N . Hence, our binomial distribution is given by $f_D(n; N)$ and is parameterized only by N .

2.4.1.1 Sexual reproduction

Let p denote the number of parents needed to produce a single offspring. The most common biological examples would be asexual ($p = 1$) and sexual ($p = 2$) offspring production, however, p can be set arbitrarily high to accommodate other scenarios.

When the population experiences neutral drift and each offspring is formed from p parents, we can model the probability that any organism has exactly n offspring as follows:

$$f_D(n; N, p) = \binom{pN}{n} \left(\frac{1}{N}\right)^n \left(1 - \frac{1}{N}\right)^{pN-n}, \quad (2.5)$$

where $0 \leq n \leq pN$ and $n \in \mathbb{Z}$. These modifications to Equation 2.4 take into account that the probability to select from the initial population² is $\frac{1}{N}$, and that there are pN opportunities to be selected, since we need p parents for each of the N offspring. These modifications can be made in conjunction with those in Section 2.4.1.2.

²Here I assume for mathematical simplicity that selection is with replacement. Although this implies that one individual can sexually recombine with itself, the probability of such an occurrence diminishes if larger populations are considered.

2.4.1.2 Changing population size

Let us now relax the assumption of constant population size.

Let $N^{(t)}$ denote the initial population size at time t , and $N^{(t+1)}$ denote the population size at time $t + 1$.

When the population experiences neutral drift and the population size changes from $N^{(t)}$ to $N^{(t+1)}$, we can model the probability that any organism has exactly n offspring as

$$f_D(n; N^{(t)}, N^{(t+1)}) = \binom{N^{(t+1)}}{n} \left(\frac{1}{N^{(t)}} \right)^n \left(1 - \frac{1}{N^{(t)}} \right)^{N^{(t+1)}-n} \quad (2.6)$$

where $0 \leq n \leq N^{(t+1)}$ and $n \in \mathbb{Z}$. These modifications to Equation 2.4 take into account that the probability to select from the initial population is $\frac{1}{N^{(t)}}$, and there are $N^{(t+1)}$ opportunities to be selected. These modifications can be made in conjunction with those in Section 2.4.1.1.

2.4.2 Overlapping Generations

When generations are allowed to overlap, the evolutionary dynamics are best described by the Moran process. Time advances in a series of updates, where one organism gives birth, and one organism dies³. The probability to select an organism as a parent is proportional to their fitness. The probability to select an organism to die is uniform (*i.e.*, $\frac{1}{N}$). Since a parent can exist in the population at the same time as their descendants, the generations are called 'overlapping'.

Because of the coexisting generations, a model of neutral drift must track ancestor-descendant relationships. To facilitate this ancestry tracking, we must label each organism at time $t = 0$ with a unique integer i from the set $1, \dots, N$, which indicates the organism's clade. When an organism reproduces, their offspring inherit the clade identifier.

Let $n_i^{(t)}$ denote the number of organisms in clade i at time t . Since the value of $n_i^{(t)}$ can be any integer in $\{0, \dots, N\}$, we can model the change of n_i over time with a Markov chain

³Here, I assume the birth-death model, where births happen before deaths. If death happens before birth, the organism that dies cannot reproduce and the values of $p_{a,b}$ change accordingly.

with states $n \in \{0, \dots, N\}$. When the population experiences neutral drift, each organism has a probability of $\frac{1}{N}$ to be chosen as a parent. Further, each organism always has a probability of $\frac{1}{N}$ to be selected for death. Since the same organism can be chosen for both events (and each event is independent) there are four cases:

1) An organism in clade i is selected for birth and an organism outside of clade i is selected for death.

2) An organism in clade i is selected for birth and an organism in clade i is selected for death.

3) An organism outside of clade i is selected for birth and an organism outside of clade i is selected for death.

4) An organism outside of clade i is selected for birth and an organism in clade i is selected for death.

Therefore, the Markov chain has a $(N + 1) \times (N + 1)$ transition matrix $\mathbf{P} = [p_{a,b}]$ where $p_{a,b}$ is

$$p_{a,b} = \begin{cases} \left(\frac{n}{N}\right) \left(1 - \frac{n}{N}\right) & \text{if } a = n \text{ and } b = n + 1 \\ \left(\frac{n}{N}\right)^2 + \left(1 - \frac{n}{N}\right)^2 & \text{if } a = n \text{ and } b = n \\ \left(1 - \frac{n}{N}\right) \left(\frac{n}{N}\right) & \text{if } a = n \text{ and } b = n - 1 \\ 0 & \text{otherwise .} \end{cases} \quad (2.7)$$

While \mathbf{P} naturally depends on the population size, I will suppress this dependence from now on. Readers familiar with the Moran process will notice that Equation (2.7) describes the same probabilities as those for tracking allele frequencies under neutral drift. This is because we are effectively assuming there are N distinct alleles (the clade labels), where the allele frequency is synonymous with the clade size.

In order to compute the drift distribution with this Markov chain, we must define the population's starting state and iterate the Markov chain the same number of times the

population has been updated. Let $\vec{x}^{(t)}$ denote a vector in \mathbb{R}^{N+1} , where $x_n^{(t)}$ represents the fraction of clades with size n at time t and $\vec{x}^{(0)}$ denote the starting condition, determined by our initial labeling of each organism with a unique clade ID. For example, if $N = 3$, then $\vec{x}^{(0)} = [0, 1, 0, 0]$, since 100% of the population has a clade of size $n = 1$ when $t = 0$.

Let u denote the number of updates that have passed since the start of the observation. The fraction of clades that become size n after u updates of drifting neutrally is given by the probability density function

$$f_D(n; N, u) = \vec{x}^{(u)} , \quad (2.8)$$

where $\vec{x}^{(t)}$ is defined recursively as

$$\vec{x}^{(t)} = \vec{x}^{(t-1)} \mathbf{P} . \quad (2.9)$$

2.4.2.1 Changing population size

Since Equation (2.7) is a function of N , any accounting for a changing population size will require a new transition matrix each time N changes. However, if we only allow the population size to change by ± 1 each update from $t \rightarrow t + 1$, and we allow at most one birth and one death to occur on any particular update, we can limit the number of extra transition matrices needed to just two.

If the population size has increased by one, then one of two events has occurred: either an organism in clade i has had an offspring, or an organism outside of clade i has had an offspring, but in either case no clade has experienced a death. Therefore, the $(N+1) \times (N+2)$ transition matrix is given by:

$$p_{a,b} = \begin{cases} \left(\frac{n}{N}\right) & \text{if } a = n \text{ and } b = n + 1 , \\ \left(1 - \frac{n}{N}\right) & \text{if } a = n \text{ and } b = n , \\ 0 & \text{otherwise .} \end{cases} \quad (2.10)$$

Similarly, if the population size decreases by one, we have a $(N + 1) \times N$ transition matrix given by:

$$p_{a,b} = \begin{cases} \left(1 - \frac{n}{N}\right) & \text{if } a = n \text{ and } b = n , \\ \left(\frac{n}{N}\right) & \text{if } a = n \text{ and } b = n - 1 , \\ 0 & \text{otherwise .} \end{cases} \quad (2.11)$$

In practice, finding the drift distribution for a population that changes size should be done by modifying Equation (2.9) to no longer suppress the dependence of P on the population size. For instance,

$$\vec{x}^{(t)} = \vec{x}^{(t-1)} \mathbf{P}(N^{(t-1)}, N^{(t)}) , \quad (2.12)$$

with the understanding that \mathbf{P} is defined by Equation (2.7) when $N^{(t-1)} = N^{(t)}$, Equation (2.10) when $N^{(t-1)} + 1 = N^{(t)}$, and Equation (2.11) when $N^{(t-1)} - 1 = N^{(t)}$. This shows that the particular sequence of births and deaths observed for a population will determine the expected neutral-drift offspring distribution.

2.5 Counting Offspring and Descendants

This section explains the process of gathering the data required to apply the selection impact metric. The method for counting offspring in the population is consistent across different reproduction dynamics, unlike the drift reference distributions outlined in Section 2.4, which can vary greatly.

2.5.1 Non-Overlapping Generations

In a Wright-Fisher process, all children are born simultaneously in time intervals called a generation. With t denoting a particular generation, and $t + 1$ denoting the subsequent generation, let $Z^{(t,t+1)} = [z_{i,j}^{(t,t+1)}]$ be the $N \times (N + 1)$ offspring production matrix for the

$t \rightarrow t + 1$ generation, where the element $z_{i,j}^{(t,t+1)}$ is given by:

$$z_{i,j}^{(t,t+1)} = \begin{cases} 1 & \text{if agent } i \text{ had } j \text{ offspring during } t \rightarrow t + 1, \\ 0 & \text{otherwise.} \end{cases} \quad (2.13)$$

In order to be compared with a drift distribution, the offspring distribution for the $t \rightarrow t+1$ generation, $f(n, Z^{(t,t+1)})$, must return the fraction of parents in the population that had n offspring. This is obtained by taking the mean of the n th column of $Z^{(t,t+1)}$ as follows:

$$f(n, Z^{(t,t+1)}) = \frac{1}{N} \sum_{i=0}^N z_{i,n}^{(t,t+1)}, \quad (2.14)$$

where N is inferred from the number of rows in Z .

In each of the following sections, Equation (2.14) is used to convert the offspring production matrix into a probability mass function.

2.5.1.1 Sexual reproduction

To account for sexual reproduction, each offspring that a parent contributes genetic material to is recorded as an offspring for that parent. Therefore, the offspring production matrix $Z^{(t,t+1)}$ becomes an $N \times (pN + 1)$ matrix, and each element $z_{i,j}$ is defined as in Equation (2.13).

2.5.1.2 Changing population size

To account for a population that changes size from $N^{(t)}$ to $N^{(t+1)}$, the offspring production matrix $Z^{(t,t+1)}$ becomes an $N^{(t)} \times (N^{(t+1)} + 1)$ matrix, and each element $z_{i,j}$ is defined as in Equation (2.13).

2.5.2 Overlapping Generations

In a Moran process, one child is born, and one organism dies in a time interval called an update.

With t denoting time, $t = 0$ denoting the initial observation, and $t = u$ the final observation, let $Z^{(u)} = [z_{i,j}^{(u)}]$ be the $N \times (N + 1)$ clade size matrix after u updates have occurred,

where the element $z_{i,j}^{(u)}$ is given by:

$$z_{i,j}^{(u)} = \begin{cases} 1 & \text{if clade } i \text{ contains } j \text{ living organisms at time } t = u \\ 0 & \text{otherwise} \end{cases} \quad (2.15)$$

2.5.2.1 Changing population size

To account for a population that changes size from $N^{(0)}$ to $N^{(u)}$, the offspring production matrix $Z^{(u)}$ becomes an $N^{(0)} \times (N^{(u)} + 1)$ matrix, and each element $z_{i,j}$ is defined as in Equation (2.15).

2.6 Empirical data

2.6.1 Drift model verification

We will first test the validity of the drift distributions described in Section 2.4 by observing a population undergoing neutral drift and recording its offspring distribution (Equation (2.14)). To this end, I simulate a population of 100 organisms that reproduce either according to a non-overlapping generational model (Figure 2.1A and C) or an overlapping generational model (Figure 2.1B). The population evolves asexually (Figure 2.1A and B) or sexually (Figure 2.1C). Every organism has the same fitness to ensure neutral evolution. In all cases, to compensate for the inherent variability among different measures of $f(n)$ under the same selection pressures, 5,000 replicate distributions are averaged into one offspring distribution, $\hat{f}(n)$. Since $f_D(n)$ represents the maximum-likelihood drift distribution, $\hat{f}(n)$ and $f_D(n)$ should be identical distributions.

Each panel of Figure 2.1 shows the drift distribution $f_D(n; N)$ (Equation (2.4)) and the average offspring distribution $\hat{f}(n)$ (Equation (2.14)) with a 99% confidence interval. Since our simulated population is drifting neutrally, we expect to see agreement between the theoretical reference and the observation, and we do. We can observe that the differences in reproduction model and number of parents change the shape of the drift offspring distribution.

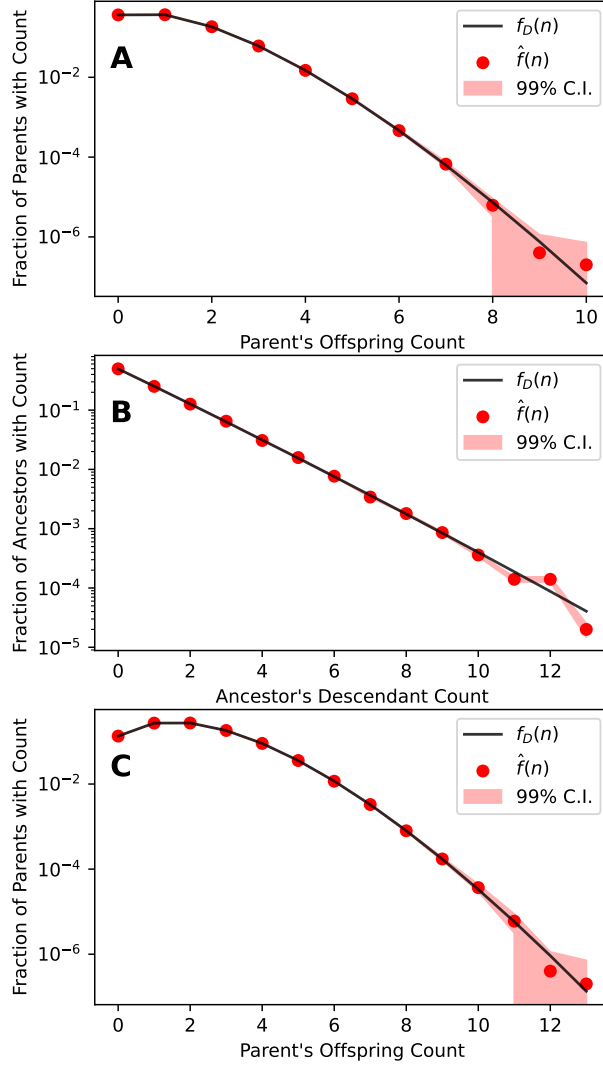


Figure 2.1: Average of 5,000 replicate offspring distributions $\hat{f}(n)$ for an *in silico* population of 100 organisms, drifting neutrally (red dots). 99% confidence interval of $\hat{f}(n)$ (red shaded area). Maximum-likelihood drift distribution $f_D(n)$ (black line). A) Wright-Fisher process with asexual reproduction ($p = 1$). B) Moran process with asexual reproduction ($p = 1$). C) Wright-Fisher process with sexual reproduction ($p = 2$).

2.6.2 Measuring the selection impact over time

In this section, I show how the selection impact metric can be used to track population dynamics over time. To do this, I collect offspring production data from a simulated population of 10,000 organisms, evolving with asexual non-overlapping generations. As opposed to the neutral control experiments I discussed in the previous section, here I will study the adaptation of a population to a fitness function with deceptive fitness valleys (see Section 2.8,

Figure 2.5). The deceptive nature of the fitness landscape provides us with the opportunity to collect data of a population undergoing stabilizing selection when trapped on a local optimum, as well as directional selection once a valley is crossed.

Figure 2.2 shows a single evolutionary history of a simulated population of 10,000 individuals using a mutation rate of $\mu = 0.005$ on the saw-tooth fitness landscape (see Section 2.8.1, Figure 2.5). Panel A shows both the average population fitness (blue line) and the highest fitness in the population (red line) for each generation. We observe periods of stasis where no change of fitness occurs, as well as periods of adaptation where beneficial mutations are discovered and fixed. Panel B shows the selection impact for each generation as a function of time. Recall that the selection impact, $S(f) = \text{EMD}(f, f_D)$, is computed as the earth-mover distance between the population’s offspring distribution $f(n, Z)$ and the drift offspring distribution for that population, $f_D(n; N)$. The selection impact is small but non-zero ($S(f) \approx 0.003$) during periods of stasis. The stasis periods correspond to stretches of time when the population has converged onto a local optimum, so stabilizing selection is acting to remove deleterious mutations in those periods. Conversely, during periods of adaptation, the selection impact increases sharply ($S(f) \approx 0.05$) as the population enters a selective sweep (directional selection) following the discovery of a beneficial mutation. The selection impact remains elevated only as long as the beneficial mutation is fixing. Once the population converges on a new local optimum, the selection impact returns to the lower value.

2.6.3 Effect of population size on the selection impact

Since population size is one of the main drivers of selection strength, it is important to study the dependence of the selection impact metric as population size changes.

In order to measure how population size affects the selection impact, I would like to test the metric on populations experiencing drift, directional selection, and stabilizing selection. I simulate the evolution of a population of organisms on five different fitness functions. The first fitness function, $w(x) = 1$, assigns all organisms the same fitness, and represents the




Figure 2.2: Fitness and selection impact of a population evolving via a Wright-Fisher process. A) The black line represents the average fitness and the red line represents the maximum fitness of each generation. The population is navigating a deceptive fitness landscape (details provided in Section 2.8.1). B) The selection impact metric of each generation of the same population.

neutral control. The functions $w(x) = x$ and $w(x) = x^3$ describe landscapes with less and more extreme gradients, respectively. The final two functions, $w(x) = -|x|$ and $w(x) = -x^2$, have a global optimum, though the fitness cost of a deleterious mutation is more extreme in the latter. For a visual representation of these functions, see Figure 2.5.

Figure 2.3 shows the average selection impact of a population when evolved on each of the five fitness function and with each of ten population sizes between 10 and 10,000. In all cases, the mutation rate is $\mu = 0.125$ (see Section 2.8). As the population size increases along the x -axis, the average selection impact of the population on the neutral landscape decreases rapidly towards zero. For all other fitness functions, the average selection impact converges to a non-zero constant.

Figure 2.3 shows that the selection impact for the hill-climb functions increases slightly as the population size increases, while the impact on the two functions with local optima decrease slightly. These small changes before convergence tell us something about the state of the population relative to the fitness landscape. The increase of the selection impact with population size (small circles and triangles) indicates that selection was acting inefficiently due to the small population. As the population size increases, more mutations become available in a single generation, which benefits hill-climbing. Conversely, a decreasing selection impact




Figure 2.3: The effect of population size on the measured selection impact of a population experiencing neutral drift (large circles), stabilizing selection (stars, crosses), and directional selection (triangles, small circles). Each data point corresponds to the mean of 5,000 generations of evolution on the respective fitness function. 99% confidence intervals are shown as a red-shaded area.

with increasing population size (stars and crosses) on fitness peaks suggests that small-population-size effects are creating a larger bias in the offspring distribution than would ordinarily occur in a larger population under the same selective pressures. That is to say, random fluctuations in offspring counts have a meaningful impact on small populations, a phenomenon that is already well known [Kimura and Ohta, 1969].

2.6.4 Mutation rate affects the selection impact

Because mutation rate is another factor known to affect the adaptation of evolving populations, I simulate the evolution of 500 organisms at different mutation rates on each of the five fitness functions listed above (described in Section 2.8).

Figure 2.4 shows the average selection impact of a population when evolved with mutation rates between 10^{-5} and one per organism per generation. When the mutation rate is very low, the selection impact is also low due to the lack of persistent non-neutral phenotypic variance; the population is effectively homogeneous most of the time, so selection cannot differentiate between individuals until mutations arise. As the mutation rate increases, the inflow of

beneficial and deleterious mutations also increases. The increased inflow of mutations creates more contrast between individuals, allowing selection to act more effectively, as is reflected in the increasing selection impact.

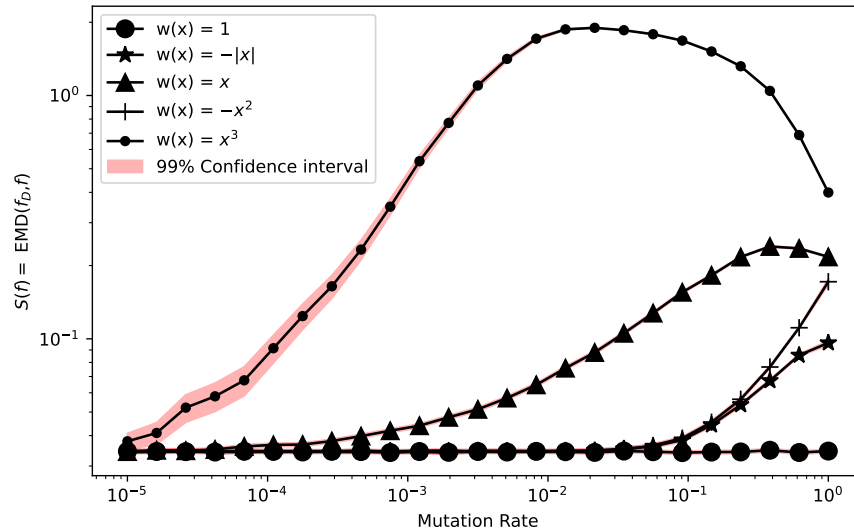


Figure 2.4: Effect of mutation rate on the measured selection impact of a population of 500 organisms experiencing neutral drift (large circles), stabilizing selection (stars and crosses), and directional selection (triangles and small circles). Each selection impact shown is a mean of 5,000 generations under the respective fitness function. Each 99% confidence interval is shown as a red-shaded area.

Figure 2.4 shows that the selection impact of populations on local-optima functions begin to increase only in response to relatively high mutation rates. At low mutation rates, selection keeps the sub-population of suboptimal mutants small, and prevents the emergence of organisms with multiple mutations. But, as we increase the mutation rate, a larger proportion of parent organisms will carry a deleterious mutation, and thus children with multiple mutations become more common too. The increased disparity between the highest- and lowest-fitness organisms explains the increased selection impact. The selection impact rises as the distribution of offspring becomes increasingly biased in favor of organisms closer to the optimum.

Conversely, Figure 2.4 shows that the selection impact of populations on a hill-climb function will increase in response to lower mutation rates. The hill-climb functions cannot be

optimized, in the sense that there is no optimal genotype. As a result, a genotype with even higher fitness can always be discovered. The mutation rate controls the discovery rate of new beneficial mutations, so as the mutation rate increases, we observe an increased selection impact that results from the increased frequency of selective sweeps. However, Figure 2.4 shows a critical mutation rate where the selection impact begins to trend downward. This reversal signals a change in the evolutionary dynamics. At sufficiently high mutation rates, there may be many organisms that have the highest fitness (or nearly the highest), but are not very competitive with one another (*i.e.*, clonal interference) [Desai and Fisher, 2007]. The clonal interference results in a more uniform allocation of offspring among those top-ranking individuals, and thus a lower selection impact.

2.7 Discussion

Beginning with a set of goals, outlined in Section 2.2, I have developed a metric for the impact that selection has on a population. The metric is calculated by comparing what we expect to occur when a population undergoes neutral drift with observations of evolution. Neutral drift is well understood and can be accurately characterized mathematically, and this allows us to construct a robust mathematical model of drift. I have shown that this mathematical model of drift correctly predicts the offspring distribution of populations drifting neutrally. I have further shown that the magnitude of the selection impact metric is sensitive to the occurrence of selective sweeps and stabilizing selection. I have provided resources on how to calculate the selection impact for a variety of circumstances, including overlapping and non-overlapping generations, asexual and sexual reproduction, constant and changing population sizes, as well as within and between sub-populations. The selection impact metric satisfies each of the requirements discussed in Section 2.2, and provides a novel way of quantifying the way selection acts on a population.

2.7.1 Comparing the selection impact to other measures of selection

The selection impact metric was designed from the start to only consider the offspring distribution of a population. This not only ensures the metric is looking at the tell-tale signature of selection acting on a population, but it also ensures that no assumptions are made about the mode or form of selection.

The mode of selection may be thought of as a kind of selection algorithm that actually makes the decision about which organisms reproduce and how many offspring they have. For example, truncation selection is a mode of selection where some percentage of the population dies, with those organisms that are the least fit dying first, and the surviving organisms reproduce with equal probability. Measures of selection strength have been formulated with this mode of selection as a foundational assumption [Haldane, 1954, Van Valen, 1965, Van Valen, 1967]. However, the measures of selection strength that assume truncation selection have been shown to under-estimate the selection strength [O’Donald, 1968]. The selection impact metric does not assume any particular mode of selection, it merely examines the after-effect of selection, the distribution of offspring.

Other measures of selection strength assume a particular relationship between fitness and traits. In particular, [O’Donald, 1968] assumes a quadratic relationship, $w(x) \propto (\theta - x)^2$, where θ is the optimal phenotype. While certainly this approximation will hold in some cases, it cannot hold in all cases, and it cannot hold when the population is in the midst of selective sweep. My selection impact metric is not limited by any such assumptions.

Like my selection impact, the Price equation [Price et al., 1970, Price, 1972] also does not assume a particular relationship between traits and fitness. However, the Price equation and extensions to it [Lande and Arnold, 1983] have received a number of criticisms that cannot go unmentioned. The ability to measure selection strength with the Price equation hinges on performing regression on the relationship between fitness and traits. However, this results in a relationship with units, namely those of the trait and of the fitness measure, and

comparisons between values with conflicting units are nonsensical [Houle, 1992, Hereford et al., 2004, Matsumura et al., 2012]. Although *ad hoc* mathematical corrections are possible, there is debate about which is the most appropriate correction to apply, if any at all. My selection impact metric, in contrast, is unitless, and so any two measurements can be compared without the need to correct for units.

2.7.2 Interpretation of the selection impact

I have referred to the metric introduced in this report as the ‘selection impact’ (as opposed to selection *strength*) because it illustrates the extent to which selection is *actually* affecting the population. ‘Selection strength’ can be thought of as a *potential* for adaptation. The selection impact does not necessarily indicate the power of this potential, but rather the effect it has on the population.

In particular, the selection impact is a function of the movement of the population through the fitness landscape, irrespective of where or how the population is moving. Further, the selection impact reflects how selection moves the organisms in the population, as opposed to the population’s center of mass. This is why in the case of stabilizing selection, we see a non-zero selection impact even though the population is in equilibrium (*e.g.*, Figure 2.2).

If there is no potential for selection (*i.e.*, no selection strength) or there is no non-neutral phenotypic variation in the population, neutral drift occurs. Therefore, the selection impact is proportional to the product of these two qualities, $S(f) \propto (\text{Pop.Var.}) \times (\text{Sel.Str.})$.

2.7.3 On the Practicality of the selection impact metric

Calculating the selection impact requires detailed information about the population (each individual’s offspring production) that is often difficult to obtain. If generations overlap, we further have to keep track of each organism’s ancestry. This level of detail is manageable for computer simulations, but rarely for field work.

Another limitation of the selection impact metric is that it cannot connect selection to specific traits, unlike other measures of selection in the literature [Haldane, 1954, Van Valen,

1965, O'Donald, 1968, Price et al., 1970, Lande and Arnold, 1983]. However, both of these limitations can be addressed by using a drift reference distribution that models the expected change in a *trait frequency distribution* under neutral drift. This would allow for the calculation of the selection impact by comparing the expectation for a drifting trait with an observation, eliminating the need to individually identify organisms or track their lineages.

If x represents a quantitative trait and $p(x, t)$ is a probability mass function showing the frequency distribution of trait values in the population at time t , we can model neutral drift by studying how $p(x, t)$ changes when mutations are applied to organisms with uniform probability. If $\mu(x, t|y)$ is a probability mass function that represents the measure of a trait with initial value y that was mutated at time t , then the expected distribution of traits at $t + 1$ is given by

$$p_D(x, t + 1) = \sum_{y=-\infty}^{\infty} p_D(y, t) \mu(x, t|y). \quad (2.16)$$

If the shape of μ does not depend on y , then p_D is the discrete convolution of p and μ

$$p_D(x, t + 1) = \sum_{y=-\infty}^{\infty} p_D(y, t) \mu(x - y, t). \quad (2.17)$$

Comparing an actual observation of $p(x, t + 1)$ with $p_D(x, t + 1)$ using the methods described in Section 2.3 will produce a measure of the selection impact based only on the distribution of traits observed in the population.

Another possible trait-based formulation of selection impact could leverage the Fokker-Planck equation [Risken, 1989],

$$\dot{p} = -\frac{\partial}{\partial x} \left(\frac{\partial U}{\partial x} \frac{\partial p}{\partial x} \right) + \frac{D}{2} \frac{\partial^2 p}{\partial x^2}, \quad (2.18)$$

where $F = -\frac{\partial U}{\partial x}$ is the selective force acting on the trait, and D is the diffusion (mutation) coefficient. The FP equation models time evolution of a probability density function $p(x, t)$. The function p can represent the fraction of individuals with phenotype x in the population

at time t . As discussed in the previous section, selection is analogous to a force generating a potential for movement. Here, the force $\frac{\partial U}{\partial x}$ comes from the gradient of the fitness landscape $U(x, t)$. If U is the neutral landscape, there is no force. If there is no variance $p(x) = \delta(x)$, then the force has nothing to act on. Variance is supplied at a constant rate D .

The FP equation can simulate neutral drift by initializing $p(x, t_{\text{initial}})$ with an observation of a population. Substituting U for the neutral landscape and evolving the equation in time gives us $p(x, t_{\text{final}})$, a prediction of the phenotype distribution after drifting neutrally. This can be compared with a second observation of the population at t_{final} using the earth-mover distance. The result is, in spirit, the selection impact since it is the difference between what actually occurred and the expected outcome of neutral drift.

2.7.4 Concluding remarks

The selection impact metric offers a new way of quantifying selection over both time and space. Since the selection impact represents the actual effect of the selection potential acting on the population, it allows us to see evolving populations as they really are. In addition, the selection impact shows promise for a number of additional applications.

In the field of evolutionary optimization, tuning an optimization algorithm to maximize the adaptation of the population is one of the hardest challenges to overcome. Using the selection impact as a diagnostic aid for tuning optimization algorithms gives direct feedback on the actual impact of the changes made to the algorithm. Sections 2.6.3, 2.6.4, and 2.7.2 provide a glimpse of how this is possible. Algorithm tuning can even be done automatically in response to the performance of the population. Even very simple methods of automatic parameter tuning have been shown to greatly improve optimization algorithms [Kramer, 2010].

The selection impact might also be used to detect the current adaptive mode of some populations. For example, whether the population is experiencing directional selection, stabilizing selection, or neutral drift can be seen at a glance when the selection impact of these three modes result in sufficiently different behavior (Section 2.6.2). This primarily

applies at large population sizes, since small populations experience these modes as more similar than different (see Figure 2.3).

Given the central role of selection in evolutionary theory, and the complex web of phenomena that influence the strength and character of selection, a measure of selection is needed that can work in any context. The selection impact metric is built on a theoretical foundation that makes no assumptions about the mode of selection, or the many mechanisms that change the strength of selection. As such, the selection impact can work in nearly any context, and offers all evolutionary scientists, biological or computational, a universal language to describe the strength of selection acting on a population.

2.8 Methods

The empirical data presented in Section 2.6 was collected using a simple evolutionary algorithm. First, a population of digital organisms is initialized. Population sizes vary and are specified in Section 2.6. Each organism has a genome given by a single integer, initially set to zero. When an organism reproduces, the genome may mutate, changing the integer value by $+1$ with probability $\mu/2$, by -1 with probability $\mu/2$, or leaving it unchanged with probability $1 - \mu$. Mutation rates vary and are specified in Section 2.6. Each organism is assigned a score according to one of the 6 fitness functions described below. Once a score is assigned to each organism, parents are chosen for reproduction with a roulette-wheel (fitness proportional) selection algorithm. The generational model (*i.e.*, overlapping versus non-overlapping generations) and the number of parents required for reproduction (*i.e.*, sexual or asexual) vary and are specified in Section 2.6.

2.8.1 Fitness functions

The following fitness functions are used in this work, where x is the organism's integer genomic state: $w(x) = 1$ for the neutral landscape, $w(x) = x$ or $w(x) = x^3$ for a non-deceptive hill-climb landscape, $w(x) = -|x|$ or $w(x) = -x^2$ for a landscape with an inescapable global optimum, and $w(x) = \text{Sawtooth}(x)$ for a deceptive hill-climb landscape. Figure 2.5 plots

these fitness functions for a subset of genomic states. The saw-tooth fitness function can be represented by⁴

$$\text{Sawtooth}(x) = \left\lfloor \frac{x}{w+1} \right\rfloor \left(r + d + \frac{d}{w} \right) - x \frac{d}{w} . \quad (2.19)$$

The w parameter is the width of the valley: the number of deleterious mutations to the right of a local optimum before encountering a beneficial mutation. d is the depth of the valley: the amount subtracted from an organism's score when it is at the right-most position in a fitness valley. All other points in the valley are interpolated between the fitness peak to the left and the bottom of the valley. r is the rise in score from one local optimum to the next. The first term of Equation (2.19) represents the cumulative advantage for crossing each valley, while the second term represents the cumulative disadvantage of crossing each fitness valley. Note that advantages increase as a step function, while disadvantages increase continuously. The sum of these two terms result in the saw-tooth shape of the fitness function.




Figure 2.5: Six fitness functions: five polynomial fitness functions and one saw-tooth function. The saw-tooth is shown as it was configured for Section 2.6.2: $d = 1$, $w = 2$, and $r = 1$.

⁴The floor function, $\lfloor \cdot \rfloor$, is a function that returns the nearest whole number, less-than or equal-to the input.

Chapter 3

Noisy fitness and selection strength

This chapter is based on the article “Connections between noisy fitness and selection strength” [Ragusa and Bohm, 2021].

3.1 Introduction

Selection is perhaps the most central concept in evolutionary theory. Selection can be described as a filtering process which changes a population over time with regard to the result of some evaluation (i.e., a fitness function). The filter can range from fully permissive, where the result of the evaluation is not considered, to maximally restrictive, where only the organism(s) with the best evaluation pass through. These two extreme cases are referred to as drift and elitism, respectively. All other ‘strengths of selection’ fall somewhere in between drift and elitism.

3.1.1 Background

We are interested in understanding the relationship between parameters of evolution and rates of adaptation. Here, we limit our investigation to the relationship between selection strength and rates of adaptation in the context of different fitness landscapes.

In evolutionary theory, the concept of a “fitness landscape” is the relationship between genotypes and fitness [Wright, 1932]. A position in a fitness landscape surrounded by regions

of lower fitness is called a “local optimum”. If a population is positioned on a local optimum, then it must cross a region of lower fitness — called a “fitness valley” — in order to escape the local optima and potentially ascend to higher fitness.

Although a population does not always require a loss in fitness to cross a fitness valley, [Iwasa et al., 2004a] deleterious mutations are often necessary to explain valley-crossing events. For example, [Covert et al., 2013] disallow deleterious mutations and observe that the adaptability of the system is hindered due to the inability to valley-cross. Similarly, it has been shown that elitism hinders adaptation on rugged fitness landscapes [Oliveto et al., 2018].

In this work, we consider three specific evolutionary parameters that each effect selection strength: **population size**, **noisy evaluation**, and **tournament size** (a parameter specific to our model).

3.1.1.1 Population Size

Population size relates to selection strength in that selection in larger populations results in less sampling error because of the greater number of sampling events when choosing parents. Small populations experience drift more readily and so experience less selection strength, while large populations are very sensitive to small fluctuations in fitness and experience strong selection.

Adjusting the population size in an evolutionary model, therefore, is a rather simple way to adjust the selection strength. While this convention is ubiquitous in evolutionary modeling, we believe it is somewhat flawed. Our concern with this approach is that altering population size changes more than selection strength. Increasing the population size increases mutation inflow and the standing variation of genotypes in the population. The increase in standing variation is a natural consequence of the diffusion of genotypes due to mutational inflow [Feller, 1951, Kimura, 1964].

3.1.1.2 Noisy Evaluation

Some recent work has shown that noisy fitness, for example due to environmental noise, can reduce the strength of selection. For example, [Wang and Zhang, 2011] and [Melbinger and Vergassola, 2015] both see an increase in drift-like behavior when fitness is noisy and the selection algorithm has difficulty comparing organisms accurately. Both publications describe the change in behavior as being similar to a reduction in the effective population size.

A third study by [Van Egeren et al., 2018] demonstrates how noisy fitness can also benefit adaptation. By creating new opportunities for deleterious mutations to hitchhike, the stochastic increases in fitness actually help populations cross fitness valleys that would have otherwise been impossible to cross. They also see an increased capability for valley crossing when the noise is increased. It is worth noting that these three studies do not conflict; The drift-like behavior, which disrupts the effectiveness of selection, is precisely how the deleterious hitchhikers survive to cross valleys.

3.1.1.3 Tournament Size

Tournament selection is a selection algorithm commonly used in computational models of evolution. Regarding our work, it has several benefits, but the most relevant is that it has an adjustable strength of selection. The strength of selection is adjusted by the ‘tournament size’ parameter, which controls how many organisms in the population compete with each other for each opportunity to reproduce. The size parameter can be set between drift and elitism, which makes it an ideal dial of selection strength.

3.1.2 Summary of our work

In this work we perform a detailed assay on the relationship between population size, noisy phenotype evaluation, and tournament size, and their effects on rates of genomic change. We run our model on nearly 4,500 different scenarios.

We observe evolution on a smooth fitness landscape as well as nine deceptive landscapes using our model. We show that for the smooth landscape it is always best to have strong

selection with noise-free fitness and a large population. For deceptive landscapes, there is an optimum configuration of tournament size and noise that balances exploration and exploitation. Population size, on the other hand, always increases genomic change when larger, because it not only increases selection strength but also maximizes mutational inflow and standing variation. We see that while these parameters for selection strength have similar effects, they each behave in unique ways. Finally, we suggest that evaluation noise is a better proxy for selection strength than the other two methods.

3.2 Methods

For this work, we implemented our evolutionary simulation using the MABE software [Bohm et al., 2017]. In order to understand the relationship between noise, tournament size and, population size, with rates of adaptation, we ran an array of nearly 4,500 scenarios. Each scenario combines a tournament size with a score noise and a population size. Each combination of parameters is run on 10 fitness functions, one non-deceptive landscape and nine deceptive landscapes. All nine deceptive fitness functions contain fitness valleys which must be crossed in order to reach a higher score; The functions differ in the depth and width of their fitness valleys. We used tournament selection for all scenarios.

3.2.1 Organism Definition

Our evolution model is an agent-based simulation where each generation of digital organisms are evaluated and selected for reproduction. Unlike more complicated digital organisms [Ofria and Wilke, 2004], we have simplified the concept considerably in order to remove as many confounding interactions from our model as possible.

In this work, the organism has a genome that is represented by a vector of 500 integer values, $g = [g_1, g_2, \dots, g_{500}]$. This genome is used to give each organism a score, which is used during selection and is inherited during reproduction. Each site in the genome may mutate during reproduction, with probability $\mu = 0.0005$ per site. When a mutation does occur at some site g_i in the genome vector, the value is mutated following the rule $g'_i = g_i \pm 1$ where

the offset $+1$ or -1 is decided randomly with equal probability.

3.2.2 Fitness Functions

An organism’s score is determined based on the organism’s genome vector g . First a sum of all sites in the genome, x , is computed:

$$x = \sum_{i=1}^{500} g_i \tag{3.1}$$

Note that since all g_i are integers, the sum x is also an integer.

The sum x is then passed through a function which maps specific values of x to a score.

We use ten fitness functions in this work. The first is a simple hill-climbing function $w(x) = x$. This function allows us to observe the dynamics of the system with no deception (i.e., no fitness valleys).

The other nine functions all include deceptive fitness valleys. There is a default fitness function, and eight alternate functions which modify the topology of the default function in order to explore how valley depth and width change the outcomes of the simulation.

All the deceptive fitness functions are periodic, meaning after crossing any fitness valley there will be another exactly like it, indefinitely so. However, the score at each of the local optima is 5 higher than the score of the previous (increasing with x) which incentivizes the population to continuously climb to ever higher peaks, similar to the simple hill-climbing function.

Due to the repeating nature of these functions, they can be defined by just one valley. Each function begins with a plateau one mutation across, creating two positions of equal fitness. Following that, as x increases, the score drops until finally jumping up to a height greater than the starting position. The difference in score between the starting position and the lowest point in the valley we call “valley depth”. The length, in x , of the valley, including the plateau, we call “valley width”. Valley depth and width are varied between the nine deceptive functions. The default function has a width and depth of 6. Shallow versions of

the function have depth 3 while deep versions have depth 9. Similarly, narrow versions of the function have width 3 while wide versions have width 12. Figure 3.1 provides a visual representation of each of the nine deceptive fitness functions.




Figure 3.1: A visual representation of the 9 deceptive fitness landscapes we use throughout this work. The functions are arranged in a 3×3 grid such that the depth of the fitness valley increases from top to bottom and the length (in mutational distance) of each valley increases from left to right. The names and dimensions of each fitness function are displayed above each plot.

3.2.3 Real-Valued Tournament Selection

In this work, we use tournament selection. The tournament selection algorithm is typically defined as shown in Algorithm 1 where T , the integer variable that determines the tournament size ($T \in \mathbf{Z}$), is restricted such that $T \geq 1$. Here, we wish to consider very weak selection strength. In order to achieve this, we extended the tournament algorithm to include tournament sizes that are between 2 and 1. Therefore, we define Algorithm 2 as a means of defining and interpreting real-valued tournament sizes.

In Algorithm 2, T is still restricted such that $T \geq 1$, but is now a real value ($T \in \mathbf{R}$). Algorithm 2 amends Algorithm 1 such that a tournament of size $\lfloor T \rfloor$ is always conducted, but an extra challenger is drawn with probability $T - \lfloor T \rfloor$ resulting in an average tournament size of T . Note that every tournament conducted by Algorithm 2 is composed entirely of integer-sized tournaments.

The idea of fractional tournament size is, as far as we are aware, a new concept. It may be challenging to conceive of the meaning of a tournament size of 1.1. The way to think of $T = 1.1$ is to see it as tournament 1 (random selection) 90% of the time and tournament 2 10% of the time. So, given a population of size 100, this would mean that 90% of the next generation would be the result of random selection and only 10% would be the result of tournament 2.

Algorithm 1: Standard Tournament Selection

Result: A digital organism destined to reproduce
count = 1;
best = chooseRandomlyFrom(population);
while *count* < *T* **do**
 challenger = chooseRandomlyFrom(population);
 if *challenger.fitness* > *best.fitness* **then**
 best = challenger;
 end
end
return best;

3.2.4 Applying Noise to Score

In some experiments, we investigate the effects of noisy evaluation on rates of adaptation. Whenever we add noise to the score of an organism as part of our experiments, it is applied as the last step before selection. N is the variable that controls the level of noise. Once an organism's score has been determined by the relevant fitness function, a random value in the range $[-N, N]$ is chosen with uniform probability and then added to the organism's score. This creates a level of uncertainty; Organisms whose scores are closer together than $2N$ may sometimes have their competitive advantages reversed by the noise applied to their score.

Algorithm 2: Real-Valued Tournament Selection

Result: A digital organism destined to reproduce
count = 1;
best = chooseRandomlyFrom(population);
whole = floor(T);
remainder = T - whole;
if *randomFloat(0,1) ≤ remainder* **then**
 | whole = whole + 1;
end
while *count < whole* **do**
 | challenger = chooseRandomlyFrom(population);
 | **if** *challenger.fitness > best.fitness* **then**
 | best = challenger;
 | **end**
end
return best;

3.2.5 Model Justification

The genetic model, selection algorithm, and fitness functions we use in this work were chosen after some consideration. We are primarily interested in the dynamics surrounding valley-crossing events and how those dynamics are affected by the parameters that we vary.

Our genomic model removes the influence of many genetic phenomena such as sexual recombination or crossover, allele dominance, gene epistasis, horizontal gene transfer, and gene translation.

Our selection algorithm only considers the relative rank of scores in the population and therefore maintains a constant strength of selection (unlike fitness-proportional selection algorithms, which experience weakened selection as the average score increases).

Our fitness functions are designed to have an unlimited number of identical valley-crossing opportunities, regardless of the population's position on the function. This allows us to characterize how well a particular configuration of the model crosses valleys. By looking at the final score of each run, we can determine how many valleys were crossed during the runtime and compare this with other configurations to easily compare rates of adaptation under different conditions.

3.2.6 System Configuration

Every experiment condition was replicated 100 times. Each replicate was run for 20,000 generations. Every generation, the parent organisms are fully replaced with the offspring they generated. Population size, tournament size, and score noise are varied between experiments and are noted in the results section.

Readers wishing to replicate the results from this paper are directed to the supplemental materials, which include files and instructions for generating the data presented in this paper (see: <http://github.com/cliff-bohm/ALIFE-2021-Connections>).

3.3 Results

3.3.1 Smooth Landscapes

Figure 3.2 shows the final scores of runs on the smooth landscape. The smooth landscape is a simple hill climb function in which each increase in x awards the same increase in score. The hill-climbing nature of the fitness function removes the explore-vs-exploit trade-off, leaving pure exploitation as the optimal strategy. This is reflected in the data; Strong selection, large populations, and no score noise result in the highest final scores.

3.3.2 Deceptive Landscapes

Next we consider deceptive fitness landscapes. We use nine fitness functions that contain fitness valleys of different widths and depths. The functions are each labeled as deep/shallow or narrow/wide versions of a default valley shape. Unlike the simple hill-climbing function, these deceptive landscapes require a balance between exploration and exploitation.

3.3.2.1 Noise vs Population size

Figure 3.3 shows the final outcomes of runs with varying population sizes and score noise amounts, with tournament size fixed at 2. If altering the population size was fundamentally the same as changing the score noise, we would expect to see a one-to-one mapping that

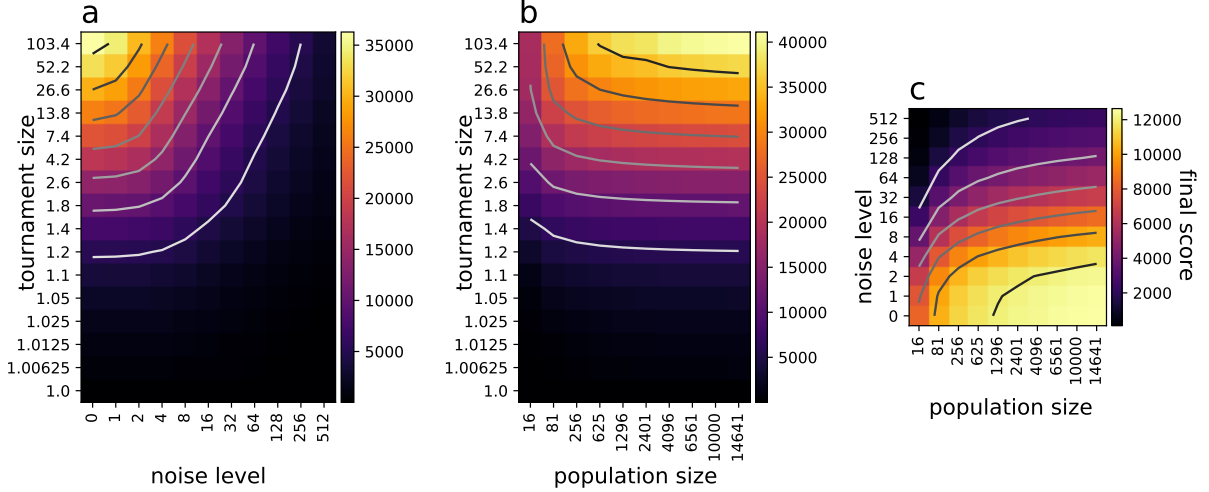


Figure 3.2: Heatmap showing the final scores of all runs on the smooth fitness function. Contour lines indicate lines of equal score. Every square represents the average score of 100 replicates. a) pairs tournament sizes with noise levels while fixing population size at 625. b) pairs tournament size with population sizes while fixing noise at 0. c) pairs noise level with population size while fixing tournament size at 2.

relates pairs of parameters to the same rate of adaptation. Instead, we see that there is no one-to-one mapping between the two parameters.

Further, we see that the narrow fitness valleys are similar to the simple hill-climbing landscape. This is most likely due to the narrow valley being so short that tunneling across via mutation is possible. Under these circumstances, low noise and large populations are advantageous. The large populations ensure that rare tunneling events happen more frequently, while the noise-free score help ensure that selection detects those individuals.

In the default and wide valley conditions, tunneling is less likely, so the effects of purifying selection become an obstacle. Under these circumstances we see that score noise creates uncertainty for selection and deleterious mutations can persist long enough to assist the population in crossing the valley. When noise is too low, the population is stuck, and we see large regions of low adaptation in Figure 3.3. We also see that when noise is too high, adaptation begins to suffer due to the decreased ability to detect differences relative to the noise-free scores; As the noise begins to dominate, the system devolves into drift-like behavior. As in the narrow conditions, large populations result in fast adaptation, presumably due to




Figure 3.3: This figure shows the final outcomes of simulations that combine Population Size and Score noise parameters with tournament size $T = 2$. Every cell is the average of 100 replicates. The color intensity indicates the number of valleys crossed during the total runtime of 20,000 generations (final score divided by 5). The nine subplots each show the results of a different fitness function.

increased standing variation (genetic diffusion) and increased mutational inflow. Apparently, these effects outweigh increases in selection strength because we see no point where population size becomes too big.

3.3.2.2 Tournament size vs Population size

Figure 3.4 shows the final outcomes of runs with varying population sizes and tournament sizes, with noise fixed at 0. As tournament size relates directly to selection strength, if altering the population size was fundamentally the same as changing the selection strength,

we would expect to see a one-to-one mapping that relates pairs of parameters to the same rate of adaptation. Instead, we see that there is no one-to-one mapping between the two parameters.

Further, we see that the narrow fitness valleys are similar to the simple hill-climbing landscape. This is most likely due to the narrow valley being so short that tunneling across via mutation is possible. Under these circumstances, high selection strength and large populations are advantageous. The large populations ensure that rare tunneling events happen more frequently, while the strong selection ensures those rare events go to fixation.

Once again, in the default and wide valley conditions tunneling is less likely, so the effects of purifying selection become an obstacle. Under these circumstances, we see that only when tournament size is low can deleterious mutations persist long enough to assist the population in crossing the valley. When tournament size is too high, the population is purifying deleterious mutations too quickly to cross valleys, and we see large regions of low adaptation in Figure 3.4. We also see that when tournament size is too low, adaptation begins to suffer due to the decreased ability to discern differences in score. As tournament size approaches 1, the system tends towards drift-like behavior. Large populations continue to be advantageous here for the same reasons indicated in the Tournament size vs Population size results.

3.3.2.3 Tournament size vs Noise

Figure 3.5 shows the final outcomes of runs with varying tournament sizes and score noise amounts, with the population size fixed at 625. If altering the tournament size was fundamentally the same as changing the score noise amount, we would expect to see a one-to-one mapping that relates pairs of parameters to the same rate of adaptation. Unlike Figures 3.3 and 3.4, we see a very strong relationship between the two, though it is not one-to-one in all cases. There are, however, large areas within each plot where the relationship is approximately linear.




Figure 3.4: This figure shows the final outcomes of simulations that combine Tournament Size and population size parameters with noise $N = 0$. Every cell is the average of 100 replicates. The color intensity indicates the number of valleys crossed during the total runtime of 20,000 generations (final score divided by 5). The nine subplots each show the results of a different fitness function.

Here again, we see that the narrow fitness valleys are similar to the simple hill-climbing landscape. This, again, is most likely due to the narrow valley being so short that tunneling

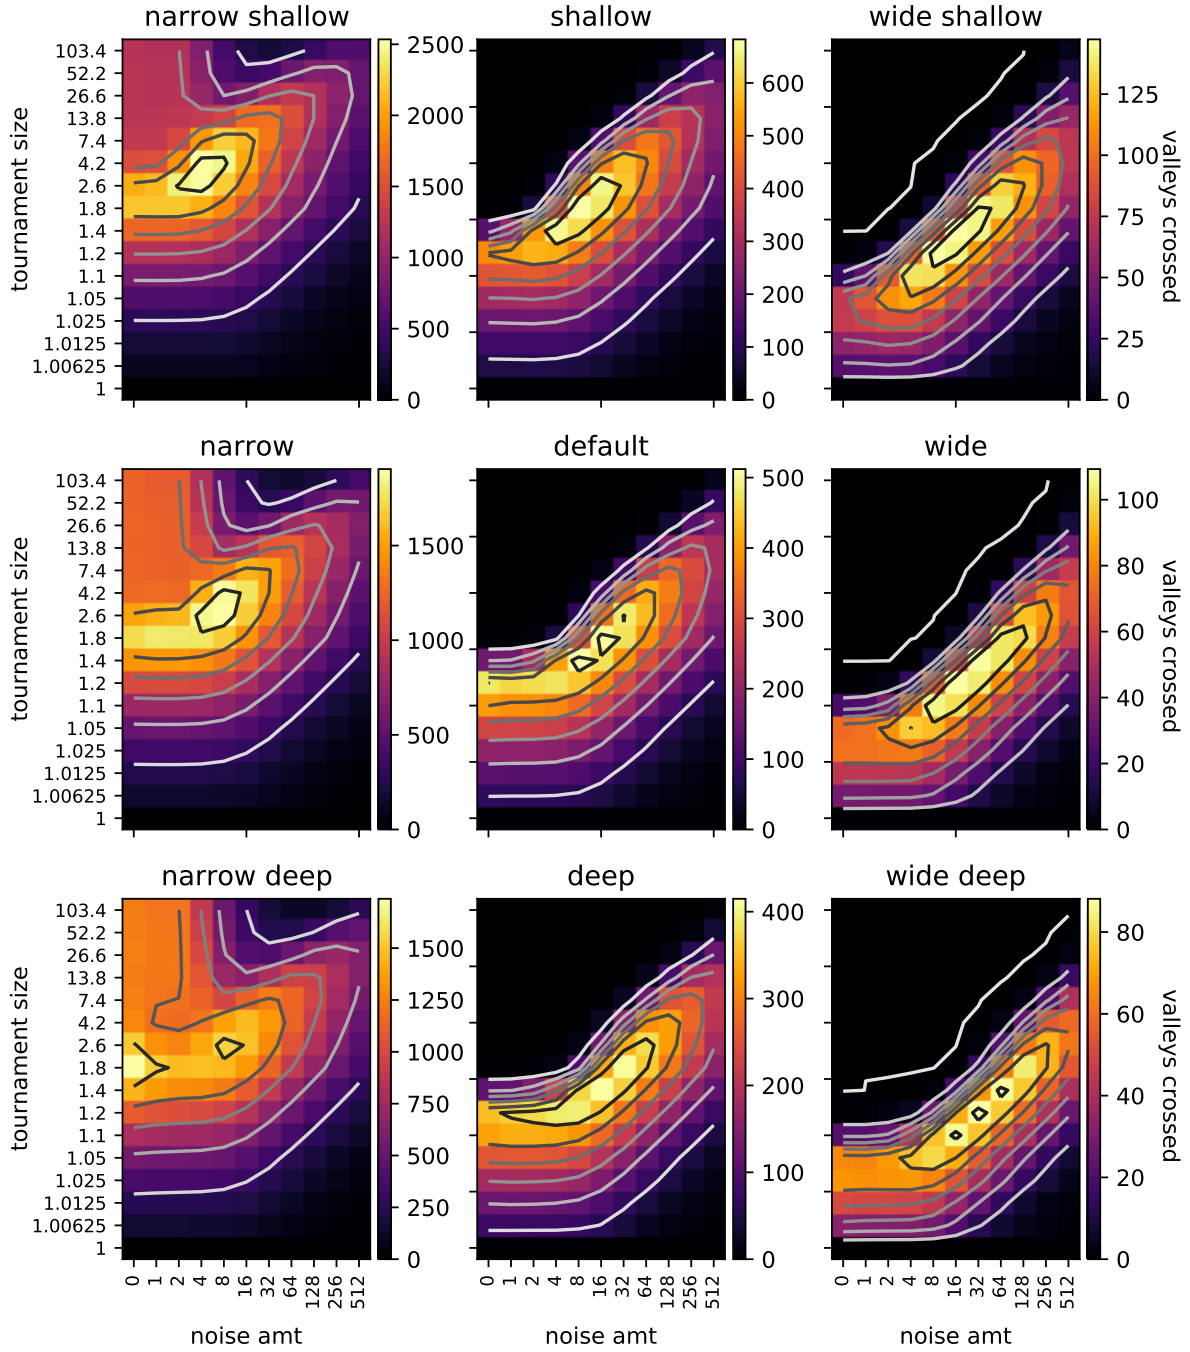


Figure 3.5: This figure shows the final outcomes of simulations that combine Tournament Size and Score noise parameters with Population size $P = 625$. Every cell is the average of 100 replicates. The color intensity indicates the number of valleys crossed during the total runtime of 20,000 generations (final score divided by 5). The nine subplots each show the results of a different fitness function.

across via mutation is possible. Under these circumstances, high selection strength and low noise are advantageous. The strong selection ensures rare high-fitness mutants go to fixation,

while the noise-free score help ensure that selection detects those individuals.

As seen before, in the default and wide valley conditions tunneling is less likely, so the effects of purifying selection become an obstacle. Under these circumstances, we see that all but the very low and very high tournament sizes are viable. The very low settings result in drift-like behavior, while the very high settings result in purifying selection that restricts valley crossing. Interestingly, the score noise seems to be capable of compensating for the increase in selection strength. As the noise increases, the optimal tournament size increases too, suggesting that although the score noise adds uncertainty, this can be overcome simply with more selection strength. However, this compensatory relationship begins to fail at larger noise strengths because once the score signal is dominated by the noise, no amount of selection strength can compensate. As the score noise increases, the system will eventually tend towards drift-like behavior.

3.4 Discussion

In this work we show the relationships between all combinations of population size, noisy evaluation, and tournament size in the context of rate of adaptation.

One key observation is that there are relationships between each of the three parameters and selection strength. However, while the behavior of these parameters are each related to selection strength, the relationship is imperfect and there is no perfect one-to-one relationships between any of the parameters we have considered. This means that each of these parameters modify selection strength in different ways.

We began our discussion of deceptive landscapes by introducing the concept of fitness valley width and depth. In our simulations, we examined valley depths between 3 and 9 and valley widths between 3 and 12. However, you could imagine setting the width or depth to zero. Doing so removes all deception from the fitness landscape. In other words, setting the width or depth to zero defines the smooth fitness landscape. With this in mind, we can see that all the results sit on a continuum of deception ranging from none to high.

In the data collected from the smooth function, we see that for each pair of parameters, the maximum rate of genomic change occurs where the parameters are tuned to maximum selection strength. The rates of change decrease when any parameter moves away from strong selection strength. The contour lines in Figure 3.2 show a trade-off between the values on each axis that indicates where strength of selection is roughly equal. As the functions become more deceptive, we see that the position of the greatest change becomes uncoupled from the highest selection strength settings.

We see an interesting alignment between the areas of the greatest change in Figure 3.5 and the contour lines in the smooth function plot (Figure 3.2.a). As deception increases, the area of the greatest change shifts to align with contour lines that are farther from the point of the highest selection strength. This illustrates the trade-off between exploration and exploitation. As the functions become more deceptive, weaker selection is needed to allow more exploration across wider and deeper valleys. In addition, the shape of the regions in Figure 3.5 indicates an almost-compensatory trade-off between tournament size and noise amount, which also aligns with the contour lines in the Figure 3.2.a.

The trade-off we identified in Figure 3.5 is not seen in Figures 3.3 and 3.4. If population size only effected selection strength, we would expect to see that very large populations would result in too much purifying selection to allow for exploration across the more deceptive valleys, but this is not the case. In these plots, we still see that a large tournament size or low noise configurations result in reductions in the rate of adaptation, but larger populations always result in a faster rate of change.

3.4.1 Population size and selection strength

Increases in population size are known to also increase the selection pressure by improving the sampling of individuals in the population during selection. However, in the data presented we do not see an optimal population size; We observe that rate of adaptation increases with population size without seeing a reversal of this trend at high values. This seems to show that the trade-off between exploration and exploitation is not invoked when increasing the

population size. Therefore, we conclude that there are other forces at work. Namely, we believe that the increase to mutational inflow and standing genetic variation that result from increasing the population size are responsible for the absence of this trade-off. Both of these forces increase the ability of the population to explore the fitness landscape. We believe these increases must be compensating for the increased selection strength that on its own would have caused exploration to be reduced due to purifying selection.

3.4.2 Proxies for selection strength

Based on our results, if someone wishes to design an experiment using an evolutionary model and wishes to vary selection strength in their model as an independent variable, they ought not to use population size as their proxy for selection strength. In the case where it is possible, it should be preferred to use a model of selection which comes with its own explicit parameter for selection strength, like tournament selection.

We find that, similar to tournament size, there is an optimum setting for noise when all other parameters are held constant. Lower than this level, the effects of selection may become too strong for the population to explore the fitness landscape effectively. Higher noise results in an unreliable signal for the selection algorithm to effectively exploit higher fitness individuals. In other words, adding noise to score creates the same trade-off as we see when we adjust the tournament size.

The relationship between tournament size and noise seems to be roughly one-to-one. This similarity suggests that evaluation noise is a rather good proxy for tournament size, which is itself a direct controller of selection strength. Ergo, noise is a rather good proxy for directly adjusting selection strength.

One benefit of using noise as an independent proxy for selection strength is that it is algorithm independent. Regardless of the selection algorithm in use, we can always apply noise to the scores it will operate on to reduce selection strength.

3.4.3 Neutral Theory in the Context of Noisy Score

The way we have applied noise to our digital organisms' fitness is somewhat abstract. The noise could be interpreted as gene expression variance, or a stochastic environment where luck plays a role in survival and reproduction. In nature, we might expect both of these sources of noise to co-occur with, possibly, even more. Therefore, natural organisms may experience a wide range of rates of adaptation even with all other factors usually attributed to controlling selection strength held constant.

This has large implications on the neutral theory of molecular evolution [Kimura, 1987]. In brief, neutral theory posits that the majority of genetic variation is neutral variation. Since noisy fitness can change the effective strength of selection, noise can change what qualifies as neutral or not. What was once neutral may become detectable if the noise on fitness is reduced, while detectable traits may become neutral in periods of additional noise. This has far-reaching consequences once you consider theories of local optima escape which require the discovery of neutral traits, for example extradimensional bypass [Cariani, 2002]. In the presence of noise, a pre-existing trait may become the neutral ridge that is used to bypass a fitness valley.

3.4.4 The effectiveness of small Tournaments

It came as a surprise that a large portion of the fast-evolving runs were using tournament sizes below $T = 2$. Seeing as how $T = 1.4$ means that more than half (60%) of the population is reproducing under drift-like conditions, it is fascinating that the population does not experience mutational meltdown. Perhaps $T = 1.4$ is like reducing the effective population size to 40%, but also including a pool of additional organisms that might be shuffled in or out between generations. Interpreted this way is not at all surprising, that $T = 1.4$ evolves. To the contrary, it would seem to have the benefit of maintaining diversity more effectively than $T \geq 2$.

The excitement over this should be tempered somewhat by acknowledging that the rates

of change seen on the smooth function (Figure 3.2) for tournament sizes at or below 1.4 are extremely low. So, while these low tournament sizes help to maximize the rate of evolution by balancing exploration and exploitation, they cannot allow evolution to proceed faster than would be the case if the landscape was smooth and non-deceptive.

Even so, real-world fitness landscapes are high dimensional and likely highly deceptive. As a consequence, it's interesting to see that very small tournament sizes can have observable long-term outcomes on evolution for deceptive landscapes. Systems experiencing near-drift behavior are still able to undergo directional selection when the time period of observation are long enough.

3.4.5 Future Work

One interesting consequence of the results presented here are the implications for multi-trait evolutionary systems. For example, consider a digital organism who's total fitness is given by summing together progress on both the shallow and narrow fitness functions. If the system is evolving under tournament selection without noise, the optimum for the narrow function is $T = 1.8$, while for the shallow function it is $T = 1.2$. Further, if the shallow function is evolved at, $T = 1.8$ it is virtually stuck. Therefore, the choice of tournament size could cause an organism on the multi-trait landscape to evolve on both functions or on just one. However, say we choose poorly ($T \geq 1.8$ which will halt progress on shallow), adding noise to the system, say $N = 8$, could move both traits into a fast-evolving scenario.

Our work has focused on results using the tournament selection algorithm. As a rank based algorithm, tournament selection is going to behave differently from more traditional fitness-proportional selection algorithms (such as roulette). Fitness proportional selection algorithms are known to have variable selection strength as a result of increasing organism scores. In preliminary work, we attempted to apply a correction to roulette selection that would fix the diminishing selection strength problem, but ultimately we chose to go with tournament selection because no such correction was necessary. Future investigations should use a fitness proportional selection method and compare the results with tournament.

In our analysis of score noise, we find that there is often a single optimal value for this parameter, given the other parameters are already chosen. There is a potentially untapped engineering technique in the use of score noise to enhance the rates of adaptation of evolutionary optimization algorithms. The single optimum makes tuning this parameter rather simple; Simple enough to automate during an evolutionary search.

Chapter 4

The role of disequilibrium in evolutionary discovery

4.1 Introduction

Empirically, the rate at which populations evolve is extremely variable [Fitch, 1995]. Among the many factors that affect the speed of evolution are the size [Ohta, 1972, Fogle et al., 2008, Handel and Rozen, 2009, Jain et al., 2011] and structure [Möller et al., 2019] of the population, the dynamics, and structure of the fitness function [Handel and Rozen, 2009, Weissman et al., 2009, Kryazhimskiy et al., 2009], mutation rates [Kimura, 1964, Kimura and Crow, 1964, Weissman et al., 2009], and the size and distribution of beneficial mutations [Gerrish and Lenski, 1998, Orr, 2005]. Getting a better understanding of these variable rates of adaptation requires us to investigate how populations shift between periods of stasis and periods of adaptation.

In a genetic fitness landscape model [Wright, 1932, Maynard Smith, 1970], genotypes specify locations on a landscape and the height differences between locations represent relative fitness: genotypes at higher elevations have an advantage over those at lower elevations. Movements on the landscape correspond to mutations from one genotype to another [Maynard

Smith, 1970]. Peaks in a fitness landscape (i.e., genotypes surrounded entirely by lower fitness) are referred to as local optima. A population may be in evolutionary equilibrium when located on a local optimum because while areas of higher fitness are present elsewhere in the landscape, the population must lose fitness in order to cross the valley between their current peak and another higher peak. The likelihood that a population is trapped on a peak is related to the extent to which populations are able to explore. Features like deep and wide valleys, weak mutation and strong selection work to limit exploration and inhibit valley crossing.

The literature has identified multiple processes that enable evolving populations to cross fitness valleys. Some theories suggest mechanisms for valley crossing that do not require the loss of fitness, such as one organism accumulating multiple mutations in a single generation, environmental changes that alter the fitness landscape [Gillespie, 1984b, Steinberg and Ostermeier, 2016], or effects that add dimensions to a fitness landscape (extradimensional bypass [Cariani, 2002]). Neutral theory [Kimura, 1983, Kimura, 1987, Kimura, 1991] suggests that local optima do not really exist and that instead there are always neutral paths to higher fitness [Gavrilets, 2003], while nearly neutral theory [Ohta, 1992, Ohta, 2002, Fay et al., 2002, Kern and Hahn, 2018] posits that some deleterious mutations may be tolerated as long as the deleterious effects are small.

Both small population size and increased mutation rates can reduce selection strength [Lynch and Conery, 2003, Lynch, 2007], thus reducing the effectiveness of purifying selection to purge deleterious mutations and enhancing the probability of valley crossing [Jain et al., 2011, Ochs and Desai, 2015, LaBar and Adami, 2016] via stochastic tunneling [Iwasa et al., 2004b, Weinreich and Chao, 2005, Weissman et al., 2009]. In addition, a reduction of selection strength during range expansion [Excoffier and Ray, 2008, Palmer et al., 2012] can also result in an enhanced probability of valley crossing [Burton and Travis, 2008].

When the population size is large and the mutation rate not too low, the equilibrium state of the population consists of many variants: a molecular quasispecies [Eigen, 1971, Eigen

and Schuster, 1979, Eigen et al., 1989, Wilke, 2005, Forster et al., 2006]. If the valley is not too deep and not too wide, the mutant cloud may include genetic variants on the other side of a local optimum, or, at least, support mutant subpopulations that are sufficiently close to a new local peak to increase the chance that this peak will be discovered by stochastic tunneling [Iwasa et al., 2004b]. The name 'stochastic tunneling' can be misleading, but it is merely when a genotype that has several mutations relative to the wild-type goes to fixation without any of the intermediate genotypes (single-mutant, double-mutant, *etc.*) themselves going to fixation.

Despite all of these mechanisms that allow for valley-crossing, which in turn allow a population in equilibrium to transition into disequilibrium, we still do not have a simple mechanism that can sustain a period of disequilibrium long enough to reproduce punctuated equilibrium.

In this chapter, we identify a new mechanism that significantly affects rates of evolution by influencing the rate at which populations can cross fitness valleys. This new mechanism can generate periods marked by rapid adaptation punctuated by stasis, and provides insights into why an evolving population switches from rapid adaptation to stasis.

To understand this new mechanism for valley-crossing, consider what happens when an evolving asexual population discovers a beneficial mutation. If the mutation provides a sufficient benefit, the genotype carrying the mutation is likely to take over the population in a selective sweep. In the case of a spatial population, as depicted in Fig. 4.1, the mutants that compete with the wild type are found on the leading edge between the mutants and the wild type. Leading-edge mutants experience a weaker strength of selection because they compete with, and have a fitness advantage over, the wild type. The boundary itself provides an additional benefit by facilitating weak competition. Reduced selection pressure on advantaged organisms occurs in well-mixed populations too, and comes as a result of maintaining above-average fitness when carrying additional deleterious mutations.

We call the phenomenon of reduced selection experienced by high-fitness mutants because

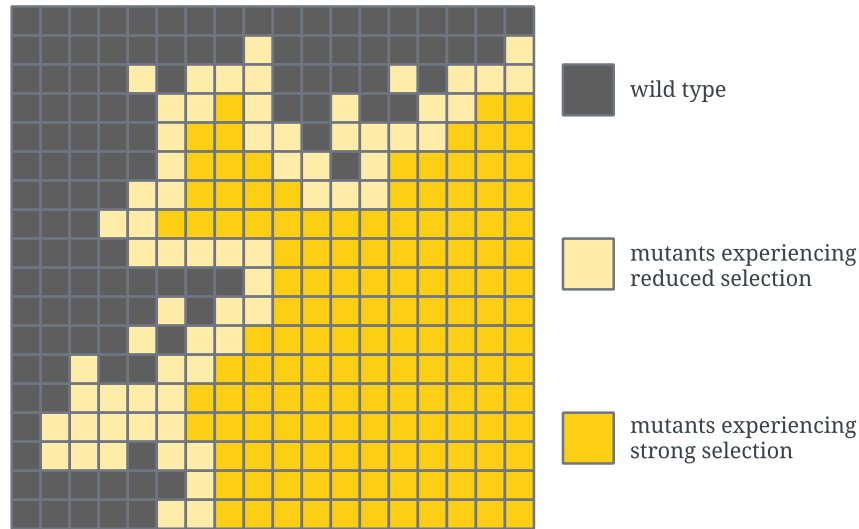


Figure 4.1: A spatial population where a wild-type population (gray) is in the process of being outcompeted by a mutant type (yellow and orange). The members of the mutant type (yellow) on the leading edge experience a reduction in selection, compared to other mutant organisms (orange), due to interactions with the wild type.

of their competitive advantage during selective sweeps, “the Free-for-All effect” (FFA). This name was chosen because the organisms in the lineages that experience reduced selection are able to accumulate deleterious mutations while still remaining competitive; for them, it is an evolutionary free-for-all. This does not imply that lineages experiencing FFA are entirely immune to the effects of selection: lethal mutations are still evolutionary dead-ends, and if the deleterious mutations are too disabling, other lineages that are less encumbered may outcompete them.

If an organism succeeds in reaching a new peak during a period of FFA (i.e., during a selective sweep), the new beneficial type will experience a fitness advantage over the remainder of the population resulting in a new selective sweep. This extends the period where FFA is active; each new valley crossing acts to keep the window of enhanced exploration afforded by FFA open, potentially leading to more discoveries. This can result in a cascading chain of discoveries that will continue until the population finally fixes on a peak, closing the FFA window.




Figure 4.2: (a): A typical fitness valley of width $v = 5$ (it takes 5 mutations from peak to peak), in a one-dimensional multiplicative fitness function (here, \log of fitness is shown). Deleterious mutations from the peak in the direction of the next higher peak have an effect $1 - s_d$. The adjacent peak has fitness $1 + s_b$. (b): Periodic fitness function as a function of the number of mutations k . In this version, valleys are $v = 5$ mutations wide. Log-fitness is shown. (c): A neutral version of (b), where the valley is replaced by a series of neutral mutations that is used when measuring T_{nd} .

4.2 Results

Free-for-all theory suggests that during selective sweeps, a portion of the mutant subpopulation will experience reduced selection pressure, which results in an increased chance of valley crossing. We can test FFA by observing variability in the rate of adaptation (i.e., an increase in the frequency of valley-crossing between the time of discovery and the time of fixation of a beneficial mutation). However, there are a number of alternative explanations that may account for similar observations, such as fitness function topology or epistasis. To ensure that we are detecting the effects of FFA, it is necessary to design experiments that limit confounding effects. To that end, we use a digital evolution system based on a simple fitness function shown in Fig. 4.2b, which are repetitions of the fitness valley described in Fig. 4.2a. Experiments are run with various population sizes, fitness function difficulties, and population structures (see Methods).

Organisms in our system consist of a genome with a single locus that encodes the mutation number k . This allele mutates with a probability μ to $k + 1$ or $k - 1$ with equal probability. The periodic function Fig. 4.2 can be written in terms of the genotype k and the valley width

v as

$$w(k) = (1 + s_b)^n (1 - s_d)^m, \quad (4.1)$$

where the value $m = k \pmod v$ determines how far the genome is within the valley, and $n = \frac{k-m}{v}$ simply numbers the peaks that have been climbed.

This fitness function ensures that the selective pressures are exactly the same on each fitness peak. Hence, every valley crossing occurs under the same conditions (no diminishing returns) so that changes in the rate of adaptation cannot be due to the fitness function.

For each experiment, we will measure the time it takes to cross the fitness valleys under various conditions, and also measure the fixation time of a beneficial mutation. The first valley-crossing time we measure is the stochastic tunneling time T_{st} , the average time it takes a clonal population, concentrated at one peak, to stochastically tunnel across a fitness valley towards a higher peak via random mutations in the presence of purifying selection. For one-dimensional fitness valleys, T_{st} is well understood and can be calculated as the mean first-passage time of an Ornstein-Uhlenbeck process [Alili et al., 2005, Yi, 2010, Artime et al., 2018] (see also [Weissman et al., 2009]). We also measure the neutral drift time T_{nd} , which is the average time to cross a plateau of the same width (i.e., $s_d = 0$, Fig. 4.2c). This time is even better understood, as it is the mean first-passage time of a Wiener process [Risken, 1989]. In addition, we will measure the actual time to cross valleys T_{obs} , by observing a freely evolving system. Finally, in addition to these valley crossing times, we measure the time that a beneficial mutation takes to fix in a population, T_{fix} . Fixation times are well studied and can be calculated easily as a function of mutation rate, population size, and mutational effect size [Kimura and Ohta, 1969].

4.2.1 Mean rate of adaptation T_{obs} changes in response to population size N

Fig 4.3a shows T_{st} , T_{obs} , T_{nd} , and T_{fix} as a function of population size (N) on a log-log scale, for a fitness landscape with valley-width $v = 5$, a deleterious mutation cost $s_d = 0.33$, and

a benefit for crossing the valley $s_b = 16.08$, with a mutation rate $\mu = 0.1$. This data was collected from populations that are structured as a linear array with a periodic boundary condition.

Fig. 4.3b and c show analogous results from a population structured as a regular grid with periodic boundaries and a well-mixed population, respectively. In Fig. 4.3b, $v = 6$, $s_d = 0.139$, $s_b = 19.107$, and $\mu = 0.1$. In Fig. 4.3c, $v = 6$, $s_d = 0.594$, $s_b = 15.0$, and $\mu = 0.1$.

Fig. 4.3a-c all show the same trend. Low population size results in rapid fixation ($T_{\text{fix}} < T_{\text{nd}}$) and $T_{\text{obs}} \approx T_{\text{st}}$ in these cases. At intermediate population sizes, when $T_{\text{nd}} < T_{\text{fix}} < T_{\text{st}}$, we see a drop in T_{obs} relative to T_{st} . Finally, around the point where $T_{\text{st}} < T_{\text{fix}}$, large population sizes result in long fixation times and T_{obs} becomes far less sensitive to changes in population size.

4.2.2 Observed valley-crossing times are bi-modal for intermediate population sizes

Figs. 4.3d-g show the distributions of T_{st} , T_{obs} , and T_{nd} , and the mean of T_{fix} for different values of N , labeled in Fig. 4.3a. Fig. 4.3d shows that when fixation is very fast compared to neutral crossing $T_{\text{fix}} < T_{\text{nd}}$, each beneficial mutation generally fixes before the next can be discovered, as evidenced by the similar shapes of the T_{st} and T_{obs} distributions (they are the same distribution aside from the number of samples collected). Figs. 4.3e-f show that when fixation is somewhere between the drift and tunneling times $T_{\text{nd}} < T_{\text{fix}} < T_{\text{st}}$, T_{obs} becomes a bimodal distribution, where one mode coincides with the T_{st} distribution and the other mode is sandwiched between T_{nd} and T_{fix} . This indicates there are *two* modes of valley-crossing under these conditions: one where circumstances match the those we created for measuring T_{st} , and another that appears to occur under alternative circumstances. Finally, Fig. 4.3g shows how, when valley-crossing happens much faster than fixation $T_{\text{st}} < T_{\text{fix}}$, the distribution of T_{obs} shifts to the right of T_{st} and is no longer bimodal.

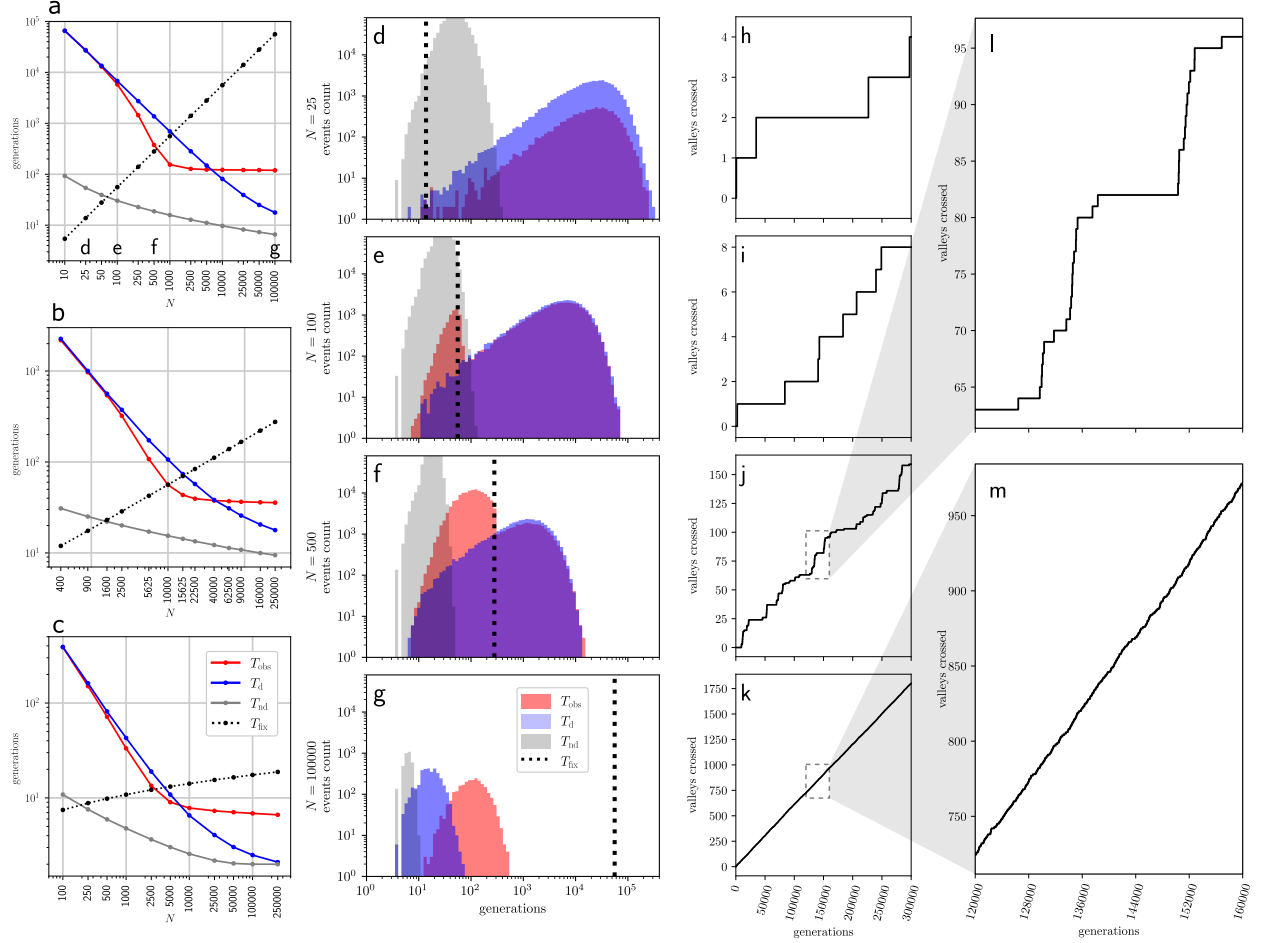


Figure 4.3: (a): Empirical fixation times T_{fix} , stochastic tunneling times T_{st} , neutral drift times T_{nd} , and observed transition time T_{obs} for fitness peaks with deleterious effect $s_d = 0.33$, and beneficial effect $s_b = 16.08$, as a function of population size, in a one-dimensional spatial population evolving at mutation rate $\mu = 0.1$. Four population sizes labeled d, e, f, and g are shown in full detail in panels (d-g). (d-g): Distribution of neutral drift times T_{nd} (gray), observed transitions T_{obs} (red), and stochastic tunneling times T_{st} (blue) for four population sizes labeled in panel (a). Overlapping red and blue distributions appear purple. The mean fixation time is indicated by a dashed vertical line. (d): Distributions at $N = 25$, in the sequential-fixation regime. (e) Distributions at $N = 100$. (f): Distributions at $N = 1,000$. (g) Distributions at $N = 100,000$, in the concurrent-mutations regime. (h-k): the number of valleys crossed over the first 300,000 generations of the experiments used in (d-g). (l-m): zoomed segment of (j-k) respectively.

4.2.3 Intermediate population sizes result in punctuated cascades of valley-crossing

Figs. 4.3h-k shows the cumulative number of valleys crossed by the population during the first 300,000 generations of the experiments that generated T_{obs} in Figs. 4.3d-g respectively. Figs. 4.3h-i each show that only a few valleys were crossed during this period. In Fig. 4.3j, we can see long cascades, periods of rapid evolution where many valleys are crossed, punctuated by periods of stasis; sometimes tens of thousands of generations pass without any change. This is further illustrated in Fig. 4.3l, which shows a detailed view of the fitness trajectory indicated by the dashed line box in Fig. 4.3j. These cascades are a result of the fastest valley-crossings being bunched together, instead of distributed uniformly. This suggests that the population may be in one of two states at a given time: stasis, or rapid change. This is, not coincidentally, reminiscent of the bimodal distribution of Fig. 4.3f; the rapid valley-crossings constitute one mode of the distribution, while the crossings that occur after periods of stasis contribute to the other mode. Finally, Fig. 4.3k (and the detail shown in Fig. 4.3m) depicts a state of constant cascade: there are no periods of stasis observed.

4.2.4 The positions of T_{obs} transitions move in response to mutation rate and valley width

We observed earlier that the points where T_{fix} crosses T_{nd} and T_{st} are important delineations that mark where T_{obs} begins to diverge from T_{st} and where T_{obs} becomes insensitive to further increases of the population size, respectively. In Fig. 4.4, we show nine plots similar to Fig. 4.3a, where T_{st} , T_{obs} , T_{nd} , and T_{fix} are graphed against population size. There are nine such plots showing all combinations of three mutation rates ($\mu = 0.05, 0.1, 0.15$, left-to-right), and three valley widths ($v = 4, 5, 6$, top-to-bottom). The same fitness advantage from the beneficial mutation is used in all nine plots ($s_b = 16.09$), but in order to maintain valley depth, the deleterious effect sizes were set to $s_d = 0.416, 0.333, 0.277$ for valley widths 4, 5, 6 respectively. The center plot contains the same data as Fig. 4.3a. Both increasing mutation

rate and decreasing valley width result in valleys that are easier to cross, and so reduce T_{st} and T_{nd} . In turn, these changes mean that T_{st} and T_{nd} cross T_{fix} at smaller population sizes (x -axis) and at shorter times (y -axis). Thus, increasing mutation rate and decreasing valley width bring the transition points between the sequential-fixation and successive-mutations regimes closer together. Conversely, decreasing mutation rate or increasing valley width moves the transition points further apart.

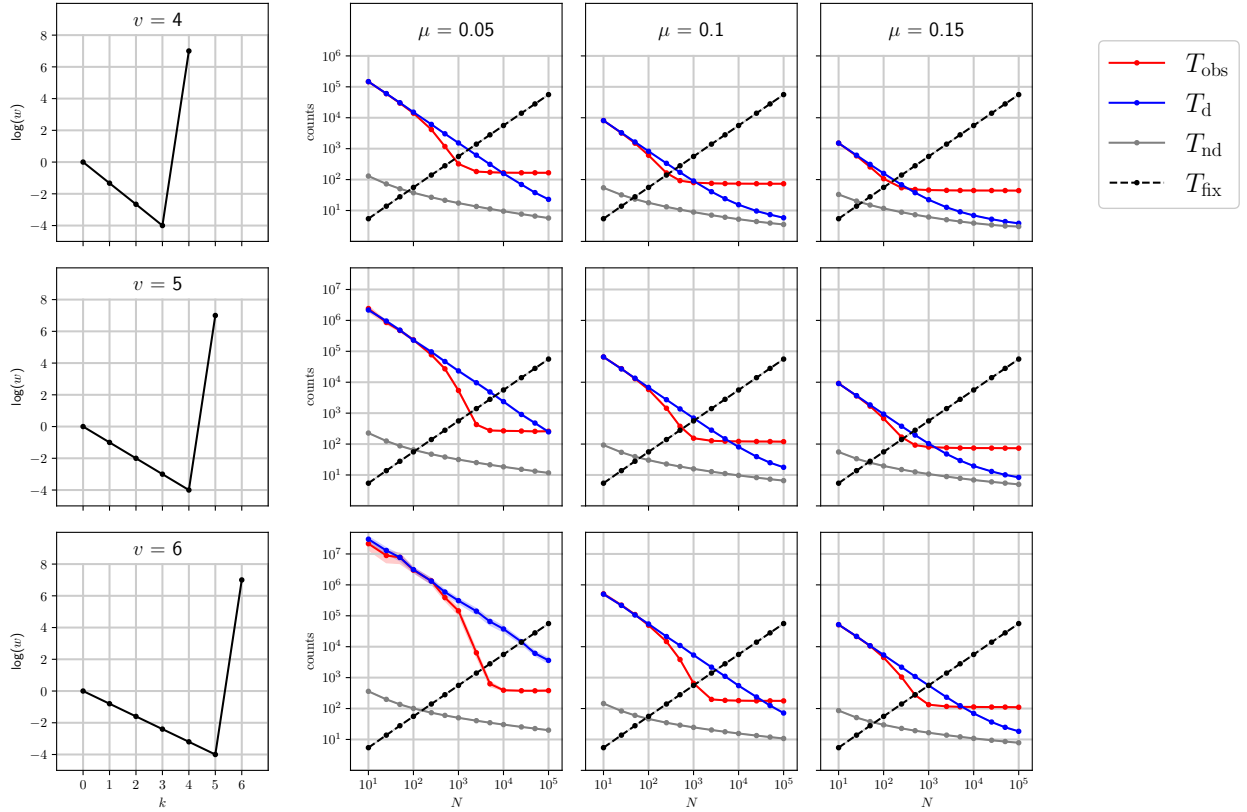


Figure 4.4: Empirical (mean) fixation times T_{fix} , stochastic tunneling times T_{st} , neutral drift times T_{nd} , and observed transition time T_{obs} for fitness peaks with width $v = 4, 5, 6$ (rows) and mutation rates $\mu = 0.05, 0.1, 0.15$ (columns). In order to keep valley depths (the difference between the valley and the peak) constant, we have adjusted the deleterious effect of a mutation so that $s_d \approx 0.42$ for the $v = 4$ landscape, while $s_d \approx 0.28$ for $v = 6$.

4.2.5 There are three modes of evolution

The primary results discussed above can be summarized succinctly with the aid of Fig. 4.5. The figure depicts an idealized picture of how population size affects T_{st} , T_{nd} , T_{fix} , and T_{obs} . The dependence of T_{st} , T_{nd} , and T_{fix} on population size is well understood: as population




Figure 4.5: Three regimes of evolutionary dynamics. The sequential-fixation region (I, shaded in red) is defined by $T_{fix} \leq T_{nd}$, while for the successive-mutations region (III, shaded in blue) $T_{fix} \geq T_{st}$. In between those regions lies region II, where selective sweeps can weaken selection ($T_{obs} < T_{st}$) so that new peaks can be discovered, possibly in cascades.

size increases, an increase in mutation in-flow combined with an expanded quasi-species will decrease T_{st} , and T_{nd} , while increasing T_{fix} . How population size affects T_{obs} is complex, however, because how T_{obs} depends on population size is controlled by the relative magnitude of the other times. Fig. 4.5 also contains a fifth valley-crossing time T_{ffa} . This is the hypothetical time to cross a valley via stochastic tunneling when selection strength is reduced due to FFA, and corresponds to the previously unexplained second mode in the distribution of T_{obs} (Figs. 4.3e-f). Therefore, we call this hypothetical time during free-for-all T_{ffa} and note that the condition $T_{nd} \leq T_{ffa} \leq T_{st}$ holds necessarily by the definition of these terms.

We previously noted that where T_{fix} crosses T_{nd} and T_{st} mark major transition points in

the behavior of T_{obs} . We can use these crossing points to divide the range of population sizes into three regions that exhibit unique evolutionary dynamics from one another.

We refer to the condition where $T_{\text{fix}} < T_{\text{nd}}$ as region I (shaded red in Figure 4.5). Region I is already well-characterized in the literature, and corresponds to the sequential-fixation (“weak mutation–strong selection” or WMSS) regime where beneficial mutations fix one at a time, in the order they are discovered [Desai and Fisher, 2007, Desai et al., 2007]. In region I, because $T_{\text{fix}} \ll T_{\text{nd}}$, FFA does not play an important role in the rate of adaptation since a new benefit would need to appear faster than the average neutral drift time, suggesting exceptionally weak selection due to FFA. Recall that $T_{\text{ffa}} \geq T_{\text{nd}}$, so although the reduction in selection strength due to FFA always occurs, the net effect in this case is moot. Therefore, as we would expect, and in accordance with the literature, we see that $T_{\text{obs}} \approx T_{\text{st}}$ under this condition.

We refer to the condition where $T_{\text{nd}} \leq T_{\text{fix}} \leq T_{\text{st}}$ as region II (depicted white in Fig. 4.5). Region II does not correspond to any previously defined mode of evolution. In region II, the reduction in selection pressure resulting from FFA *can* aid in valley crossing (unlike region I). As T_{fix} increases towards T_{ffa} in response to the population size, we see that the probability of cascades increases, which reduces T_{obs} (the average of T_{st} and T_{ffa} , weighted by occurrence) relative to T_{st} . This is reflected in Figs. 4.3a-c, e-f, j, and l, as well as in Fig. 4.4.

As the population size continues to increase, and T_{fix} approaches T_{st} , we see that the system approaches a state of nearly-constant selective sweeps because the discovery of beneficial mutations occurs more quickly than the fixation.

We refer to the condition where $T_{\text{st}} < T_{\text{fix}}$ as region III (shaded blue in Fig. 4.5). Region III is also well-characterized in the literature and is the concurrent-mutations regime (also called the strong mutation–weak selection (SMWS) regime) [Rouzine et al., 2003, Wilke, 2004, Desai and Fisher, 2007, Desai et al., 2007, Rouzine et al., 2008, Brunet et al., 2008, Fogle et al., 2008, Park et al., 2010, Fisher, 2013, Good et al., 2012]. Here, beneficial mutations are discovered at a faster rate than they fix, resulting in competing subpopulations containing different beneficial

mutations. In this regime, evolution slows down due to clonal interference, and the speed of adaptation becomes largely independent of population size [Gerrish and Lenski, 1998, Kim and Stephan, 2003, Wilke, 2004, de Visser and Rozen, 2005, Campos and de Oliveira, 2004]. In region III, because $T_{\text{ffa}} \ll T_{\text{fix}}$, the system is in a constant state of selective sweeps and therefore constantly experiencing the reduced selection from FFA. However, the substantial decrease in the rate of adaptation due to clonal interference eventually overshadows the speed-up due to FFA. As we approach “infinite population size” all valley-crossing times, theoretical and actual, approach T_{nd} due to the infinitely widening quasispecies and the vanishing effects of purifying selection. This is of little interest for discussions regarding valley-crossing, since these unrealistic theoretical populations never get stuck on local optima.

4.3 Discussion

It is well-known that rates of adaptation, as inferred from the fossil record [Conway Morris, 1998, Erwin et al., 1987, John Sepkoski Jr, 1998, Niklas et al., 1983, Niklas, 1997] and experimentation [Wolf et al., 2006, Elena et al., 1996], are not constant. Multiple mechanisms have been proposed to explain the variability in the tempo of evolution, but these mechanisms are usually specific to the time and place within which they are observed. Here, we have presented a new dynamic, reduced selection on advantaged lineages during fixation — the free-for-all effect (FFA) — that *universally* affects the rate of evolution. Under some conditions, FFA predicts that evolving populations will shift between modes of stasis and change. In particular, FFA provides a simple explanation for long periods of stasis separated by brief periods of rapid adaptation (a.k.a., punctuated equilibrium).

4.3.1 Selective sweeps lead to reduced selection via FFA

Because FFA is a consequence of selective advantages, the effect is always present during selective sweeps. When $T_{\text{nd}} < T_{\text{fix}} < T_{\text{st}}$ and landscapes are deceptive, FFA predicts that evolving populations will cross fitness valleys faster than expected if selection strength would not change during selective sweeps. In other words, FFA predicts $T_{\text{obs}} < T_{\text{st}}$ when the

conditions are right. Figs. 4.3e and f show examples of the valley-crossing times observed during both stasis and selective sweeps. The distribution of T_{obs} is bimodal, where the right mode of the distribution corresponds to T_{st} and consists of valley crossings that occur when the population is in stasis on a peak. The anomalous left mode of the distribution, comprised of shorter times, is upper-bound by T_{fix} and suggests that these valley crossings must occur before fixation, that is, while the reduced selection pressure of FFA is available. Heuristically, the rate of adaptation T_{obs} is described by a Markov process with two states, equilibrium and non-equilibrium, where the state of the population determines the expected speed of valley-crossing.

While FFA can be clearly observed under some conditions, it can be difficult to detect for a number of reasons. In particular, some conditions make it impossible for FFA to change the rate of adaptation, while other conditions do not affect FFA directly, but result in other concurrent changes to the rate of adaptation that mask the effects of FFA. For example, we do not see a change in the rate of adaptation when beneficial mutation tends to fix before a new one is discovered (the SSWM regime). Thus, even though FFA is reducing selection pressure, the time required to valley-cross is still more than the fixation time, and as a result, FFA does not result in more valley-crossing. In Fig. 4.3a-c, the absence of a noticeable change in the rate of adaptation is evidenced by $T_{\text{obs}} \approx T_{\text{st}}$, and indeed the distributions of T_{st} and T_{obs} in Fig. 4.3d are very similar. In short, FFA is imperceptible in the SSWM regime.

When fixation times are so large that they exceed the stochastic tunneling time $T_{\text{st}} < T_{\text{fix}}$ (the WSSM regime), multiple beneficial mutations are discovered before they can fix, leading to clonal interference and the accompanying loss of some (or many) beneficial mutations. As a consequence, it becomes difficult for the rate of adaptation to increase further [de Visser et al., 1999, Desai and Fisher, 2007, Rouzine et al., 2008] even though FFA reduces the selection strength on the advantaged lineages. In Fig. 4.3a-c, this appears as T_{obs} becoming decoupled from T_{st} (and the distributions for those times in Fig. 4.3g becoming incongruous).

Finally, in order to observe a reduction in T_{obs} relative to T_{st} resulting from FFA, fitness

valleys are required. A reduction in selection strength can only increase the rate of adaptation if strong selection was preventing progress to begin with. If there are no valleys, strong selection does not present a barrier to progress and, in fact, any reduction of selection strength in a non-deceptive landscape will slow the rate of adaptation [Ragusa and Bohm, 2021].

4.3.2 FFA depends on the relationship between fixation and valley-crossing time

We have previously asserted that the relationship between T_{fix} , T_{st} , and T_{nd} can predict if FFA results in accelerated adaptation. If this is true, then any parameter that affects these times will impact FFA dynamics. To support this claim, we show in Fig. 4.4 that changing valley width or mutation rate can affect the relationship between T_{fix} , T_{st} , and T_{nd} , and show how these changes correlate with changes in rates of adaptation T_{obs} . Conditions that make the valleys easy to cross limit the observable impact of FFA because valley-crossing already occurs rather frequently without reduced selection. Conversely, conditions that make valleys harder to cross (lower mutation rate and/or wider valleys) intensify the degree to which FFA appears to accelerate adaptation because valley-crossing occurs less frequently without the aid of FFA.

4.3.3 FFA does not require spatially-structured populations

FFA is easiest to understand in the context of spatial populations, where the leading edge of a sweep experiences reduced selection, as in Fig. 4.1. However, in Fig. 4.3c, we also observe the effects of FFA in well-mixed populations where there is no leading edge. Indeed, selective sweeps in well-mixed populations are slightly different from sweeps in spatial populations.

The underlying mechanism of FFA in well-mixed and spatial populations is fundamentally the same; FFA is the reduced selection pressure on those members of the population that carry a beneficial mutation, as a result of their competitive advantage over others in the population. In a spatial population, this reduction in selection happens along the leading edge,

because this is precisely where the competition occurs. In well-mixed populations, we must abandon the notion of *where* competition happens; competition happens between everything, everywhere, all at once. Instead of leveraging a spatial intuition, we can understand FFA in well-mixed populations by considering that any descendants of an advantaged genotype that acquire a deleterious mutation may still have a selective advantage over the average fitness for a limited time. During this window of opportunity, the remaining advantage can allow for further deleterious mutations to be accumulated on the lineage, so long as these multi-mutants continue to have an advantage relative to the average fitness. As the unmutated beneficial genotype goes to fixation, the average fitness will rise, and the temporary selective advantage of a multi-mutant is lost. However, if before this happens, one of the multi-mutants discovers another beneficial mutation (i.e., crosses a fitness valley) these dynamics will repeat as the new beneficial type goes to fixation.

Generally speaking, well-mixed populations experience dramatically faster fixation times than their spatial counterparts when controlling for population size, and for this reason the FFA window in well-mixed populations is comparatively short. In order for FFA to result in accelerated rates of adaptation, T_{ffa} must be correspondingly short so that fitness valleys can be crossed before the window of opportunity closes.

4.3.4 FFA is a non-equilibrium phenomenon

Because FFA is a condition created by a population undergoing a selective sweep, FFA is strictly a non-equilibrium effect of population genetics. While there are mathematical tools to investigate non-equilibrium dynamics in populations (notably, the Fokker-Planck equation [Risken, 1989]), for all but the most trivial dynamics (such as neutral drift), only stationary solutions to the Fokker-Planck equations are available in closed form. Likewise, most computational simulations of evolution (for example, replicator-mutator equations) rely on methods that assume a constant strength of selection (see also [Weissman et al., 2009]). Those methods will also fail to observe FFA. In agent-based simulations, on the other hand, the selective forces on an organism are an emergent property: they are determined by the

immediate competitors (in a well-mixed or structured group) and can change dramatically during selective sweeps.

4.3.5 Range expansion is a special case of FFA

Range expansions occur when a population discovers and expands into new geographic territory. It has been suggested that range expansion can result in reduced selection on the leading edge and enhance valley-crossing [Burton and Travis, 2008], similar to our observations of FFA during selective sweeps. However, range expansion differs from selective sweeps in two ways. First, a range expansion requires unpopulated territory for a population to expand into, while FFA results from selective sweeps. In a selective sweep, the more fit genotypes overtake the existing population, so no new territory is necessary. Second, because the enhanced ability for valley crossing of a range expansion ends when the new territory is full, cascades of genetic improvement (like those seen in Fig.4.3j) are not possible with range expansions alone, except under very special conditions. One such condition is an environment that is structured such that multiple ranges, each more difficult to inhabit than the last, are available to the population, as in the MEGA plate experiment [Baym et al., 2016].

Despite these differences, there are clear similarities between FFA and range expansion dynamics. In fact, if we consider empty geographic territory to be occupied by a zero-fitness “virtual wild-type” (i.e., the environment presents no resistance to expansion), we can see immediately that FFA predicts the dynamics of range expansion. Thus, FFA appears to be a general explanation for the reduced selection during range expansions that also includes cases where the wild-type fitness is non-zero. In addition, if a valley is crossed during a range expansion, then the resulting selective sweep causes reduced selection via FFA and could result in a FFA cascade.

4.3.6 FFA can lead to intermittent cascades of adaptation separated by periods of stasis

We have shown that under certain conditions, FFA can result in irregular rates of adaptation, as shown in Figs. 4.3h-m. In particular, Figs. 4.3j and l show a punctuated rate of evolution where some valley crossings occur as single events, but others are clustered, temporally, in cascades. The phenomenon of long cascades, with long waiting times in between, is often described as *punctuated equilibrium* (PE) [Gould and Eldredge, 1977].

FFA describes how strength of selection shifts when the population moves between periods of equilibrium and periods of non-equilibrium. Another concept that invokes shifts in selection strength is *Shifting Balance Theory* (SBT), based on work by Wright [Wright, 1932, Wright, 1982], that explains how isolated subpopulations experiencing reduced selection are more likely to cross fitness valleys. Gould and Eldredge [Gould and Eldredge, 1977] proposed SBT as a potential explanation for the variable rates of evolution observed in natural systems during PE. However, SBT has been criticized for being so specific a mechanism that the circumstances that allow it to occur must be rare [Gavrilets, 2003]. In a recent article, it has been suggested that the temporary shift in selection strength that occurs during range expansions should be used to extend SBT [Johnson, 2008]. Along these lines, we suggest that FFA resulting from selective sweeps should also be considered as part of SBT. With range expansion and FFA as additional explanations for temporary shifts in selection strength, SBT not only becomes a more general theory, but also becomes a far more reasonable hypothesis for PE.

4.3.7 The complexity of natural fitness functions does not preclude FFA

In the simulations shown here, gene interactions and multiple mutations on a single offspring (in a single event) are not possible, and diminishing returns are not present in the fitness function. While our system is ideal to demonstrate that FFA exists, such a system will never

be found in nature. Still, the simulation does not have any properties that could *not* be found in a naturally occurring instance of evolution. In fact, the dissimilarity between the simulation presented here, and natural systems stems from the *lack* of elements that are commonly expected in natural systems. In the remainder of this section, we consider how evolutionary dynamics of FFA would be affected by the inclusion of some features common to natural populations.

Allowing multiple mutations on the same organism in a single mutational event will reduce T_{st} , T_{ffa} , and T_{nd} because it makes crossing a valley easier. As a result, permitting simultaneous mutations (or any change that affects valley-crossing times), will affect the location of the boundaries between regions in Fig. 4.5, but will not otherwise affect the dynamics of FFA, since FFA is the result of newly discovered fitness advantages, not specifically valley-crossing.

Populations evolving on any fitness function (not just the one we have used here) will still experience FFA dynamics. However, testing evolution on a landscape other than the one we used in my simulations prevents us from distinguishing valley crossings that are shortened by FFA from the regular valley crossings that occur when the population is in equilibrium. For example, if we were to use a fitness function with diminishing returns, then because the strength of FFA depends on the fitness gains of beneficial mutations, there will be a diminishing effect of FFA over evolutionary time.

Multidimensional (multi-trait) fitness functions provide additional opportunities for FFA in the form of more mutational degrees of freedom, and more opportunities to discover beneficial mutations overall. In this work, the fitness function is one-dimensional (a single trait), so there is only ever one valley available for exploration. However, in a multidimensional fitness function many directions for exploration are always present, and each may present a different fitness challenge. Interestingly, FFA resulting from a discovery on one trait will provide a reduction in selection strength that could allow another possibly unrelated trait to drift toward a new peak. This is particularly interesting if we consider a scenario where crossing an easy valley on one dimension could provide a large enough benefit to cross more

prohibitive fitness valleys on other dimensions. Thus, a FFA cascade that accumulates a large net benefit before fixation could eventually result in crossing a valley that would be strictly impassable in the absence of FFA.

One consequence of multidimensionality is the possibility for gene interactions (epistasis), the condition in which the allele of one gene alters the fitness effect of an allele of another gene. A system with gene interactions will still experience FFA, but the gene interactions will create confounding signals that make detecting FFA challenging. When two beneficial mutations have a combined fitness benefit larger than the sum of their individual benefits, the resulting fitness advantage will strengthen FFA; the larger benefit will cause a bigger reduction of the selection strength. The opposite can also happen: two beneficial mutations together may have a weaker fitness advantage than the sum of their benefits. In this case, FFA is weakened by the loss of fitness. These kinds of gene interactions, and others, make clearly observing FFA harder, but they do not preclude FFA from occurring.

In this work, we use asexual populations, and so the rate of adaptation in region III is reduced due to clonal interference. Sexual populations could, in theory, avoid clonal interference via recombination, and thus experience quicker fixation times in region III than asexual populations. Sex has also been shown to increase the strength of purifying selection, making valley-crossing harder overall [Kondrashov, 1988, Rice and Chippindale, 2001] relative to asexual populations. These changes may affect the observed rates of adaption in sexual populations, but will not affect the dynamics of FFA.

4.3.8 Observing FFA in natural systems

We have argued that FFA dynamics are an unavoidable feature of evolutionary dynamics, even though its ramifications are not always felt. We thus suspect that some natural populations experience cascades of adaptation mediated by FFA. However, it is likely very difficult to disambiguate FFA-induced spurts of adaptation from confounding effects, such as multiple adaptations made possible by a change in the external environment. To make an unambiguous case, it is probably necessary to have detailed genomic and fitness data on the line of descent,

documenting that the population crossed valleys that standard population genetics would deem uncrossable. However, it is not hard to imagine scenarios where FFA could play a major role.

In the evolution and spread of cancer, for example, multiple somatic mutations appear in rapid succession. It is not unreasonable to hypothesize that a single (or several) “driver” mutations [Martínez-Jiménez et al., 2020] create a non-equilibrium environment of reduced selection that leads to the accumulation of multiple somatic mutations that, in turn, can cross valleys that otherwise could not be crossed.

Viral dynamics provides another possible venue to look for manifestations of FFA. Pandemics are often started by a beneficial mutation that results in a selective sweep, often with global reach (as in the SARS-Cov2 pandemic) and multiple waves of variants can sweep through local populations. Indeed, a punctuated pattern of evolution has been observed in influenza evolution [Wolf et al., 2006]. Given that, in the case of SARS-Cov2, tens of mutations can accumulate within a single immunocompromised individual (where selection is relaxed) [Pedro et al., 2021, Cele et al., 2022], it is not unreasonable to imagine that during the initial period of a strong wave, FFA might sufficiently reduce selection on the sweeping variant such that deleterious mutations can rapidly accumulate and lead to new variants.

4.4 Conclusion

We have introduced the Free-for-All effect: a new non-equilibrium component of evolutionary dynamics that affects the tempo and mode of evolution. We have discussed the conditions that make it possible to detect the impact of FFA on evolutionary rates, and demonstrated how FFA can result in cascades of discoveries of beneficial mutations. The FFA effect, and the cascades that can result from FFA, provide a novel explanation for the observation of a punctuated pattern of evolution (a.k.a., punctuated equilibrium).

If a very large fitness advantage is discovered, as a result of a FFA cascade or any other means, FFA can become so strong that selection is removed almost completely. Valleys of

unfathomable depth could be crossed, and entirely new body plans could conceivably emerge that are unimaginable under standard evolutionary dynamics, leading to further explosions of diversity [Gould, 1989]. At the same time, the diminishing quantity of beneficial mutations in the vicinity of local optima make fixation ever more likely, gradually bringing all cascades to an end eventually. Thus, FFA provides an elegant explanation for the observation of punctuated evolution: FFA explains why cascades occur (reduced selection strength during selective sweeps), and why cascades come to an end (the return to strong purifying selection after fixation). Furthermore, since FFA is active during all selective sweeps, it offers a mechanism for shifts in selection strength that does not carry the burden of requiring rare circumstances or environmental effects.

Free-for-all suggests that any time a system undergoes change, it enters an unstable state that may result in more change.

While this work studied the impact of FFA on the evolution of a single trait, it is immediately clear that mult-trait landscapes will only enhance this effect, as a selective sweep due to the improvement of one trait can reduce selection on *another* trait. Cascades that occur *in parallel* (rather than serially) are expected to lead to larger fitness gains, more powerful reductions in selection strength, and therefore to even more pronounced periods of rapid adaptation. These will be studied in future work.

4.5 Methods

The experiments in this work were conducted using the MABE framework, a modular research tool that supports the development and execution of agent-based digital evolution research [Bohm et al., 2017].

4.5.1 Computational model

The experiments in this work are designed to prohibit gene interactions, multiple mutations on a single offspring, and diminishing returns in the fitness function.

The digital organisms in this work consist of a genome with a single locus that encodes

the mutation number k . This allele mutates with a probability μ ; mutations result in $k + 1$ or $k - 1$ with equal probability. We use a periodic fitness function (Fig. 4.2) that can be written in terms of k and the valley width v as

$$w(k) = (1 + s_b)^n (1 - s_d)^m, \quad (4.2)$$

where the value $m = k \pmod{v}$ determines how far the genome is within the valley, and $n = \frac{k-m}{v}$ simply numbers the peaks that have been climbed. This is how the function is presented in the main text. For computational purposes, an alternative definition is used. The alternative form of the fitness function is given by

$$w(k) = r^{Bn - Dm}, \quad (4.3)$$

where r is a parameter that controls the strength of selection. B and D are defined in relation to s_b and s_d : $s_b = r^B - 1$ and $s_d = 1 - r^{-D}$. All the experimental configurations for the fitness function and evolutionary algorithm are provided below in terms of v, r, B , and D .

We carry out evolution experiments using three population structures: a one-dimensional array with periodic boundary conditions, a regular two-dimensional grid with periodic boundary conditions, and a well-mixed population. In the one-dimensional population structure, each agent is in the center of a 1×3 neighborhood and has two neighbors. In the two-dimensional population structure, each agent is in the center of a 3×3 neighborhood and has eight neighbors. The well-mixed populations are structured identically to the one-dimensional populations, with the modification that once per generation immediately before parent selection the location of each agent is randomized to mix the population, effectively removing the dimensionality. We still only allow each organism to have at most two offspring (as in binary fission) to avoid the explosive growth rates that usually accompany a truly fitness-proportional, well-mixed population. We use a generational evolutionary model where the population size is fixed, and the entire population is replaced every generation. During re-

production, a parent is selected from each neighborhood using fitness-proportional selection. The offspring is placed at that neighborhood’s center location in the next generation.

4.5.2 Data collection

We measure four properties of the evolutionary system: the time it takes a stationary clonal population to cross a fitness valley T_{st} , the time to cross the same valley without purifying selection T_{nd} , the fixation time of a beneficial mutation T_{fix} , and the time to cross fitness valleys when the system is allowed to evolve for a long period T_{obs} . To record T_{st} , we initialize a population at a local optimum. When we detect that an agent has discovered the next higher peak and has progeny after an additional 100 generations (i.e., it has established) we record the time in generations of the discovery. We use the same process to record T_{nd} , but with an altered fitness function where the valley has no depth ($s_d = 0$). To record T_{fix} , $N - 1$ agents (where N is the population size) are initialized on a local optimum, with the remaining agent initialized on the next higher peak and tagged with a heritable marker. The simulation is run with $\mu = 0$, until either all organisms have the marker (fixation) or none do (extinction). Upon fixation, the time in generations is recorded. On extinction, we disregard the sample. For T_{st} , T_{nd} , and T_{fix} each time a recording is made the population is re-initialized and more data is collected. To record T_{obs} , we initialize a population on a local optimum and let it evolve. We use the same method of detecting the discovery and establishment of valley-crossing events used when recording T_{st} and, T_{nd} , however, rather than re-initializing the population after each valley crossing, we record the event and allow the simulation to continue.

4.5.3 Experimental parameters

Experiments are run for a fixed number of CPU hours (4 hours unless otherwise noted). Table 4.1 organizes all the experimental conditions into a single reference.

Population structure	1-dimensional array w/ periodic boundary
Population Sizes	10^1 to 10^5
Mutation rates	0.05, 0.1, 0.15
Selection strength (r)	1.5
Fitness function parameters	$B = 7, D = \frac{4}{v-1}, v = [4, 5, 6]$

Population structure	2-dimensional grid w/ periodic boundary
Population Sizes	4×10^2 to 25×10^5
Mutation rates	0.1
Selection strength (r)	1.35
Fitness function parameters	$B = 10, D = 0.5, v = 6$

Population structure	Well-mixed
Population Sizes	10^2 to 25×10^5
Mutation rates	0.08
Selection strength (r)	2
Fitness function parameters	$[B = 4, D = 1.3, v = 3]$

Table 4.1: Experimental conditions used to generate all data in this report.

4.5.4 Distribution plots

In the Section 4.2, we show plots that show distributions of T_{st} , T_{nd} , T_{obs} , and T_{fix} . We create the distribution plots by binning the data using variable-sized bins with ranges defined by Eqn. (4.4) and then plotting the resulting histograms on a log/log scale.

$$\text{bin}_i = [1.15^i, 1.15^{i+1}) \text{ for } i \in [0, 100] \quad (4.4)$$

Chapter 5

Augmenting evolutionary algorithms with Super explorers

This chapter is based on the article “Augmenting Evolution with Bio-Inspired ‘Super Explorers’” [Ragusa and Bohm, 2022].

5.1 Introduction

All search and optimization algorithms, including both natural selection and computer models of evolution, are subject to the fundamental limitations of the no-free-lunch theorems [Ho and Pepyne, 2002], and particularly to the explore-exploit tradeoff [Millidge et al., 2021]. Managing this tradeoff is typically a main concern in the development of evolutionary algorithms. In this work, we introduce the “super-explorer method” for augmenting selection methods used in evolutionary algorithms. The super-explorer method allows us to tune the trade-off between exploration and exploitation to better align with the ruggedness of a fitness landscape.

The super-explorer method augments other selection methods by adding agents (the “super explorers”), that are allowed to drift (are not subject to selection). Discoveries made by super explorers are immediately available for refinement by the underlying algorithm.

In essence, super explorers allow for the simultaneous exploitation of new discoveries while exploration continues in parallel. We compare three implementations of the super-explorer method across a wide range of configurations. We test each implementation on two different kinds of fitness landscapes: an NK-fitness landscape [Kauffman et al., 1993], and a saw-tooth fitness landscape [Ragusa and Bohm, 2021]. We show that for both types of fitness function, there are configurations where evolution augmented with super explorers preforms better than without.

5.1.1 Biological inspiration

Natural selection does not act with consistent strength across time and space; Shifting balance theory [Wright, 1932, Wright, 1982], range expansion [Slatkin and Excoffier, 2012, Peischl et al., 2013, Peischl and Excoffier, 2015, Peischl et al., 2015, Gilbert et al., 2017, Burton and Travis, 2008], environmental noise [Wang and Zhang, 2011, Van Egeren et al., 2018, Ragusa and Bohm, 2021], population size changes [Jain et al., 2011, Ochs and Desai, 2015, Rozen et al., 2008], sexual selection [Bohm et al., 2019], and mass-extinction events [Mathias and Ragusa, 2016, Engholdt and Mathias, 2021] all describe scenarios where selection strength changes over time, space, or both.

As a consequence of these fluctuations, natural populations experience both increased and decreased effectiveness at exploring the fitness landscape over time. Furthermore, during the periods of increased exploration, the population may discover new fitness peaks that, under stronger selection, would have been unlikely or impossible to discover. Range expansion events are one particular example of a process that can cause a reduction in selection strength. As a species enters new territory, a lack of competition can result in an accumulation of deleterious mutations in organisms on the leading edge of the expansion for as long as uncolonized territory remains [Burton and Travis, 2008]. The continued accumulation of mutations in one lineage can result in adaptation via valley-crossing, a process known to be critical for evolutionary adaptation [Covert et al., 2013, Oliveto et al., 2018].

In this chapter, we present a method designed to augment existing optimization algorithms

with the exploration-boosting power of range expansion, without requiring that populations have spatial structure. In addition, our method is able to maintain the exploratory advantages of the leading edge of a range expansion indefinitely. The key insight underlying our method is that for sufficient exploration to occur, mutations must be able to accumulate in some—but not all—lineages before purifying selection acts on them. To achieve this, we introduce super explorers to the population. These agents artificially experience the same mutation-accumulation as an organism at the leading edge of a range expansion.

5.1.2 Similar algorithms

Super explorers are agents that are set aside from the main population and always have exactly one offspring every generation. Since the lineages of super explorers propagate regardless of fitness, they experience sustained genetic drift. The super explorer method shares some similarities with other algorithms that ignore the fitness gradient to boost the effectiveness of a search algorithm.

Particle Swarm Optimization (PSO) [Poli et al., 2007] is a population based search algorithm that models agents as particles and simulates forces that attract the particles to the highest observed fitness (analogous to selection). However, the individual particles in PSO do not move directly to the higher gradients, but instead chaotically trend towards it, often exploring areas of the fitness landscape more distant from known optima.

Lexicase selection [Helmuth et al., 2014, La Cava et al., 2016] and real-valued Tournament selection [Ragusa and Bohm, 2021] are examples of evolutionary algorithms that probabilistically ignore some or all fitness gradient information. Novelty search [Lehman and Stanley, 2011] is an evolutionary algorithm that foregoes following the fitness gradient altogether and instead focuses on collecting a catalog of functionally distinct solutions, evaluating them for relevance to the fitness function *ex post facto*.

Systems augmented with super explores are different from the examples above in key ways. While super explorers do experience a form of occasional selection during the decay-replace process (described in Methods), they do not individually experience a fitness gradient, even in

the form of selection for novelty. While some other methods, such as lexicase and real-valued Tournament, can occasionally allow for several generations of drift on some lineages, super explorers are maintained separately from the main population, ensuring their lineages are repeatedly free of selection.

5.2 Methods

5.2.1 Augmenting evolution with super explorers

Super explorers can augment any well-mixed discrete-generation agent-based selection method (compare Algorithm 3 with Algorithm 4). Super explorers are added to a selection method by dividing a population into two mutually exclusive pools, the *active selection pool* (or *ASP*) and *super-explorer pool* (or *SEP*) (see Fig. 5.1). The ASP evolves via the rules of an externally defined selection method (such as, roulette, tournament, lexicase, etc.), except that parents can be drawn not only from the ASP, but also from the SEP. On the other hand, the agents in the SEP each produce one offspring, with the standard mutation load, regardless of fitness. While lineages in the ASP survive and perish following the rules of the selection method in use, lineages in the SEP end only by decay. Whenever an agent in the SEP is about to reproduce, there is a chance (the *decay rate*) that the agent will be removed and replaced by some other agent in the population. The number of agents in the ASP and SEP, the decay rate, the method used to replace decayed SEP agents, the mutation settings, and the selection method used in the ASP (and related settings) are all parameters established by the user.

Because the super explorers reproduce regardless of their fitness, they are able to accumulate mutations without consequence. The decoupling of survival from fitness allows super explorers to explore the fitness landscape in directions that selection might prohibit; they may descend into, and possibly cross, fitness valleys. The mutational paths of super explorers are, in fact, random walks (i.e., undirected and unfocused), so they are an inefficient search algorithm on their own. However, by introducing super explorer decay and replacement to

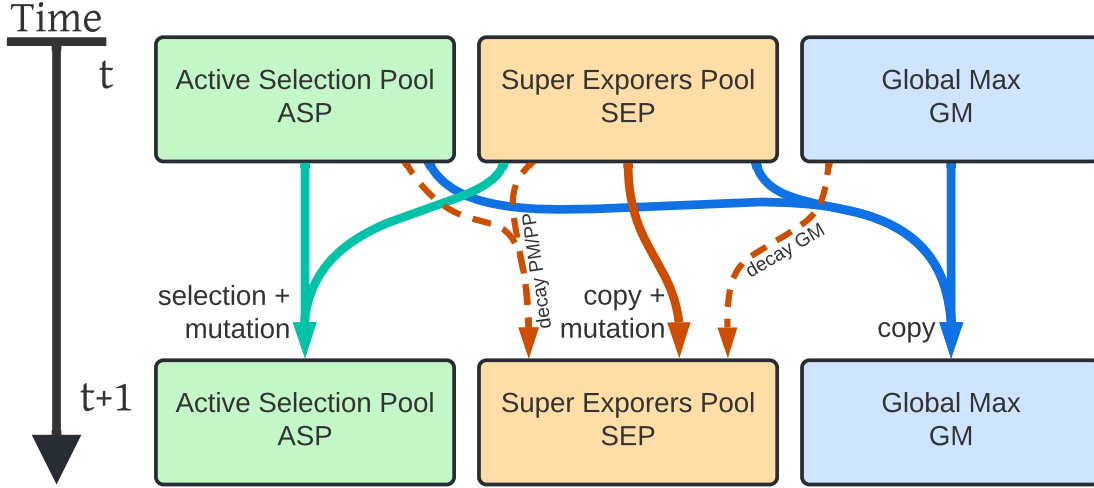


Figure 5.1: A diagram of an evolving system augmented with super explorers. Here, the population is divided into the *active selection pool* (ASP) and the *super-explorer pool* (SEP). While selection in the ASP (green arrows) is determined by a standard selection method (such as roulette, tournament, lexicase, etc.) with parents drawn from the ASP and SEP, agents in the SEP are free from selection and simply accumulate mutations (red solid arrow). From time to time, agents in the SEP decay and are replaced. Replacements (red dashed arrows) are drawn from either the ASP and SEP, or the global max, depending on the replacement method, where the global max maintains a copy of the highest scoring agent seen to date (blue arrows).

the algorithm, we provide a degree of focus to help keep the search process “on track.” The search pattern of the super explorers can be seen as a genotypic radiation where the center of the radiation is determined by the replacement method, the radiation distance is determined by the decay rate, and the density of the search is determined by the number of agents in the SEP.

Since the selection method used to choose parents in the ASP has access to both the ASP and the SEP, super explorers can migrate into the ASP if their fitness is competitive (i.e., if a super explorer drifts across a valley, they can be discovered by and incorporated into the ASP). Thus, a selection method augmented with super explorers has the advantage of simultaneously exploring freely (in the SEP) while exploiting discoveries (in the ASP), compared to a non-augmented method that tries to balance the explore-exploit trade-off with only an ASP.

Algorithm 3: A typical Evolutionary Algorithm

Result: A population of adapted digital organisms.
population = [Organism() **for** _ **in** range(popSize)];
for g **in** range(numGenerations) **do**
 for org **in** population **do**
 | evaluateFitness(org);
 end
 newPopulation = [];
 while len(newPopulation) < popSize **do**
 | newPopulation.append(mutate(selectFrom(population)));
 end
 population = newPopulation;
end
return population;

5.2.2 Tunable parameters of super explorer systems

When employing the super explorer method, the user must first choose a population size, selection method and mutation scheme. These choices, as well as the tuning of any associated parameters, will have significant consequences on rates of adaptation. Regardless, these choices have no bearing on the super explorer parameters themselves. As long as the selection method is agent-based, well-mixed, and has discrete-generations, it can be augmented with super explorers as presented in this work. Some extensions of super explorers, which allow for the relaxation of some of these constraints, are addressed in the discussion.

There are only three parameters that relate directly to super explorers: SEP size, decay rate, and replacement method, which we describe below.

5.2.2.1 Parameter: super-explorer pool size

SEP size can be set to any integer value, from 0 to population size, and controls the intensity of the search conducted by the super explorers. Note that the ASP size is not a parameter and is simply defined as population size minus SEP size.

Algorithm 4: An Evolutionary Algorithm Augmented with Super Explorers

```
Result: A population of adapted digital organisms.
population = [Organism() for _ in range(popSize)];
superExplorers = [Organism() for _ in range(sepSize)];
for  $g$  in range(numGenerations) do
    for  $org$  in [population+superExplorers] do
        | evaluateFitness(org);
    end
    newPopulation = [];
    while len(newPopulation) < popSize do
        | newPopulation.append(mutate(selectFrom([population+superExplorers])));
    end
    population = newPopulation;
    newSEP = [];
    for  $org$  in superExplorers do
        | if random()  $\leq$  decayProbability then
            | newSEP.append(mutate(selectFrom([population+superExplorers])));
        | end
        | else
            | newSEP.append(mutate(org));
        | end
    end
    superExplorers = newSEP;
end
return population;
```

5.2.2.2 Parameter: decay rate

The decay rate is the per-generation probability that each agent in the SEP is subject to replacement, and directly controls the explore-exploit trade-off of the SEP. The decay rate can be set to any value r , such that $0 \leq r \leq 1$. The decay rate determines the average amount of time, $\bar{t} \sim 1/r$, that a super-explorer lineage persists before replacement, which in turn determines how far super explorers can drift (i.e., explore). While lower decay rates can result in more exploration, they can also be wasteful if the valleys in the fitness landscape do not require such distant explorations.

5.2.2.3 Parameter: replacement method

The replacement method defines the process used to replace a decayed agent in the SEP. In this work, we consider three replacement methods, shown in Fig. 5.1, though many others are possible.

The first method, the *PopSelect* (PS) method, replaces decayed SEP agents with a new agent generated using the same selection process used for reproduction in the ASP.

The second method, the *PopMax* (PM) method, replaces decayed SEP agents with the max of the population (i.e., the max of all the agents in the combined ASP and SEP).

The third method is the *GlobalMax* (GM) method. In addition to the ASP and SEP, this method requires that the global max, a copy of the highest fitness agent seen to date in either pool, is maintained and used to replace decayed SEP agents.

5.2.3 Fitness functions

We use two types of fitness functions to collect performance data for the parameter sweeps. First, we chose a typical NK-landscape to provide an intuition regarding the performance of super explores in a familiar context. Then, we turn to a saw-tooth function to provide a different perspective on the operation of the super-explorer method under more controlled conditions.

5.2.3.1 Fitness function: NK-landscape

The first fitness function is the well studied NK-landscape [Kauffman et al., 1993]. This landscape has a rugged topology and a global optimum. The NK function, $\text{NK_eval}(i, g_i|N, K)$, decides the fitness contribution of the i -th gene, g_i . The fitness landscape is randomly generated with parameters $N = 20$ and $K = 5$ which control genome size, $G = \{g_1, \dots, g_N\}$, and strength of epistasis (in our implementation, $K = 1$ represents no epistasis), respectively. The score assigned to each organism by NK is given by

$$\text{score}(G) = \frac{1}{N} \sum_{i=1}^N \text{NK_eval}(i, g_i|N, K) \quad (5.1)$$

5.2.3.2 Fitness function: saw-tooth landscape

The second fitness function, the saw-tooth fitness function, [Ragusa and Bohm, 2021] defines a saw-tooth mapping function from genomic values to scores along an infinite set of ever higher fitness peaks (like the teeth of a saw) as shown in Fig. 5.2.

The saw-tooth landscape has two special properties. First, it has no global optimum. Second, every fitness valley is self-similar; regardless of absolute position on the fitness landscape, the difficulty of descending into and the benefit of crossing every valley is the same. These two properties together result in a landscape that presents a constant challenge to an evolving population, without diminishing returns (an idea associated with Fisher’s geometric model [Orr, 2005]). In this landscape, agents have 10 independent genes, $G = \{g_1, \dots, g_{10}\}$, where each gene is a single integer value. Applying the mapping function to each gene results in 10 gene scores that are summed to produce the agent’s overall score. The saw-tooth function $\text{saw}(x|w, p, b)$ is specified by a valley width (w), a fitness penalty per mutation into each valley (p), and a fitness benefit per valley crossed (b):

$$\text{score}(G) = \sum_{i=1}^{10} \text{saw}(g_i|w, p, b) \quad (5.2)$$

Here we test four versions of saw-tooth functions. In all four, $p = -0.05$ and the $b = 1.0$. We varied w to create four versions of the function with different difficulties, from $w = 4$ (easy) to $w = 7$ (hard). These functions are shown in Fig. 5.2.

5.2.4 Roulette selection

We used roulette selection in the ASP. Roulette selection is known to suffer from a diminishing selection pressure as the absolute value of agents’ scores increases during evolution. In order to ensure that equivalent relative increases to score result in the same relative increase in offspring production, the scores generated from both the NK-landscape and the saw-tooth landscape are exponentiated before selection occurs.

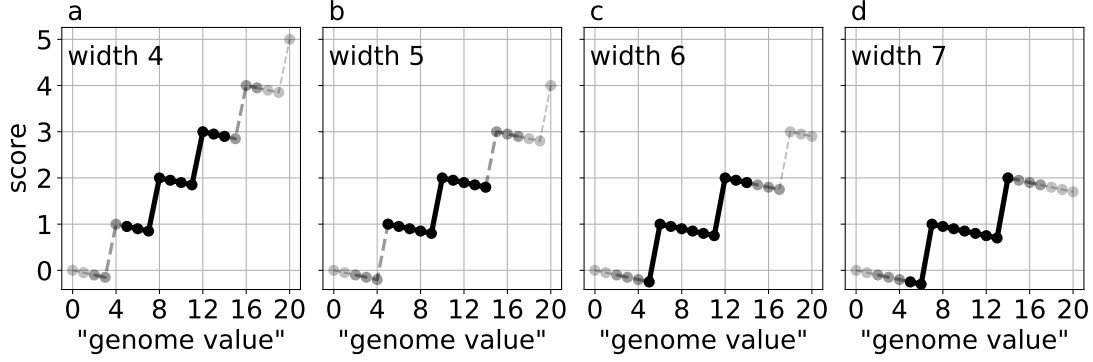


Figure 5.2: The saw-tooth functions that map gene values to “scores”, used in Fig. 5.4. Panels [a] through [d] show the valley widths (w) 4 through 7, respectively. The values for penalty ($p = -0.05$) and benefit ($b = 1.0$) are the same in all 4 functions.

$$f(G) = \exp(\text{score}(G)) \quad (5.3)$$

5.2.5 Experiment conditions

In order to determine the effectiveness of super explorers, we compare roulette selection with and without super-explorer augmentation on two classes of fitness function (NK and saw-tooth). We use a total population (ASP + SEP) size of 1024 in all experiments. For each combination of selection regime and function, we run a three-dimensional sweep of SEP size, decay rate, and replacement method. We run decay rates $\{0.001, 0.002, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1.0\}$, pool sizes $\{0, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024\}$, and replacement methods $\{\text{GM}, \text{PM}, \text{PS}\}$.

We run 101 replicates for each condition using the NK-fitness landscape ($N = 20$ and $K = 5$) for 10,000 generations and record the final maximum fitness from each replicate. Each replicate evolves on the *same* NK landscape, which is randomly generated *a priori*. A bit genome was used where mutations are introduced with a bit-flip mutation operator, with a 0.0005 per-site mutation rate (i.e., a 0.01 per-agent mutation rate).

We run one replicate for each condition using the saw-tooth landscape for 40,000 generations and record the number of valleys crossed. In addition to the parameter sweep described

above, we used four saw-tooth functions with $w = \{4, 5, 6, 7\}$ (shown in Fig. 5.2). An integer genome was used where mutations are introduced with a point-offset mutation operator that modifies a genome value by ± 1 , with a 0.05 per-site mutation rate (i.e., a 0.5 per-agent mutation rate, as each agent has 10 loci).

The MABE evolution framework was used to run experiments [Bohm et al., 2017].

5.3 Results

5.3.1 NK landscape

Figure 5.3 displays the data collected from 101 replicate evolutionary experiments on an NK-landscape with parameters $N = 20$ and $K = 5$. The figure presents three color maps (labeled [1], [2], and [3]) each showing results generated using a different replacement method (GM, PM, and PS). The cells in each color map show the averages, across replicates, of the maximum fitness detected at the end of each replicate. Note that values in each plot associated with SEP size = 0 are the same, since decay rate only matters if there are super explorers. The fitness differences between the max-fitness, min-fitness, and control configurations of each panel are computed and checked for significance with a two sample z -test (shown in Table 5.1).

The NK-landscape results show that the GM and PM methods provide similar results that are quite different from the PS data. In the GM and PM results, we see that the addition of any super explorers improves performance under almost all conditions. In addition, as SEP size increases, low decay rates tend to improve adaptation while high decay rates tend to have the opposite effect. In fact, a population made up entirely of super explorers (SEP size = 1024) combined with decay rate = 1.0, shows little to no signs of improvement over a non-augmented population.

We note a band of yellow color in the PS data, highlighting the conditions with the greatest final scores. This band appears to show a trade-off between SEP size and decay rate. To the right and above the band, we see values that conform closely to the non-augmented

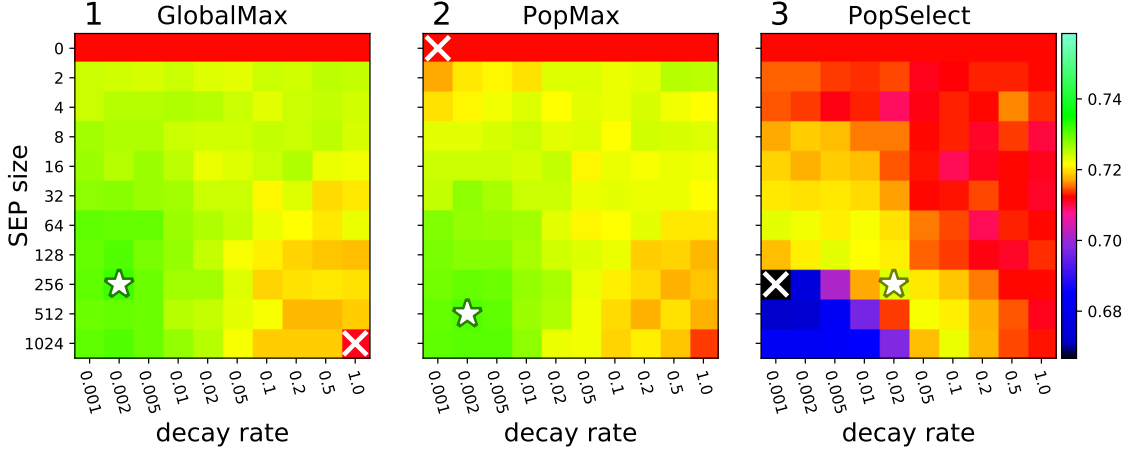


Figure 5.3: Average maximum final fitness on an NK landscape from 101 replicates using super explorers with three replacement methods: [1] GM, [2] PM, and [3] PS. The color range is set so that red is associated with the control condition; SEP size = 0 (i.e., no super explorers). The stars and crosses indicate the highest-fitness and lowest-fitness configuration of each panel, respectively. The largest increase and decrease of fitness, relative to the control, are shown in Table 5.1.

system, while to the left and below the band, we see that final recorded scores are lower than the non-augmented system. This area of low values correlates with a large SEP size and low decay rate.

Note that the values in the column associated with decay rate = 1.0 in the PS data all match the values associated with SEP size = 0. In these conditions, SEP agent replacement happens every generation and uses the same replacement method used in the ASP. As a result, the two pools act as a single ASP (small fluctuations are the result of sampling noise).

Decay Method	Largest Increase	Largest Decrease
GM	+2.76% *	−0.15%
PM	+2.61% *	±0.00%
PS	+1.51% *	−6.44% *

Table 5.1: The fitness differences between the max-fitness, min-fitness, and control configurations of the NK-landscape data in Figure 5.3. ‘*’ indicates a p-value of $p < 1. \times 10^{-6}$ (two sample z -test).

5.3.2 Saw-tooth landscape

Figure 5.4 shows color maps illustrating the number of valleys crossed at the end of the 40,000 generations, with agents evolved on saw-tooth fitness functions. Each of the letters [a] through [d] indicate a different valley width, from 4 to 7. The numbers [1] through [3] indicate the replacement method used (GM, PM, and PS). Note that values in each panel associated with SEP size = 0 are the same, since decay rate only matters if there are super explorers. As we saw in the NK data, the saw-tooth landscape results show that the GM and PM methods provide similar results that are quite different from the PS data.

Across all panels of the saw-tooth data, there are vertical trends, associated with SEP size, and horizontal trends, associated with decay rate. As we move from top to bottom and introduce a higher ratio of super explorers (larger SEP size) relative to the size of the ASP, we generally see a smooth transition from SEP size = 0 to SEP size = 1024. Conversely, as we move from low decay rates to higher decay rates, particularly larger SEP sizes, we see that the change in performance is not as simple to describe. In all GM and PM results, while we see that large SEP combined with a high decay rate reduces rates of adaptation relative to the non-augmented system, the effect of large SEP and low decay rate is not constant. Rather, in the data for large SEP and low decay rate, the final scores flip from worse than non-augmented to better as we move from less deceptive to more deceptive landscapes.

In the PS results Fig. 5.4[a,b,c,d][3], we see a different trend. As mentioned in the NK data analysis, when decay rate = 1.0 agents in the SEP experience the same selection as agents in the ASP (i.e., replaced every generation using the selection method from the ASP). Therefore, as in the NK data, the two pools act as a single ASP (small fluctuations are the result of sampling noise). Unlike the NK data, we see that larger SEP sizes perform as well as or better than the non-augmented system, except in the lower left of [a, b].[3] where SEP is very large, and decay rate is very low. The highest rates of valley crossing appear to correlate with large SEP sizes and middling decay rates (i.e., around 0.05 to 0.1).

Fig. 5.4[3] shows the saw-tooth PS result. As in the NK data, the column associated with




Figure 5.4: Total number of valleys crossed after 40k generations on saw-tooth landscapes. Horizontal panel groups [a], [b], [c], and [d] show data for $w = \{4, 5, 6, 7\}$ saw-tooth functions, respectively. Vertical panel groups [1], [2], and [3] show data for replacement methods GM, PM, and PS, respectively. The stars and crosses indicate the highest-fitness and lowest-fitness configuration of each panel, respectively. Note: the scales on each color bar (i.e., for each row) are different.

decay rate = 1.0 uses the same replacement method used in the ASP. Since the replacements occur every generation, the two pools act as a single ASP (small fluctuations are the result of sampling noise).

5.4 Discussion

In this work, we show that augmenting selection (in this case roulette-wheel selection) with super explorers can improve rates of adaptation across a range of landscapes and parameter settings. In this discussion, we are primarily interested in describing the super-explorer augmentation method.

The conditions where SEP size = 0 show the behavior of the non-augmented system. In these conditions, there are no super explorers, so decay rate has no effect. Hereafter we will refer to these conditions as the 'control' and they will serve as the baseline for all of our comparisons.

In our experiments, we test conditions with different SEP sizes, but we do not alter the number of agents in the total population ($ASP + SEP = 1024$) or the mutation inflow per agent. Therefore, the populations always experiences the same number of evaluations (i.e., uses the same amount of computational power) and have the same potential for discovery in the form of mutations. As a result, differences in performance between the control and other conditions must be accounted for by other means. Perhaps the observed differences can be explained in terms of a trade-off between exploration and exploitation.

5.4.1 The explore-exploit trade-off of super explorers

In evolutionary systems (natural or computational) the explore-exploit trade-off explains how the conditions that allow for effective navigation on smooth fitness landscapes hinder navigation on deceptive fitness landscapes and *vice versa*. On a smooth landscape the paths towards higher fitness are easy to discover, so unfocused distant explorations are wasteful: it is prudent to spend energy exploiting the local fitness gradient. Conversely, on deceptive landscapes the paths to higher fitness often require crossing fitness valleys, so focusing on

local search only results in short-term superficial gains: it is prudent to spend time on long shots.

Our initial inspiration for the super explorer method was the biological phenomenon of range expansion. There, we observed that the freedom from selection along the leading edge of a range expansion creates lineages in which mutations could accumulate, while the remainder of the population experiences purifying selection. To mimic this effect, we designed a method that divides a population into two pools: the SEP that specializes in exploration, and the ASP that specializes in exploitation. In practice, we see that either pool can be set up to specialize in exploration or exploitation, since both the selection method used in the ASP, and the super-explorer decay rate parameter used in the SEP, can be tuned to maximize exploration or exploitation.

5.4.1.1 Decay rate

While it may seem intuitive that low decay rates enhance exploration and inhibit exploitation, and that high decay rates do the opposite, it turns out that explaining why this is the case is not trivial. There are two effects to consider: 1) the amount of genetic change that a lineage in the SEP can accumulate before it decays (and is replaced) and 2) the strength of selection resulting from the rate of replacement.

Generally speaking, the amount of genetic change that any lineage can accumulate is related to the strength of selection (which inhibits change), and how long the lineage can avoid extinction. SEP lineages (unbroken phylogenies in the SEP) are special in that they are free from selection, and so their potential for genetic change is only limited by how long they persist, which in turn depends on the decay rate. Meanwhile, the decay rate *also* establishes how often replacements occur and, because SEP lineages are free from selection, replacement is the only selective force at work in the SEP. Hence, decay rates control the frequency of selection events, and this frequency correlates directly with selection strength. Note that the conditions likely to enhance exploration, high drift potential and low selection strength, are conditions generated by low decay rates while the conditions likely to enhance exploitation,

low drift potential and strong selection, are conditions generated by high decay rates.

5.4.1.2 Replacement method

The three replacement methods we chose to test correlate with a range of selection strengths. The GM method provides the strongest selection strength because it always maintains the best solution. The GM method insures that small fitness changes (potentially undetectable by the ASP's selection method) are always exploited, but this comes at a cost: the inability to forget can inhibit escape from local optima. In fact, when the GM method is combined with an SEP size equal to population size and decay rate equal to 1.0, the result is elitist behavior; every agent tests one mutation, and then either becomes the new global max or is forgotten. Compared to the GM method, the PS method results in weaker selection, which cannot exceed the selection strength of the ASP's selection method. This is because at decay rate = 1.0, all agents in the population are replaced using the ASP selection method every generation (regardless of SEP size). The PS method represents a middle ground in terms of selection strength. However, since we are using a relatively large population size relative to decay rate, forgetting high quality solutions is also unlikely, so the GM and PM methods generate similar results.

5.4.1.3 SEP size

The SEP size parameter determines how much each pool (ASP and SEP) drives the system. As SEP size increases, particularly when decay rates are low, there are more low quality solutions in the total population, which increases the chance that low quality solutions will be selected during ASP reproduction. As a result, larger SEP sizes result in weaker selection in the ASP. This provides another mechanism that helps explain why low decay rates can enhance exploration.

5.4.2 Analyzing the NK data

Fig. 5.3[1,2] show the results of using the GM and PM methods to evolve populations on the NK-landscape. In these panels, the lower-left corner (SEP size = 1024, decay rate =

0.001) maximizes exploration, while the lower-right corner (SEP size = 1024, decay rate = 1.0; i.e., elitism) maximizes exploitation, as described above. In the GM and PM results, we see that enhancing exploration improves performance, while maximizing exploitation results in performance that is not significantly different from the control. The fact that the results of elitism are similar to the control suggest that the ASP is already generating elitist-like behavior that hinders exploration. It follows then that improved adaptation associated with lowered decay rates results from increased exploration.

The sudden jump in fitness from $\text{SEP} = 0$ to $\text{SEP} = 2$ in Fig. 5.3[1,2] occurs because the GM and PM methods select max *perfectly* when making replacements into the SEP. Using the GM method, it is *impossible* to forget a high-value solution once it has been discovered; forgetting is only very unlikely when using the PM method. As a result, the GM and PM methods are able to identify and exploit fitness improvements that the ASP selection method alone may not detect. Finally, the decrease in fitness along the right-most column (decay rate = 1.0), from $\text{SEP} = 2$ all the way to $\text{SEP} = 1024$, shows that the behavior of the ASP alone is not exactly the same as elitism. While the ASP can allow for the accumulation of small (i.e., nearly-neutral) deleterious mutations, a system experiencing elitism can never move in a direction that results in *any* loss of fitness. As a result, the highest fitness occurs at decay rate = 1.0 and SEP size = 2, indicating that the best performance is generated by a large ASP (that has the ability to explore, if only locally), augmented by a SEP that does not forget.

Fig. 5.3[3] shows the results of the PS method used to evolve populations on the NK-landscape. In this panel, as in the GM and PM methods, the lower-left corner maximizes exploration, while the lower-right corner maximizes exploitation. However, in these results, the entire right side of the panel evolves similarly to the control. The low fitness recorded in the low left is the result of weak selection that is unlikely to find effective solutions, and it is likely to forget what it does find. The most interesting feature in this panel is the yellow band—indicating the highest scores—that appears to show a trade-off between SEP

size and decay rate. SEP size = 1024 and decay rate = 0.1 marks the bottom of this band, and suggests that lineages that survive for about 10 generations are optimal when the entire population is in the SEP. As we decrease the size of the SEP, we see that lower decay rates, (i.e., drifting lineages that persist for longer) are optimal. This is because the larger ASP results in more exploitation and less exploration. Apparently, the trade-off exists because when there are fewer SEP agents they need more time to drift in order to make discoveries, but larger ASPs are better at exploiting beneficial fitness gradients.

5.4.2.1 Analyzing the saw-tooth data

We ran four saw-tooth landscapes, representing different levels of deception. Across all panels in Fig. 5.4 we see results that correlate with maximized exploration on the bottom left and maximized exploitation on the bottom right. As in the NK-landscape data, the top row (SEP = 0) shows the non-augmented system (the control). Fig. 5.4[a].[1,2] show the results of the GM and PM methods on a saw-tooth function with only limited deception. Here we see that small SEP sizes generate the best performance, while larger SEP size degrades performance. This indicates that the ASP selection settings are well tuned for this function, and evolution does not improve with the addition of super explorers. Moreover, elitism results in the lowest levels of performance, since valley crossing is required to optimize this function. In Fig. 5.4[b,c,d].[1,2] we find the results from the other three saw-tooth landscapes, each with an increasing level of deception. As the functions become harder to adapt to, we see that the control does not produce the best performance because it does not allow enough exploration. As a consequence, larger SEP sizes become more effective. In Fig. 5.4[d] (the hardest function), the best performance occurs when the entire population is SEP agents (SEP size = 1024) and SEP lineages persist for about 200 generations (decay rate = 0.005). This suggests that navigation on this landscape requires a high degree of exploration. The fact that the highest performance values are not at the lowest decay rates means that too much exploration is wasteful. It is worth noting that in Fig. 5.4[c,d][1,2], SEP size = 1024 with decay rate = 0.05 performs about as well as the control. We theorize that the explore-

exploit trade-off in these conditions is similar. We believe this hypothesis can be extended to state that any conditions that share performance values likely also experience similar explore-exploit trade-offs.

The results generated by the saw-tooth landscape using the PS method, Fig. 5.4[a,b,c,d][3] differ significantly from the NK-landscape PS method results (we address this difference in the next section). As expected, the values along the top and right side of each panel are the same and correlate with the control condition. Except in the case of the least deceptive function ([a]) the best performance across the three replacement methods is found in the PS data. This supports the idea that weakened selection is beneficial for valley crossing on the harder saw-tooth functions.

5.4.3 The effect of diminishing returns

The NK-landscape suffers from diminishing returns, which means that the ratio of beneficial to deleterious mutations and the fitness gain per beneficial mutation decreases as fitness increases. As a result, the optimal explore-exploit ratio changes over the course of evolution. Strong exploitation initially maximizes the rate of adaptation, but later is likely to result in getting stuck on local optima. Conversely, stronger exploration, though slower at first, can achieve higher final performance in the long run, because it is able to escape local optima.

In order to disambiguate the effects of diminishing returns from other dynamics, we included the saw-tooth landscapes, which do not exhibit diminishing returns and allow us to study the behavior of a system with a constant difficulty.

5.4.4 Super explorers alone

The bottom row in each panel in Figures 5.3 and 5.4 corresponds to a configuration where the entire population is SEP agents; there are no agents in the ASP. The high performance observed in some SEP-only configurations demonstrates that the combination of drift and replacement is a viable search process in its own right. In fact, for several of the fitness functions, a configuration with $SEP = \text{population size}$ achieves maximum or nearly maximum

performance. We call this new drift-and-replace search process “Drift Pool Optimization” or DPO.

There is a surprising similarity between DPO and particle swarm optimization (PSO). Both DPO and PSO operate as a swarm of particles, as opposed to implementing evolution by natural selection. The main difference is that while the particles in PSO tend to converge towards known optima, the particles in DPO originate near current optima and are allowed to drift. In both cases, the degree to which solutions can diverge from known optima is critical to success: too little divergence will stifle innovation, while too much diffusion will result in chaotic behavior incapable of effectively making discoveries.

5.4.5 Super explorers and recombination

In addition to the experiments shown in this work, we tested the saw-tooth functions with recombination (results not shown). When producing offspring in the ASP and when picking replacements using the PS method, we select two parents and perform three-point recombination to generate offspring (while GM and PM method replacements are still asexual). We found that while scores with recombination tended to be higher, the trends in the data were almost identical.

5.4.6 Extending the super explorer method

In this work, we presented the super-explorer method, an augmentation that can be used to enhance other selection methods. Here we present some potential modifications and extensions. While a super-explorer augmented system is already able to simultaneously support enhanced exploration (in the SEP) and exploitation (in the ASP), the algorithm could be upgraded so that it automatically modifies both decay rate and SEP size based on observed rates of evolution in an adaptive manner, for example by monitoring the time the population spends in stasis.

Another possible modification to the algorithm would be to encourage diversity in the SEP, which would ensure that the SEP explores genetic space more uniformly. Both genetic

diversity and phenotypic diversity could be investigated. In addition, a method like MAP-Elites [Mouret and Clune, 2015] could replace the global max pool and offer a more diverse alternative to elitism.

An interesting alternative to super-explorer augmentation would be to simply use another selection method in place of the SEP. This would allow two selection methods with different behaviors to synergize. We are particularly interested in investigating augmenting Particle Swarm Optimization (PSO) with super explorers, and also in using PSO in place of the SEP to augment other systems.

There is no reason that a system must be limited to only two selection methods. A population could be subdivided into any number of pools, each acting as an ASP or SEP with unique selection methods and parameterization. Automation in the form of pool-size balancing and parameter adjustments could also be considered.

5.4.7 Using the super explorer method to study evolution

The super explorer method could be used to further our understanding of general evolutionary processes related to the trade-off between exploration and exploitation. Since the behavior of the SEP can be related directly to the explore-exploit trade-off, super explorers could be used to gauge the relative explore-exploit trade-off of various selection methods. In this work, we only considered a single selection method for the ASP: roulette selection. Other configurations of roulette-wheel selection, representing different selection strengths, as well as other selection methods should be tested. This approach could allow us to ask a number of different questions, such as: how does tournament size affect the explore-exploit trade-off?

5.5 Conclusion

In this work, we introduced a new bio-inspired optimization technique called “the super explorer method” and we demonstrated the efficacy of the method on a number of deceptive fitness landscapes. There is a wealth of literature exploring how population size, selection strength, and mutation rate affect rates of adaptation. In addition to these, we argue that

the frequency at which selection is applied and the distance a lineage can drift before it is evaluated also affect the success of an evolutionary algorithm. Other evolutionary algorithms have experimented with changing the frequency with which selection is applied (e.g., lexicase and real-valued tournament) or ignoring fitness altogether (e.g., novelty search), but we believe we are the first to design a method that intentionally protects lineages in order to promote discovery by drift.

Chapter 6

Conclusions

In this final chapter, I will summarize the main results of this dissertation and discuss the significance of each result for computer scientists and biologists. I also provide advice on the design and tuning of evolutionary algorithms based on the findings of my work. I identify topics that should be the subject of future research, and I discuss the potential I see in pursuing them. At the end of the chapter, I give my final remarks on the importance of the work, and highlight the most important lessons learned.

6.1 Measuring the impact of selection

In Chapter 2, I introduced a novel measure of selection strength — the selection impact — which I define as a population’s ‘distance from neutral drift’. The new metric examines a population’s offspring production and compares it with the offspring production of a neutrally drifting control population. The selection impact measures the influence that selection has on the reproductive success of organisms in a population. The new metric is dimensionless and makes no assumptions about the type or mode of selection, allowing for easy comparisons between any two measurements, even when comparing between species or modes of selection. The selection impact is strongest when both the strength of selection is strong and the total non-neutral phenotypic variation within the population is high. In one test case, I used

the selection impact to optimize the mutation rate of an evolutionary algorithm that was optimizing simple fitness functions; I found the peak impact occurred when the mutation rate placed the population just on the boundary between the sequential-fixation regime and the successional-mutations regime [Desai and Fisher, 2007].

The evolutionary algorithm engineers and population geneticists each have a set way of thinking about selection that reflects their goals and typical courses of action. Especially for algorithm engineers, selection is seen as a component, a part of an algorithm that can be tuned, swapped out, and re-designed, or in a word subordinate to the engineer. This way of thinking omits the fact that selection is about more than just the selection algorithm; selection is influenced by many factors that I have enumerated earlier during my review of shifting balance theory and related phenomena. The selection impact metric is based on a different way of thinking. The selection impact metric views neutral drift as the default condition of a population of replicators. Selection, therefore, is anything and everything that moves the dynamics of the population away from neutral drift. This is, as far as I am aware, a novel perspective of selection that opens the door for further analysis. I believe the discussion surrounding a new metric of selection strength, based on this way of thinking, serves as the starting point for a discussion about selection in general.

The selection impact metric relies on the experimentalist to first derive the neutral drift reference distribution for their population of interest. I derived several of these reference distributions for Wright-Fisher and Moran generational models. There are similarities between the Wright-Fisher and Moran processes in terms of the way I calculate their drift reference distributions. For the Moran process, it is necessary to track ancestors and descendants, while the Wright-Fisher process tracks parents and offspring; one is clearly a special case of the other. Unfortunately, my method of deriving the Moran drift references is computationally intensive. Ideally, I would have been able to simplify the calculation of the Moran drift reference. It may be possible to simplify the Moran calculations by more closely examining the similarities between the two generational models, though I expect to find that the Wright-

Fisher generational model is simpler because it is the special case.

I chose to define the selection impact as a distance from neutral drift, but I could just as well have been justified to create a metric that was a distance from the strongest possible selection. Preliminary work on this approach showed that the pair of metrics, one a distance from drift and the other a distance from elitism, are not replacements for each other: the relationship between the two distance metrics is not simply that one is x and the other is $1 - x$; there is a non-linear relationship between them. For this reason, I chose not to focus on the elitism distance in order to refine my understanding of the drift distance. However, exploring both metrics simultaneously may end up being a fruitful path to take. When the two metrics are plotted against each other, the fixation of a new beneficial mutation traces out a loop: the two metrics are not even in phase with one another. This is not true of all pairs of selection strength metrics, and investigating the metrics that do form loops along with the structure of the loops themselves may provide insights into the fixation dynamics of a population.

Other measures of selection strength, like the ubiquitous Price equation and its generalizations, relate selection strength to the change of macroscopic traits in the population. During my work on the selection impact, I began looking for ways to incorporate macroscopic traits into the metric. The foundation of the selection impact is that neutral drift is the null hypothesis, and this must remain the same while other changes are considered. One solution I considered is adopting the Fokker-Planck equation (FPE) [Risken, 1989] as the underlying computational framework for the selection impact. Instead of deriving a drift reference from first principles, we can instead use an observation of a population together with the FPE to predict the future state of the population, while assuming only neutral drift. In this technique, the population is first observed and a distribution of trait measurements is compiled. Then, using the FPE to simulate neutral drift on the trait frequencies, the distribution of traits is advanced forward in time by Δt . Finally, a second measurement of the population is taken after Δt has elapsed in real time. The final measurement and the predicted outcome

can then be compared using the same method of comparing distributions as the original selection impact. The difference between the prediction, which assumes neutral drift, and the observation is therefore the selection impact on the trait in question.

6.2 Altering selection strength with fitness noise

In Chapter 3, I compared three methods of controlling selection strength: adding noise to each organism's fitness score, changing the population size, and changing the tournament size of the selection algorithm. By analyzing the rate of adaptation on a special fitness function, we found that adding noise to fitness changed the rate of adaptation in a manner similar to changing the tournament size. We also discovered that each method for controlling the selection strength affected some aspect of the information transfer from the environment to reproductive success: fitness noise hinders the detection of *truly* fit organisms, population size can increase or reduce sampling error, and tournament size increases or decreases the strictness of the selection filter. However, increasing the population size has an additional effect: increasing the inflow of mutations. Consequently, I now believe population size is an inadequate parameter for controlling the selection strength because it affects more than one property of the system. Despite their differences, the similarities of each method prompted us to consider compensatory changes that involve two or more methods. For example, if increasing the population size is desirable to increase mutation inflow, but selection needs to be kept low to improve valley-crossing, then adding noise alongside the increase in population size could compensate for and counteract the increase in selection strength.

In Chapter 1 I asked, does noisy fitness (among other things) merely *look like* changing the strength of selection or is it *actually* changing the strength of selection, and is there a meaningful distinction to be made? I have come to believe that there is no real distinction between appearing to change the selection strength, and actually doing so. Each phenomenon that appears to affect the strength of selection plays some role in the transference of information from the environment to the population. This has led me to wonder if there isn't

a more central theory that could tie all of these selection-altering phenomena together. I was inspired by the work of [Rivoire and Leibler, 2011, Montanez, 2013, Rivoire and Leibler, 2014] and others who represent evolution as a transfer of information from the environment (or fitness function) to an organism’s genotype and eventually the entire population. This seems to me to be the right way forward for understanding selection strength in a broader, unified way. With a sufficiently detailed model of the information transfer, each phenomenon that alters the strength of selection will have a place in that model, and clearly show the loss of information that results in the weakening of selection. Paradoxically, we know that perfect selection hinders adaptation, so we can infer that a lossless transfer of information is suboptimal for evolution. However, this raises the question: what is the optimal type of information loss and at what stage along the chain of transmission is the information loss best placed?

During my work on noisy fitness, I devised a simple way of defining a tournament size between two integer values. I called this technique real-valued tournaments (RVT). RVT can be defined as follows: if $T \in \mathbb{R}$ and $T \geq 1$ then T is a real-valued tournament size. A real-valued tournament size can be used in the tournament selection algorithm as follows: conduct an integer-valued tournament of size¹ $\lfloor T \rfloor$ with probability $1 - (T - \lfloor T \rfloor)$ or a tournament size of $\lfloor T \rfloor + 1$ with probability $T - \lfloor T \rfloor$. The average tournament size will be T . RVT has utility beyond being a selection algorithm. For example, we can calculate the *effective tournament size* of a mutant-type group undergoing a selective sweep to reveal the weaker selection due to the free-for-all effect. First, let x denote the density of the mutant type in a well-mixed population. If we assume the mutant type always wins in a tournament against a wild type (even when a few minor deleterious mutations are hitchhiking), all wild/mutant tournaments are effectively size $T = 1$ and all mutant/mutant tournaments are effectively size $T = 2$. Thus, if we take a weighted average of these two types of encounters, we will

¹The floor function, $\lfloor \cdot \rfloor$, is a function that returns the nearest whole number, less-than or equal-to the input.

have the effective tournament size of the mutant-type group.

$$T_{\text{eff}} = \frac{2x(1-x) + 2x^2}{2x(1-x) + x^2} = \frac{2}{2-x} \quad (6.1)$$

This is perhaps one of the easiest ways to understand the free-for-all effect: by comparing the probability of entering a tournament you are assured to win, versus a tournament that actually tests the organisms involved. Working with tournament selection in this way provides a simple and intuitive way to think about selection. This methodology merits further development.

6.3 The evolutionary free-for-all effect

In chapter 4 I introduced a new evolutionary dynamic, the free-for-all effect (FFA). FFA is a reduction of selection pressure on clades that have a competitive advantage. The reduction in selection strength is proportional to the fitness advantage, and can lead to a self-perpetuating period of accelerated evolution called a FFA cascade. FFA naturally shifts the balance of selection during selective sweeps, and so fits nicely into a general shifting balance theory, like I have discussed in Chapter 1. For these reasons, FFA is also high on the list of phenomena that can help explain the pattern of punctuated equilibrium.

Free-for-all does not guarantee additional valley-crossings will occur, it merely reduces the average stochastic tunneling time from what it would be during a period of equilibrium. Therefore, when fixation times are longer, there is more time to allow stochastic tunneling to occur while FFA is having an effect. This tells us that longer fixation times can accelerate evolution, merely by making time for mutations to accumulate in the population while conditions are more favorable to exploration. Longer fixation times also introduce a greater probability of clonal interference in asexual populations, so there is a trade-off to consider for researchers looking to leverage FFA in their evolutionary algorithms. Clonal interference can be mitigated by using sexual recombination, which is already fairly standard in evolutionary computing, however sex creates stronger selection [Kondrashov, 1988, Rice and Chippindale,




Figure 6.1: A 1-dimensional ring population structure with N organisms. Organisms with a beneficial mutation are colored red. In each generation, the number of organisms with the beneficial mutation can increase by at most 2. The fixation of a beneficial mutation takes at least $\frac{N}{2}$ generations.

2001], so this too may need to be compensated for.

While investigating free-for-all in different population structures, we observed an important difference between well-mixed and spatial populations: spatial populations have a maximum growth rate. For example, consider a 1-dimensional ring population structure where organisms can only reproduce into adjacent locations, as shown in Figure 6.1. Because offspring must be placed near their parent, the growth rate of an advantaged genotype is limited, regardless of its fitness advantage over other genotypes. This also places a limit on the minimum fixation time of a beneficial mutation, keeping the free-for-all window open at least that long.

In general, a d -dimensional spatial population where offspring are placed adjacent to their parent has a minimum fixation time of

$$t_{\min} \propto \frac{\sqrt{d}}{2} N^{1/d}. \quad (6.2)$$

In contrast, an advantaged genotype in a well-mixed population can grow arbitrarily fast, limited only by the fitness advantage itself. In theory, a sufficiently beneficial mutation could sweep a well-mixed population in a single generation.

Since FFA has a larger effect when fixation is slow and when the fitness benefit is large, well-mixed populations can prohibit their own free-for-all effect by rushing high-fitness genotypes to fixation. In contrast, spatial populations offer a way for high-fitness genotypes to experience a long FFA window. This is particularly important when FFA cascades occur, since each new discovery in the cascade increases the fitness advantage of the mutant type. In a spatial population, the additional fitness only serves to enhance FFA. In a well-mixed population, these additional benefits would accelerate the fixation event and close the FFA window faster. Below, I discuss the possibility of studying other population structures and their relationship to FFA.

The importance of the free-for-all effect when discussing the tempo and mode of evolution cannot be understated. FFA is a shift of the strength of selection that occurs during *every* selective sweep (range expansions included as a special case). As such, FFA must be kept in mind for nearly all of evolutionary science. FFA creates the opportunity for extended periods of rapid adaptation that might not occur otherwise, and is a crucial addition to the many other factors that can explain punctuated equilibrium.

FFA also brings into focus the importance of deep clades for the exploration of the fitness landscape. Valley-crossing requires at least one deleterious mutation to occur on the same lineage before a benefit can be discovered. Since during a selective sweep only those lineages with higher-than-average fitness can accumulate deleterious mutations and continue to grow, the deleterious mutations are concentrated in these above-average lineages. These lineages can travel a further distance in the fitness landscape, cross wider fitness valleys, and explore hard to reach areas. In contrast, deleterious mutations accumulated during periods of equilibrium are easily removed by selection because the lineages they accumulate on are likely brought below the average fitness. Thus, it is important to 'protect' some lineages from the fitness penalty of their deleterious mutations in order to increase the exploratory power of the population. Spatial population structures naturally protect the lineages at the boundary of a selective sweep, while also maintaining the strength of purifying selection

within the boundary. For this reason, and the reasons given above, I strongly encourage the use of spatial population structures in all EAs where well-mixed populations are not strictly necessary.

As of today, free-for-all is a theoretical discovery, by which I mean we have not yet found conclusive evidence for the effect in observations of natural populations. Finding such evidence is difficult for a number of reasons. Primarily, there are many confounding factors which may either look like FFA, or mask FFA with a counteracting effect. For example, in our work with well-mixed populations, we observed a kind of 'refractory period' that takes place after a selective sweep. This refractory period is essentially the time during which a population recovers from the loss of diversity that accompanies a selective sweep. This loss of diversity makes valley-crossing harder because the population cannot explore the fitness landscape effectively. Thus, the loss of diversity appears to lengthen the average time between valley-crossing events because additional time is needed to rebuild the molecular quasispecies [Eigen and Schuster, 1979, Wilke, 2005, Biebricher and Eigen, 2006]. FFA decreases the average time between crossings, and so the two effects together can cancel each other out when the impact of FFA on the dynamics of adaptation is weak, like in well-mixed populations. This 'canceling out' only happens when examining the average rate of adaptation, but both the refractory period and FFA can be clearly seen in the distribution of valley-crossing times: when compared to the equilibrium crossing-time distribution, the refractory period creates as a truncated left tail and FFA creates bi-modality. When designing our experiments, we made sure to exclude as many factors that could obfuscate the detection of FFA as possible, including epistasis, multiple traits, sex, diminishing returns, negative frequency dependence, and ecology. Nature systems, on the contrary, has each of these all at once, so conclusively demonstrating FFA in a natural system would require either a model organism with none of these confounding factors, or a method for discounting the effect of each of these factors from the data. Neither of these paths seem particularly easy, and thus warrant an entire investigation of their own.

To understand the real impact of FFA, it is important to investigate how FFA affects populations of organisms with multiple traits. In all the investigations of FFA that I have shown, we have used a one-dimensional fitness function. This means that the FFA cascades we have observed so far have all been sequential: each fitness valley that is crossed is on the same fitness dimension and so must be crossed one at a time. In a multi-trait system, the discovery of a beneficial mutation on one trait relaxes the selection pressure on the entire organism, so several traits may valley-cross in parallel with one another. These parallel FFA cascades have two important consequences: larger fitness improvements and the circumvention of trade-offs. Parallel cascades are more likely to result in very large fitness improvements, as the parallel discoveries combine their fitness advantages. We have already shown that FFA reduces selection strength proportional to the net fitness advantage, so parallel cascades and the large advantages they bring could dramatically reduce selection strength beyond anything we have seen so far. Additionally, multi-trait systems often have trade-off relationships between two or more traits, where one trait cannot improve without another trait losing fitness (see [Deb et al., 2002] for more information about evolutionary optimization with trade-offs). Parallel FFA cascades create opportunities for traits linked by trade-offs to all improve at the same time, since the loss of fitness due to a trade-off may not be lethal during the period of reduced selection strength.

In preliminary work, we found that, indeed, each trait in an organism benefits from the reduced selection during a selective sweep. Therefore, FFA induced by the improvement of one trait may indirectly help other ostensibly unrelated traits cross a fitness valley. The increased probability to valley-cross on every trait drastically increases the likelihood of a FFA cascade. Furthermore, we found that systems with sufficiently many traits experience extremely long cascades lasting up to 80,000 generations and crossing up to 160 fitness valleys with only 6 traits. These very long cascades lend further support to the idea that FFA plays a significant role in the explanation of the observed, often dramatic, changes in the frequency of innovations, known as punctuated equilibrium. The interplay of multiple traits and FFA

is the subject of ongoing research. Preliminary findings about multi-trait systems give rise to a very simple evolutionary computation augmentation: simply adding 'dummy traits' to the fitness function that increase the dimensionality of the fitness landscape, making it possible for FFA to be triggered by a dummy trait and assist the main optimization task. Investigating ways to leverage FFA to speed up adaptation led to the development of super explorers (Chapter 5), and still provides opportunities for additional developments.

When searching for other factors that may affect a population's ability to escape its current local optimum, we found that the *shape* of the optimum itself is such a factor. For example, in my work on the selection impact metric the more steeply-sloped fitness landscapes showed a higher selection impact (all other factors being the same), indicating selection doing more work. This has led to the hypothesis that populations tend to get stuck on local optima with gentler slopes and rounder surfaces for longer periods of time than on steeper, jagged peaks, because the latter create a more turbulent equilibrium and thus create a kind of standing FFA effect. To help illustrate this point, imagine a selective sweep occurring that never quite ends, but instead perpetually climbs and slips down the top of the local optimum, like Sisyphus's boulder in the eponymous Greek myth. These small reoccurring Sisyphean selective sweeps are essentially what occurs when the population is equilibrated on a local optimum and selection is acting against the inflow of deleterious mutations. Thus, we predict that stabilizing selection results in a small FFA effect that, over very long time spans, should see steeper local optima escaped faster than gentler ones. This hypothesis has not yet been tested.

We have investigated FFA in well-mixed populations, as well as one- and two-dimensional population structures. FFA occurs in each of them, but the size of the effect and the manner in which it manifests is slightly different for the reasons discussed above. Clearly, a general study of how population structure affects FFA is needed. I am particularly inspired by the work of [Mühlenbein, 1991, Lieberman et al., 2005, Mühlenbein, 2009, Askari and Samani, 2015, Möller et al., 2019, Tkadlec et al., 2019, Kuo et al., 2021] who have extensively studied

how population structure can affect fixation rates and probabilities, among other things. We have already observed that population structures that promote Fisher waves [Fisher, 1937], slow down fixation rates, and decouple growth rate from fitness advantage to promote a stronger free-for-all effect. We have also already observed that the number of individuals that experience FFA depends on the population structure. For example, in the 1D structure in Figure 6.1, only the two individuals at the boundary between the mutant and wild-type subpopulations experience FFA, while in the analogous 2D population the size of the boundary changes over the course of fixation. There is substantial potential to improve evolutionary computation with further research into how FFA is affected by population structure.

6.4 Super Explorers: protected, sustained, neutral drift

In Chapter 5, I introduced a new method called “super explorers” as a way to augment evolutionary algorithms. The method is based on previous research that showed that the free-for-all effect in spatial populations results in reduced selection pressure and deeper clades. Super explorers are designed to reproduce regardless of their fitness, resulting in deep clades that can explore fitness valleys without penalty. Occasionally, a super explorer will be replaced by an organism from the rest of the evolving population, putting a soft bound on the depth of the clades they produce. I showed that the addition of super explorers to an evolutionary algorithm increases the exploration ability of the algorithm, even for modest super explorer counts. Super explorers provide a benefit comparable to lowering the selection strength of the entire population, without the risk of losing good solutions to Muller’s ratchet and neutral drift. The resulting increase in adaptive power is analogous to the effects of FFA, as intended.

The development of super explorers led to a novel method of joining two or more populations together that I am now calling *horizontal concatenation*, depicted in Figure 6.2. This method allows for the combination of an arbitrary number of population-based search methods and encourages rapid exchange of information between populations, unlike an island model that emphasizes isolation.




Figure 6.2: A general schematic of the horizontal concatenation method. An arbitrary number of population-based search algorithms and archives (referred to collectively as pools) are joined together in parallel. Each pool has read-only access to the other pools and can make use of them to update itself, depicted here as arrows from t to $t + 1$. There is no requirement that a pool read from other pools, and the technical limitations of some pools may prevent them from connecting with other pools. See Chapter 5 for more information.

The horizontal concatenation method was inspired by the architecture of Markov brains (Fig. 3, [Hintze et al., 2017]). Each pool manages a container of organisms. For example, a typical evolutionary algorithm manages the evolving population by choosing parents and creating offspring, while an archive manages a container of organisms by deciding which organisms to include, and which to exclude. In this way, each pool can be thought of as an independent, self-contained unit with its own set of organisms that it alone is responsible for. Each pool must decide for itself what constitutes a single update of its container, from time t to time $t + 1$. For example, an evolving population can have overlapping generations or non-overlapping generations, and the definition of a single update depends on the choice of generational model. With this description of a pool in mind, the horizontal concatenation method can be defined simply as allowing each pool to read from the containers of other pools for the purposes of informing their own update process. It is in this way that the horizontal concatenation method resembles the structure of a Markov brain: the set of all containers is updated in parallel from t to $t + 1$, and each pool can read from any other container, but can only write to its own container. Since an arbitrary number of pools can be linked together (concatenated) and all pools update in parallel (so are arranged horizontally), I call this method of parallel search “horizontal concatenation”.

Horizontal concatenation has several advantages: it allows us to track which pool an

organism and its offspring are in without ambiguity, and it allows us to add more pools seamlessly. The ability to unambiguously track which pool an organism is in mostly serves to make analysis of the line of descent easier. The ability to add more pools seamlessly, without the need to modify the internal logic of other pools, makes sharing information between pools a far easier and more powerful than it seems at first. To illustrate this point, first consider the setup we used in the original work: we had a normally evolving population, a pool of super explorers, and an additional pool that contained just the organism with the highest fitness ever discovered, the global max. In that work, we sometimes replace a decayed super explorer with the global max. However, we could instead replace the global max pool with a more sophisticated archive of solutions, like MAP-Elites [Mouret and Clune, 2015], providing a diverse range of choices to replace a decayed super explorer. But we can do even better than simply replacing one pool with another; we can add more pools with drastically different search properties. For example, we can concatenate an evolving population and an implementation of particle-swarm optimization [Poli et al., 2007]. Preliminary work on this particular combination has shown that interfacing the two search algorithms is rather easy; the inclusion of a global-max pool helps link the two searches together. Furthermore, two evolutionary searches with different selection algorithms or different generational models could be combined to leverage the differences between them. The possibilities are too numerous to list. Future work should investigate the limitations of this method, which are as yet unknown. One obvious limitation is the increased computational overhead of adding more pools, which is already a familiar limitation of evolutionary algorithms. It is unclear, however, if the increased computational effort will be worth the search dynamics of a complex hybrid algorithm. Horizontal concatenation shares some similarities to other parallel evolutionary algorithms, like island models and demes, and this warrants comparative analysis to determine how it fits into that area of research.

6.5 Final Remarks

I believe that it is a mistake to regard evolutionary computation as separate from evolutionary science. Even though evolutionary algorithms are inspired by evolution, the relationship between the two does not end there. In a very real sense, natural evolution and evolutionary algorithms are both instantiations of the same fundamental process. The differences are merely details. This is why I believe it is paramount that evolutionary science and evolutionary computation work together more collaboratively.

Across all of these projects, I consistently found that well-mixed population models were hard to tune for good performance compared to spatial models, only weakly exhibited free-for-all and so made little use of its benefits, and create spooky competition-at-a-distance interactions to boot. For these reasons, I have come to believe that spatial population models, even simple one-dimensional ring populations, alleviate nearly all the problems often associated with well-mixed populations. As such, I recommend the use of one-dimensional ring populations or other simple spatial population models as the go-to population model for evolutionary algorithms. Favoring spatial models is not actually a new revelation, but something that has been discussed before [Mühlenbein, 1991, Mühlenbein, 2009]. Spatial populations, in my opinion, deserve far more attention in evolutionary optimization than they currently get.

We find that in nature, selection strength is a highly fluid property of the evolving system. The idea of naturally shifting selection strength is an old idea, but it has been re-discovered many times to be the consequence of a diverse set of mechanisms. I believe we should start to view selection as more of an emergent property of the entire evolutionary system, with some aspects of the system contributing more heavily to selection and others lightly.

I have investigated selection, and shifting selection strength, in order to find new ways of improving evolutionary algorithms. I found that selection strength can be thought of as the degree to which an evolving system is dissimilar to neutral drift. This perspective opens the door to accept all phenomena that affect the strength of selection as part of a unified theory

of selection. Such a theory is likely built on the idea that evolution is the transmission of information from the fitness landscape to the population.

Spatial structure in the population is an important re-occurring theme in much of this thesis, and offers a minimally invasive way to improve an evolutionary algorithm. Spatial structures allow for local selection strengths to weaken while keeping the global average strength high, allowing for enhanced exploration at little cost to exploitation.

I have presented a new evolutionary dynamic: the free-for-all effect, namely the reduction of selection strength on organisms with higher-than-average fitness during a selective sweep or range expansion. Shifting selection strength seems to be the best contender to explain the punctuated nature of evolution, and FFA appears to be one of the most likely ways of causing this shift to occur.

The discovery of FFA led to the design of a new evolutionary search method called super explorers. Super explorers serve as the first example of a new technique in evolutionary computing, called horizontal concatenation, which has the potential to combine drastically different search algorithms in a hybrid parallel search.

Natural evolution inspires evolutionary algorithms. This initial inspiration has led to an entire field of evolutionary optimization techniques. However, like all optimization techniques, evolution is not a silver bullet, and must be tuned for each specific situation. Natural evolution, while theoretically bound by the same constraints, seems to adapt continuously, sometimes slowly and sometimes rapidly. In order to design the next generation of evolutionary optimization algorithms, we should continue to look back to nature for inspiration.

Attributions

Chapter 3

This chapter is based on work done in collaboration with Clifford Bohm [Ragusa and Bohm, 2021]. Both authors contributed equally to the final text of the original manuscript, data collection and analysis, and experimental design. Ragusa provided the final figures and pseudo-code. Bohm maintains the github repository.

Chapter 4

This chapter is based on unpublished work done in collaboration with Clifford Bohm, Christoph Adami, Charles Ofria, and Richard E. Lenski. Ragusa and Bohm contributed equally to the original text, data collection and analysis, and experimental design. Bohm provided the final figures. Adami made contributions to the original text.

Chapter 5

This chapter is based on work done in collaboration with Clifford Bohm [Ragusa and Bohm, 2022]. Both authors contributed equally to the final text of the original manuscript, data collection and analysis, and experimental design. Bohm provided the final figures. Ragusa provided the pseudo-code, and initial design of super explorers and the horizontal concatenation method.

BIBLIOGRAPHY

- [Adami, 1998] Adami, C. (1998). *Introduction to Artificial Life*. Springer Science & Business Media.
- [Adami, 2006] Adami, C. (2006). Digital genetics: unravelling the genetic basis of evolution. *Nature Reviews Genetics*, 7(22):109–118.
- [Alili et al., 2005] Alili, L., Patie, P., and Pedersen, J. L. (2005). Representations of the first hitting time density of an Ornstein-Uhlenbeck process. *Stochastic Models*, 21(4):967–980.
- [Anderson, 1962] Anderson, T. W. (1962). On the distribution of the two-sample Cramer-von Mises criterion. *The Annals of Mathematical Statistics*, 33(3):1148–1159.
- [Arttime et al., 2018] Arttime, O., Khalil, N., Toral, R., and San Miguel, M. (2018). First-passage distributions for the one-dimensional fokker-planck equation. *Physical Review E*, 98(4):042143.
- [Askari and Samani, 2015] Askari, M. and Samani, K. A. (2015). Analytical calculation of average fixation time in evolutionary graphs. *Physical Review E*, 92(4):042707.
- [Bak and Boettcher, 1997] Bak, P. and Boettcher, S. (1997). Self-organized criticality and punctuated equilibria. *Physica D: Nonlinear Phenomena*, 107(2):143–150.
- [Bak et al., 1987] Bak, P., Tang, C., and Wiesenfeld, K. (1987). Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters*, 59(4):381–384.
- [Baym et al., 2016] Baym, M., Lieberman, T. D., Kelsic, E. D., Chait, R., Gross, R., Yelin, I., and Kishony, R. (2016). Spatiotemporal microbial evolution on antibiotic landscapes. *Science*, 353(6304):1147–1151.
- [Biebricher and Eigen, 2006] Biebricher, C. K. and Eigen, M. (2006). What is a quasispecies? *Quasispecies: concept and implications for virology*, pages 1–31.
- [Bininda-Emonds et al., 2007] Bininda-Emonds, O. R. P., Cardillo, M., Jones, K. E., MacPhee, R. D. E., Beck, R. M. D., Grenyer, R., Price, S. A., Vos, R. A., Gittleman, J. L., and Purvis, A. (2007). The delayed rise of present-day mammals. *Nature*, 446:507–512.
- [Bohm et al., 2019] Bohm, C., Ackles, A. L., Ofria, C., and Hintze, A. (2019). On sexual selection in the presence of multiple costly displays. In *ALIFE 2019: The 2019 Conference on Artificial Life*, pages 247–254. MIT Press.
- [Bohm et al., 2017] Bohm, C., C G, N., and Hintze, A. (2017). MABE (modular agent based evolver): A framework for digital evolution research. In *Artificial Life Conference*

Proceedings, pages 76–83. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info

- [Brunet et al., 2008] Brunet, E., Rouzine, I. M., and Wilke, C. O. (2008). The stochastic edge in adaptive evolution. *Genetics*, 179:603–20.
- [Burton and Travis, 2008] Burton, O. J. and Travis, J. M. (2008). The frequency of fitness peak shifts is increased at expanding range margins due to mutation surfing. *Genetics*, 179(2):941–950.
- [Bush et al., 1977] Bush, G. L., Case, S. M., Wilson, A. C., and Patton, J. L. (1977). Rapid speciation and chromosomal evolution in mammals. *Proceedings of the National Academy of Sciences*, 74(9):3942–3946.
- [Bäck, 1994] Bäck, T. (1994). Parallel optimization of evolutionary algorithms. In Davidor, Y., Schwefel, H.-P., and Männer, R., editors, *Parallel Problem Solving from Nature — PPSN III*, Lecture Notes in Computer Science, page 418–427, Berlin, Heidelberg. Springer.
- [Campos et al., 2004] Campos, P. R., Adami, C., and Wilke, C. O. (2004). Modelling stochastic clonal interference. *Modelling in Molecular Biology*, pages 21–38.
- [Campos and de Oliveira, 2004] Campos, P. R. A. and de Oliveira, V. M. (2004). Mutational effects on the clonal interference phenomenon. *Evolution*, 58:932–7.
- [Cariani, 2002] Cariani, P. A. (2002). Extradimensional bypass. *Biosystems*, 64(1-3):47–53.
- [Carson, 1975] Carson, H. L. (1975). The genetics of speciation at the diploid level. *The American Naturalist*, 109(965):83–92.
- [Cele et al., 2022] Cele, S., Karim, F., Lustig, G., San, J. E., Hermanus, T., Tegally, H., Snyman, J., Moyo-Gwete, T., Wilkinson, E., Bernstein, M., Khan, K., Hwa, S.-H., Tilles, S. W., Singh, L., Giandhari, J., Mthabela, N., Mazibuko, M., Ganga, Y., Gosnell, B. I., Karim, S. S. A., Hanekom, W., Van Voorhis, W. C., Ndung’u, T., COMMIT-KZN Team, Lessells, R. J., Moore, P. L., Moosa, M.-Y. S., de Oliveira, T., and Sigal, A. (2022). Sars-cov-2 prolonged infection during advanced hiv disease evolves extensive immune escape. *Cell Host Microbe*, 30:154–162.e5.
- [Charlesworth et al., 1982] Charlesworth, B., Lande, R., and Slatkin, M. (1982). A neo-darwinian commentary on macroevolution. *Evolution*, 36(3):474–498.
- [Chow, 2004] Chow, S. S. (2004). Adaptive radiation from resource competition in digital organisms. *Science*, 305(5680):84–86.
- [Conrad, 1990] Conrad, M. (1990). The geometry of evolution. *BioSystems*, 24(1):61–81.

- [Conway Morris, 1998] Conway Morris, S. (1998). *The crucible of creation: the Burgess Shale and the rise of animals*. Oxford University Press.
- [Covert et al., 2013] Covert, 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.
- [Davidson and Erwin, 2006] Davidson, E. H. and Erwin, D. H. (2006). Gene regulatory networks and the evolution of animal body plans. *Science*, 311(5762):796–800.
- [De Jong, 2012] De Jong, K. (2012). *Generalized Evolutionary Algorithms*, page 625–635. Springer, Berlin, Heidelberg.
- [de Visser and Rozen, 2005] de Visser, J. A. G. M. and Rozen, D. E. (2005). Limits to adaptation in asexual populations. *J Evol Biol*, 18:779–88.
- [de Visser et al., 1999] de Visser, J. A. G. M., Zeyl, C. W., Gerrish, P. J., Blanchard, J. L., and Lenski, R. E. (1999). Diminishing returns from mutation supply rate in asexual populations. *Science*, 283:404–6.
- [Deb et al., 2002] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197.
- [Desai and Fisher, 2007] Desai, M. M. and Fisher, D. S. (2007). Beneficial mutation–selection balance and the effect of linkage on positive selection. *Genetics*, 176(3):1759–1798.
- [Desai et al., 2007] Desai, M. M., Fisher, D. S., and Murray, A. W. (2007). The speed of evolution and maintenance of variation in asexual populations. *Current biology*, 17(5):385–394.
- [Eigen, 1971] Eigen, M. (1971). Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften*, 58:465–523.
- [Eigen et al., 1989] Eigen, M., McCaskill, J., and Schuster, P. (1989). The molecular quasi-species. *Adv. Chem. Phys.*, 75:149–263.
- [Eigen and Schuster, 1979] Eigen, M. and Schuster, P. (1979). *The Hypercycle—A Principle of Natural Self-Organization*. Springer-Verlag, Berlin.
- [Eldredge, 1972] Eldredge, N. (1972). Punctuated equilibria: an alternative to phyletic gradualism. *Models in paleobiology*, page 82–115.
- [Elena et al., 1996] Elena, S. F., Cooper, V. S., and Lenski, R. E. (1996). Punctuated

- evolution caused by selection of rare beneficial mutations. *Science*, 272(5269):1802–1804.
- [Engholtdt and Mathias, 2021] Engholtdt, K. and Mathias, H. D. (2021). A biologically-inspired model for mass extinction in genetic algorithms. In *2021 IEEE Congress on Evolutionary Computation (CEC)*, pages 1078–1085. IEEE.
- [Erwin and Valentine, 1984] Erwin, D. H. and Valentine, J. W. (1984). “hopeful monsters,” transposons, and metazoan radiation. *Proceedings of the National Academy of Sciences*, 81(17):5482–5483.
- [Erwin et al., 1987] Erwin, D. H., Valentine, J. W., and Sepkoski, J. J. (1987). A comparative study of diversification events: The early paleozoic versus the mesozoic. *Evolution*, 41(6):1177–1186.
- [Excoffier and Ray, 2008] Excoffier, L. and Ray, N. (2008). Surfing during population expansions promotes genetic revolutions and structuration. *Trends in Ecology & Evolution*, 23(7):347–351.
- [Fay et al., 2002] Fay, J. C., Wyckoff, G. J., and Wu, C.-I. (2002). Testing the neutral theory of molecular evolution with genomic data from drosophila. *Nature*, 415(6875):1024–1026.
- [Feller, 1951] Feller, W. (1951). Diffusion processes in genetics. In *Proceedings of the second Berkeley Symposium on Mathematical Statistics and Probability*, pages 227–246. University of California Press.
- [Fisher, 2013] Fisher, D. S. (2013). Asexual evolution waves: Fluctuations and universality. *J. Stat. Mech.: Theory and Experiment*, 2013:P01011.
- [Fisher, 1937] Fisher, R. A. (1937). The wave of advance of advantageous genes. *Annals of Eugenics*, 7(4):355–369.
- [Fitch, 1995] Fitch, W. M. (1995). *Tempo And Mode In Evolution: Genetics And Paleontology 50 Years After Simpson*. National Academies Press (US).
- [Fogle et al., 2008] Fogle, C. A., Nagle, J. L., and Desai, M. M. (2008). Clonal interference, multiple mutations and adaptation in large asexual populations. *Genetics*, 180(4):2163–2173.
- [Fontdevila, 1992] Fontdevila, A. (1992). Genetic instability and rapid speciation: are they coupled? *Genetica*, 86(1):247–258.
- [Forster et al., 2006] Forster, R., Adami, C., and Wilke, C. O. (2006). Selection for mutational robustness in finite populations. *J Theor Biol*, 243:181–90.
- [Gavrilets, 1999] Gavrilets, S. (1999). Evolution and speciation in a hyperspace: the roles

- of neutrality, selection, mutation and random drift. *Evolution*, page 1.
- [Gavrilets, 2003] Gavrilets, S. (2003). Perspective: models of speciation: what have we learned in 40 years? *Evolution*, 57:2197–215.
- [Gerrish and Lenski, 1998] Gerrish, P. J. and Lenski, R. E. (1998). The fate of competing beneficial mutations in an asexual population. *Genetica*, 102:127–144.
- [Gilbert et al., 2017] Gilbert, K. J., Sharp, N. P., Angert, A. L., Conte, G. L., Draghi, J. A., Guillaume, F., Hargreaves, A. L., Matthey-Doret, R., and Whitlock, M. C. (2017). Local adaptation interacts with expansion load during range expansion: maladaptation reduces expansion load. *The American Naturalist*, 189(4):368–380.
- [Gillespie, 1984a] Gillespie, J. H. (1984a). The molecular clock may be an episodic clock. *Proceedings of the National Academy of Sciences*, 81(24):8009–8013.
- [Gillespie, 1984b] Gillespie, J. H. (1984b). Molecular evolution over the mutational landscape. *Evolution*, 38:1116–1129.
- [Good et al., 2012] Good, B. H., Rouzine, I. M., Balick, D. J., Hallatschek, O., and Desai, M. M. (2012). Distribution of fixed beneficial mutations and the rate of adaptation in asexual populations. *Proc Natl Acad Sci U S A*, 109:4950–5.
- [Gould, 1977] Gould, S. J. (1977). Evolution’s erratic pace. *Natural History*, 86(5):12–16.
- [Gould, 1987] Gould, S. J. (1987). *Is a New and General Theory of Evolution Emerging?*, page 113–130. Life Science Monographs. Springer US.
- [Gould, 1989] Gould, S. J. (1989). *Wonderful Life: the Burgess Shale and the Nature of History*. WW Norton & Company.
- [Gould and Eldredge, 1977] Gould, S. J. and Eldredge, N. (1977). Punctuated equilibria: The tempo and mode of evolution reconsidered. *Paleobiology*, 3(2):115–151.
- [Gould and Eldredge, 1993] Gould, S. j. and Eldredge, N. (1993). Punctuated equilibrium comes of age. *Nature*, 366(6452):223–227.
- [Hahn, 2008] Hahn, M. W. (2008). Toward a selection theory of molecular evolution. *Evolution*, 62(2):255–265.
- [Haldane, 1954] Haldane, J. B. S. (1954). The measurement of natural selection. *Proceedings of the 9th International Congress of Genetics*, 1:480–487.
- [Handel and Rozen, 2009] Handel, A. and Rozen, D. E. (2009). The impact of population size on the evolution of asexual microbes on smooth versus rugged fitness landscapes. *BMC*

Evolutionary Biology, 9(1):236.

- [Helmuth et al., 2014] Helmuth, T., Spector, L., and Matheson, J. (2014). Solving uncompromising problems with lexica selection. *IEEE Transactions on Evolutionary Computation*, 19(5):630–643.
- [Hereford et al., 2004] Hereford, J., Hansen, T. F., and Houle, D. (2004). Comparing strengths of directional selection: How strong is strong? *Evolution*, 58(10):2133–2143.
- [Hintze et al., 2017] Hintze, A., Edlund, J. A., Olson, R. S., Knoester, D. B., Schossau, J., Albantakis, L., Tehrani-Saleh, A., Kvam, P., Sheneman, L., Goldsby, H., Bohm, C., and Adami, C. (2017). Markov brains: A technical introduction. (arXiv:1709.05601). arXiv:1709.05601 [cs, q-bio].
- [Ho and Pepyne, 2002] Ho, Y.-C. and Pepyne, D. L. (2002). Simple explanation of the no-free-lunch theorem and its implications. *Journal of optimization theory and applications*, 115(3):549–570.
- [Hoffmann and Parsons, 1997] Hoffmann, A. A. and Parsons, P. A. (1997). *Extreme Environmental Change and Evolution*. Cambridge University Press.
- [Holland, 1992] Holland, J. H. (1992). Genetic algorithms. *Scientific American*, 267(1):66–73.
- [Houle, 1992] Houle, D. (1992). Comparing evolvability and variability of quantitative traits. *Genetics*, 130(1):195–204.
- [Iwasa et al., 2004a] Iwasa, Y., Michor, F., and Nowak, M. A. (2004a). Stochastic tunnels in evolutionary dynamics. *Genetics*, 166(3):1571–1579.
- [Iwasa et al., 2004b] Iwasa, Y., Michor, F., and Nowak, M. A. (2004b). Stochastic tunnels in evolutionary dynamics. *Genetics*, 166(3):1571–1579.
- [Jain et al., 2011] Jain, K., Krug, J., and Park, S.-C. (2011). Evolutionary advantage of small populations on complex fitness landscapes. *Evolution*, 65(7):1945–1955.
- [John Sepkoski Jr, 1998] John Sepkoski Jr, J. (1998). Rates of speciation in the fossil record. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 353(1366):315–326.
- [Johnson, 2008] Johnson, N. (2008). Sewall wright and the development of shifting balance theory. *Nature Education*, 1(1):52.
- [Kantorovich, 1960] Kantorovich, L. V. (1960). Mathematical methods of organizing and planning production. *Management Science*, 6(4):366–422.

- [Kauffman et al., 1993] Kauffman, S. A. et al. (1993). *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA.
- [Kern and Hahn, 2018] Kern, A. D. and Hahn, M. W. (2018). The neutral theory in light of natural selection. *Molecular biology and evolution*, 35(6):1366–1371.
- [Kim and Stephan, 2003] Kim, Y. and Stephan, W. (2003). Selective sweeps in the presence of interference among partially linked loci. *Genetics*, 164:389–98.
- [Kimura, 1964] Kimura, M. (1964). Diffusion models in population genetics. *Journal of Applied Probability*, 1(2):177–232.
- [Kimura, 1983] Kimura, M. (1983). *The neutral theory of molecular evolution*. Cambridge University Press.
- [Kimura, 1987] Kimura, M. (1987). Molecular evolutionary clock and the neutral theory. *Journal of Molecular Evolution*, 26(1):24–33.
- [Kimura, 1991] Kimura, M. (1991). The neutral theory of molecular evolution: a review of recent evidence. *The Japanese Journal of Genetics*, 66(4):367–386.
- [Kimura and Crow, 1964] Kimura, M. and Crow, J. F. (1964). The number of alleles that can be maintained in a finite population. *Genetics*, 49:725–738.
- [Kimura and Ohta, 1969] Kimura, M. and Ohta, T. (1969). The average number of generations until fixation of a mutant gene in a finite population. *Genetics*, 61:763–71.
- [Kondrashov, 1988] Kondrashov, A. S. (1988). Deleterious mutations and the evolution of sexual reproduction. *Nature*, 336(6198):435–440.
- [Koza, 1994] Koza, J. R. (1994). Genetic programming as a means for programming computers by natural selection. *Statistics and computing*, 4:87–112.
- [Kramer, 2010] Kramer, O. (2010). Evolutionary self-adaptation: a survey of operators and strategy parameters. *Evolutionary Intelligence*, 3(2):51–65.
- [Kreitman, 2000] Kreitman, M. (2000). Methods to detect selection in populations with applications to the human. *Annual Review of Genomics and Human Genetics*, 1(1):539–559.
- [Kryazhimskiy et al., 2009] Kryazhimskiy, S., Tkačik, G., and Plotkin, J. B. (2009). The dynamics of adaptation on correlated fitness landscapes. *Proceedings of the National Academy of Sciences*, 106(44):18638–18643.
- [Kullback and Leibler, 1951] Kullback, S. and Leibler, R. A. (1951). On information and

- sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86.
- [Kuo et al., 2021] Kuo, Y. P., Arrieta, C. N., and Carja, O. (2021). A theory of evolutionary dynamics on any complex spatial structure. *bioRxiv*, pages 2021–02.
- [La Cava et al., 2016] La Cava, W., Spector, L., and Danai, K. (2016). Epsilon-lexicase selection for regression. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, pages 741–748.
- [LaBar and Adami, 2016] LaBar, T. and Adami, C. (2016). Different evolutionary paths to complexity for small and large populations of digital organisms. *PLoS Comput Biol*, 12:e1005066.
- [Lande and Arnold, 1983] Lande, R. and Arnold, S. J. (1983). The measurement of selection on correlated characters. *Evolution*, 37(6):1210–1226.
- [Langley and Fitch, 1974] Langley, C. H. and Fitch, W. M. (1974). An examination of the constancy of the rate of molecular evolution. *Journal of Molecular Evolution*, 3(3):161–177.
- [Lehman and Miikkulainen, 2015] Lehman, J. and Miikkulainen, R. (2015). Extinction events can accelerate evolution. *PLOS ONE*, 10(8):e0132886.
- [Lehman and Stanley, 2011] Lehman, J. and Stanley, K. O. (2011). Novelty search and the problem with objectives. In *Genetic Programming Theory and Practice IX*, pages 37–56. Springer.
- [Lenski et al., 2003] Lenski, R. E., Winkworth, C. L., and Riley, M. A. (2003). Rates of DNA sequence evolution in experimental populations of *Escherichia coli* during 20,000 generations. *Journal of Molecular Evolution*, 56(4):498–508.
- [Lewis, 1962] Lewis, H. (1962). Catastrophic selection as a factor in speciation. *Evolution*, 16(3):257–271.
- [Lewontin et al., 1974] Lewontin, R. C. et al. (1974). *The Genetic Basis of Evolutionary Change*, volume 560. Columbia University Press, New York.
- [Lieberman et al., 2005] Lieberman, E., Hauert, C., and Nowak, M. A. (2005). Evolutionary dynamics on graphs. *Nature*, 433(7023):312–316.
- [Lynch, 2007] Lynch, M. (2007). *The origins of genome architecture*. Sinauer Associates Sunderland.
- [Lynch et al., 2016] Lynch, M., Ackerman, M. S., Gout, J.-F., Long, H., Sung, W., Thomas, W. K., and Foster, P. L. (2016). Genetic drift, selection and the evolution of the mutation rate. *Nature Reviews Genetics*, 17(1111):704–714.

- [Lynch and Conery, 2003] Lynch, M. and Conery, J. S. (2003). The origins of genome complexity. *Science*, 302(5649):1401–1404.
- [Malamud et al., 1998] Malamud, B. D., Morein, G., and Turcotte, D. L. (1998). Forest fires: An example of self-organized critical behavior. *Science (New York, N.Y.)*, 281(5384):1840–1842.
- [Martínez-Jiménez et al., 2020] Martínez-Jiménez, F., Muiños, F., Sentís, I., Deu-Pons, J., Reyes-Salazar, I., Arnedo-Pac, C., Mularoni, L., Pich, O., Bonet, J., Kranas, H., et al. (2020). A compendium of mutational cancer driver genes. *Nature Reviews Cancer*, 20(10):555–572.
- [Mathias and Ragusa, 2016] Mathias, H. D. and Ragusa, V. R. (2016). An empirical study of crossover and mass extinction in a genetic algorithm for pathfinding in a continuous environment. In *2016 IEEE Congress on Evolutionary Computation (CEC)*, pages 4111–4118. IEEE.
- [Matsumura et al., 2012] Matsumura, S., Arlinghaus, R., and Dieckmann, U. (2012). Standardizing selection strengths to study selection in the wild: a critical comparison and suggestions for the future. *BioScience*, 62(12):1039–1054.
- [Maynard Smith, 1970] Maynard Smith, J. (1970). Natural selection and the concept of a protein space. *Nature*, 225:563.
- [Mayr, 1954] Mayr, E. (1954). Change of genetic environment and evolution. *Evolution as process.*, page 157–180.
- [Melbinger and Vergassola, 2015] Melbinger, A. and Vergassola, M. (2015). The impact of environmental fluctuations on evolutionary fitness functions. *Scientific Reports*, 5(11):15211.
- [Millidge et al., 2021] Millidge, B., Tschantz, A., Seth, A., and Buckley, C. (2021). Understanding the origin of information-seeking exploration in probabilistic objectives for control. *arXiv preprint arXiv:2103.06859*.
- [Möller et al., 2019] Möller, M., Hindersin, L., and Traulsen, A. (2019). Exploring and mapping the universe of evolutionary graphs identifies structural properties affecting fixation probability and time. *Nature Communications Biology*, 2:1–9.
- [Montanez, 2013] Montanez, G. D. (2013). Information transmission through genetic algorithm fitness maps. In *2013 IEEE Congress on Evolutionary Computation*, pages 756–763. IEEE.
- [Mouret and Clune, 2015] Mouret, J.-B. and Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.

- [Muller, 1964] Muller, H. J. (1964). The relation of recombination to mutational advance. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis*, 1(1):2–9.
- [Möller et al., 2019] Möller, M., Hindersin, L., and Traulsen, A. (2019). Exploring and mapping the universe of evolutionary graphs identifies structural properties affecting fixation probability and time. *Nature Communications Biology*, 2(11):1–9.
- [Mühlenbein, 1991] Mühlenbein, H. (1991). Parallel genetic algorithms, population genetics and combinatorial optimization. In Becker, J. D., Eisele, I., and Mündemann, F. W., editors, *Parallelism, Learning, Evolution*, Lecture Notes in Computer Science, page 398–406, Berlin, Heidelberg. Springer.
- [Mühlenbein, 2009] Mühlenbein, H. (2009). *Evolutionary Computation: Centralized, Parallel or Collaborative*, volume 1 of *Intelligent Systems Reference Library*, page 561–595. Springer Berlin Heidelberg, Berlin, Heidelberg.
- [Nevo, 2001] Nevo, E. (2001). Evolution of genome–phenome diversity under environmental stress. *Proceedings of the National Academy of Sciences*, 98(11):6233–6240.
- [Niklas, 1997] Niklas, K. J. (1997). *The Evolutionary Biology of Plants*. University of Chicago Press.
- [Niklas et al., 1983] Niklas, K. J., Tiffney, B. H., and Knoll, A. H. (1983). Patterns in vascular land plant diversification. *Nature*, 303(5918):614–616.
- [Ochs and Desai, 2015] Ochs, I. E. and Desai, M. M. (2015). The competition between simple and complex evolutionary trajectories in asexual populations. *BMC evolutionary biology*, 15(1):1–9.
- [Ofria and Wilke, 2004] Ofria, C. and Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, 10(2):191–229.
- [Ohta, 1972] Ohta, T. (1972). Population size and rate of evolution. *Journal of Molecular Evolution*, 1(4):305–314.
- [Ohta, 1992] Ohta, T. (1992). The nearly neutral theory of molecular evolution. *Annual Review of Ecology and Systematics*, 23(1):263–286.
- [Ohta, 2002] Ohta, T. (2002). Near-neutrality in evolution of genes and gene regulation. *Proceedings of the National Academy of Sciences*, 99(25):16134–16137.
- [Oliveto et al., 2018] Oliveto, P. S., Paixão, T., Pérez Heredia, J., Sudholt, D., and Trubenová, B. (2018). How to escape local optima in black box optimisation: When non-elitism outperforms elitism. *Algorithmica*, 80(5):1604–1633.

- [Orr, 2005] Orr, H. A. (2005). The genetic theory of adaptation: a brief history. *Nature Reviews Genetics*, 6(2):119–127.
- [Østman and Adami, 2014] Østman, B. and Adami, C. (2014). Predicting evolution and visualizing high-dimensional fitness landscapes. *Recent advances in the theory and application of fitness landscapes*, pages 509–526.
- [O’Donald, 1968] O’Donald, P. (1968). Measuring the intensity of natural selection. *Nature*, 220:197–198.
- [Palmer et al., 2012] Palmer, S. A., Clapham, A. J., Rose, P., Freitas, F. O., Owen, B. D., Beresford-Jones, D., Moore, J. D., Kitchen, J. L., and Allaby, R. G. (2012). Archaeogenomic evidence of punctuated genome evolution in gossypium. *Molecular Biology and Evolution*, 29(8):2031–2038.
- [Park et al., 2010] Park, S., Simon, D., and Krug, J. (2010). The speed of evolution in large asexual populations. *J. Stat. Phys.*, 138:381–410.
- [Pedro et al., 2021] Pedro, N., Silva, C. N., Magalhães, A. C., Cavadas, B., Rocha, A. M., Moreira, A. C., Gomes, M. S., Silva, D., Sobrinho-Simões, J., Ramos, A., Cardoso, M. J., Filipe, R., Palma, P., Ceia, F., Silva, S., Guimarães, J. T., Sarmiento, A., Fernandes, V., Pereira, L., and Tavares, M. (2021). Dynamics of a dual sars-cov-2 lineage co-infection on a prolonged viral shedding covid-19 case: Insights into clinical severity and disease duration. *Microorganisms*, 9:300.
- [Peischl et al., 2013] Peischl, S., Dupanloup, I., Kirkpatrick, M., and Excoffier, L. (2013). On the accumulation of deleterious mutations during range expansions. *Molecular ecology*, 22(24):5972–5982.
- [Peischl and Excoffier, 2015] Peischl, S. and Excoffier, L. (2015). Expansion load: recessive mutations and the role of standing genetic variation. *Molecular ecology*, 24(9):2084–2094.
- [Peischl et al., 2015] Peischl, S., Kirkpatrick, M., and Excoffier, L. (2015). Expansion load and the evolutionary dynamics of a species range. *The American Naturalist*, 185(4):E81–E93.
- [Poli et al., 2007] Poli, R., Kennedy, J., and Blackwell, T. (2007). Particle swarm optimization. *Swarm intelligence*, 1(1):33–57.
- [Price, 1972] Price, G. R. (1972). Extension of covariance selection mathematics. *Annals of Human Genetics*, 35(4):485–490.
- [Price et al., 1970] Price, G. R. et al. (1970). Selection and covariance. *Nature*, 227:520–521.
- [Ragusa and Bohm, 2022] Ragusa, V. and Bohm, C. (2022). Augmenting evolution with

- bio-inspired “super explorers”. In *ALIFE 2022: The 2022 Conference on Artificial Life*. MIT Press.
- [Ragusa and Bohm, 2021] Ragusa, V. R. and Bohm, C. (2021). Connections between noisy fitness and selection strength. In *ALIFE 2021: The 2021 Conference on Artificial Life*. MIT Press.
- [Rechenberg, 1984] Rechenberg, I. (1984). The evolution strategy. a mathematical model of darwinian evolution. In *Synergetics—From Microscopic to Macroscopic Order: Proceedings of the International Symposium on Synergetics at Berlin, July 4–8, 1983*, pages 122–132. Springer.
- [Rice and Chippindale, 2001] Rice, W. R. and Chippindale, A. K. (2001). Sexual recombination and the power of natural selection. *Science*, 294(5542):555–559.
- [Riskin, 1989] Riskin, H. (1989). *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer Verlag, Berlin, Heidelberg, and New York, 2nd edition.
- [Rivoire and Leibler, 2011] Rivoire, O. and Leibler, S. (2011). The value of information for populations in varying environments. *Journal of Statistical Physics*, 142(6):1124–1166.
- [Rivoire and Leibler, 2014] Rivoire, O. and Leibler, S. (2014). A model for the generation and transmission of variations in evolution. *Proceedings of the National Academy of Sciences*, 111(19):E1940–E1949.
- [Rouzine et al., 2008] Rouzine, I. M., Brunet, E., and Wilke, C. O. (2008). The traveling-wave approach to asexual evolution: Muller’s ratchet and speed of adaptation. *Theor Popul Biol*, 73:24–46.
- [Rouzine et al., 2003] Rouzine, I. M., Wakeley, J., and Coffin, J. M. (2003). The solitary wave of asexual evolution. *Proc Natl Acad Sci U S A*, 100:587–92.
- [Rozen et al., 2008] Rozen, D. E., Habets, M. G., Handel, A., and de Visser, J. A. G. (2008). Heterogeneous adaptive trajectories of small populations on complex fitness landscapes. *PloS one*, 3(3):e1715.
- [Sahney et al., 2010] Sahney, S., Benton, M. J., and Falcon-Lang, H. J. (2010). Rain-forest collapse triggered carboniferous tetrapod diversification in euramerica. *Geology*, 38(12):1079–1082.
- [Sastry et al., 2005] Sastry, K., Goldberg, D., and Kendall, G. (2005). *Genetic Algorithms*, page 97–125. Springer US, Boston, MA.
- [Shir, 2012] Shir, O. M. (2012). Niching in evolutionary algorithms. In Rozenberg, G., Bäck, T., and Kok, J. N., editors, *Handbook of Natural Computing*, page 1035–1069, Berlin,

Heidelberg. Springer.

- [Simpson, 1944] Simpson, G. G. (1944). *Tempo and Mode in Evolution*. Columbia University Press.
- [Sites and Moritz, 1987] Sites, Jack W., J. and Moritz, C. (1987). Chromosomal evolution and speciation revisited. *Systematic Biology*, 36(2):153–174.
- [Slater, 2013] Slater, G. J. (2013). Phylogenetic evidence for a shift in the mode of mammalian body size evolution at the cretaceous-palaeogene boundary. *Methods in Ecology and Evolution*, 4(8):734–744.
- [Slatkin and Excoffier, 2012] Slatkin, M. and Excoffier, L. (2012). Serial founder effects during range expansion: a spatial analog of genetic drift. *Genetics*, 191(1):171–181.
- [Smalley Jr. et al., 1985] Smalley Jr., R. F., Turcotte, D. L., and Solla, S. A. (1985). A renormalization group approach to the stick-slip behavior of faults. *Journal of Geophysical Research: Solid Earth*, 90(B2):1894–1900.
- [Smit and Eiben, 2009] Smit, S. and Eiben, A. (2009). Comparing parameter tuning methods for evolutionary algorithms. In *2009 IEEE Congress on Evolutionary Computation*, page 399–406.
- [Steinberg and Ostermeier, 2016] Steinberg, B. and Ostermeier, M. (2016). Environmental changes bridge evolutionary valleys. *Sci Adv*, 2:e1500921.
- [Takahata, 1987] Takahata, N. (1987). On the overdispersed molecular clock. *Genetics*, 116(1):169–179.
- [Templeton, 1980] Templeton, A. R. (1980). The theory of speciation via the founder principle. *Genetics*, 94(4):1011–1038.
- [Tkadlec et al., 2019] Tkadlec, J., Pavlogiannis, A., Chatterjee, K., and Nowak, M. A. (2019). Population structure determines the tradeoff between fixation probability and fixation time. *Communications Biology*, 2(11):1–8.
- [Turcotte et al., 1985] Turcotte, D. L., Smalley, R. F., and Solla, S. A. (1985). Collapse of loaded fractal trees. *Nature*, 313:671–672.
- [Uhlenbeck and Ornstein, 1930] Uhlenbeck, G. E. and Ornstein, L. S. (1930). On the theory of the Brownian motion. *Physical Review*, 36(5):823–841.
- [Van Egeren et al., 2018] Van Egeren, D., Madsen, T., and Michor, F. (2018). Fitness variation in isogenic populations leads to a novel evolutionary mechanism for crossing fitness valleys. *Communications Biology*, 1(11):1–9.

- [Van Valen, 1965] Van Valen, L. (1965). Selection in natural populations. III. Measurement and estimation. *Evolution*, 19(4):514–528.
- [Van Valen, 1967] Van Valen, L. (1967). Selection in natural populations. 6. Variation genetics, and more graphs for estimation. *Evolution*, 21(2):402–405.
- [Wang and Zhang, 2011] Wang, Z. and Zhang, J. (2011). Impact of gene expression noise on organismal fitness and the efficacy of natural selection. *Proceedings of the National Academy of Sciences*, 108(16):E67–E76.
- [Weinreich and Chao, 2005] Weinreich, D. M. and Chao, L. (2005). Rapid evolutionary escape by large populations from local fitness peaks is likely in nature. *Evolution*, 59:1175–82.
- [Weissman et al., 2009] Weissman, D. B., Desai, M. M., Fisher, D. S., and Feldman, M. W. (2009). The rate at which asexual populations cross fitness valleys. *Theoretical population biology*, 75(4):286–300.
- [Wilke, 2004] Wilke, C. O. (2004). The speed of adaptation in large asexual populations. *Genetics*, 167:2045–53.
- [Wilke, 2005] Wilke, C. O. (2005). Quasispecies theory in the context of population genetics. *BMC Evol Biol*, 5:44.
- [Wolf et al., 2006] Wolf, Y. I., Viboud, C., Holmes, E. C., Koonin, E. V., and Lipman, D. J. (2006). Long intervals of stasis punctuated by bursts of positive selection in the seasonal evolution of influenza a virus. *Biology Direct*, 1(1):34.
- [Wolpert and Macready, 1995] Wolpert, D. and Macready, W. (1995). No free lunch theorems for search. *Technical Report SFI-TR-95-02-010*.
- [Wolpert and Macready, 1997] Wolpert, D. and Macready, W. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82.
- [Wolpert and Macready, 2005] Wolpert, D. and Macready, W. (2005). Coevolutionary free lunches. *IEEE Transactions on Evolutionary Computation*, 9(6):721–735.
- [Wright, 1932] Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the Sixth International Congress of Genetics*, 1:356–366.
- [Wright, 1982] Wright, S. (1982). The shifting balance theory and macroevolution. *Annual review of genetics*, 16(1):1–20.
- [Yi, 2010] Yi, C. (2010). On the first passage time distribution of an Ornstein–Uhlenbeck

process. *Quantitative Finance*, 10(9):957–960.