

Death and Progress: How Evolvability is Influenced by Intrinsic Mortality

Frank Veenstra¹, Pablo González de Prado Salas², Kasper Stoy², Josh Bongard³
Sebastian Risi²

February 13, 2020

Frank Veenstra (corresponding author): f.veenstra@napier.ac.uk ; +4407708406824 ;

¹Edinburgh Napier University, Edinburgh

Pablo González de Prado Salas: pablogps86@gmail.com ; +34 661 73 65 09 ; ²IT University of
Copenhagen, Copenhagen, 2300

Kasper Stoy: ksty@itu.dk ; +4572185000 ; ²IT University of Copenhagen, Copenhagen, 2300

Josh Bongard: jbongard@uvm.edu ; +1 8026564665 ; ³UVM, Burlington, VT 05405

Sebastian Risi: sebr@itu.dk ; +45 7218 5018 ; ²IT University of Copenhagen, Copenhagen,
2300

Abstract

Evolvability is discussed extensively in evolutionary biology as well as evolutionary computation. There are many factors influencing evolvability and this paper illustrates how intrinsic mortality (death induced through internal factors) in an evolving population contributes favorably to evolvability on a fixed deceptive fitness landscape. To study the relationship between mortality and evolvability, we use the hierarchical if-and-only-if (H-IFF) function as our deceptive fitness landscape together with a steady state genetic algorithm (SSGA) with a variable mutation rate and intrinsic mortality rate. The results display a tight correlation between the mutation rate and an indiscriminate intrinsic mortality rate in regard to being able to find the global maximum on H-IFF. Additionally, a spatial model which includes extrinsic mortality factors is also employed, and shows a similar relationship

between mortality and mutation rate. To consider if a mortality rate is potentially valuable to evolutionary computation, we also compared the performance of the optimal mutation and mortality rate to a state of the art evolutionary algorithm called age-fitness pareto optimization (AFPO). The main results indicate that the intrinsic mortality rate and mutation rate cause random genetic drift that can allow a population to find the global optimum on the proposed landscapes. Intrinsic mortality in evolutionary algorithms convey how evolvability could be induced in natural systems as well as potentially improving the evolvability for existing evolutionary algorithms. This paper gives an overview of how intrinsic mortality influences the evolvability of a population and thereby supports the premise that the programmed death of individuals has a beneficial effect on the evolvability of the entire population.

Index terms— Intrinsic Mortality, Evolvability, Evolutionary Computation, Spatial Models, Senescence

1 Introduction

Considering classical Darwinism, senescence, the deterioration of function with age, directly opposes the evolutionary advantage longevity can have. The older an individual can become, the more offspring it could potentially produce. Hence, the personal fitness of an individual is directly increased by living longer. Nature is riddled with mechanisms that seem to restrict longevity, where senescence is prevalent in a multitude of species. Though the individual benefit is usually related to fitness, Darwin (1872) already noted that longevity is likely a product of the complex interactions between a species and its environment. Therefore, could longevity itself somehow be a determinant for the evolutionary trajectories we see in nature?

Most octopuses are semelparous, reproducing only once in their lifetime. This is an observation already mentioned in Aristotle's History of Animals (Aristotle, 1910): They live young and die fast (O'Dor & Webber, 1986). The process of senescence in the *Enteroctopus dofleini* for example, is regulated by secretions from an endocrine gland that normally causes death by starvation after reproduction (Anderson et al., 2002). In this case, the octopus suddenly stops foraging and instead takes care of its eggs and hatchlings, followed by the eventual death of the octopus parent. However, by simply removing this endocrine gland, octopuses seemingly live significantly longer than usual, and they can reproduce more than once (Wodinsky, 1977). Senescence in octopuses is particularly elusive, and the true advantage of this type of senescence might be caused by various phenotypic traits and selection pressures. Does a short life have an evolutionary advantage, or is the decreased lifespan a byproduct of the mechanism inhibiting foraging behavior, which enables the protection of offspring with the dire side effect of mortality?

Other animals, such as particular salmon, undergo a similar process of senescence: dying after having laid eggs. Some spiders are cannibals and kill their male counterpart after copulation (cannibalism is also a feature exhibited by some octopuses). Elephants run out of teeth, a form of mechanical senescence, while some turtles express negligible senescence (not showing aging symptoms). Naked mole rats can become significantly older than other rodents. Artificial selection of drosophila can allow them to live 50% longer after a few generations (Nusbaum & Rose, 1994), and the proteins DAF-2 and DAF-16 directly regulate life span in *C. elegans* (Lin et al., 2001). Furthermore, it has been shown that long-lived yeast mutants are outcompeted

by shortlived wildtypes (Kyryakov et al., 2016). Longevity and aging thus seem to emerge differently across species. But what is the evolutionary value of senescence if there is any? Considering recent publications by Kowald & Kirkwood (2016), and Goldsmith (2016), which discuss whether aging is programmed or not, it is relevant to test whether mortality poses any benefit for evolutionary algorithms (EAs) that could support any of the existing theories on aging.

One of the theories of senescence suggests that senescence can promote the evolvability of a population (Goldsmith, 2014). Evolvability is the *population's* ability to traverse the fitness landscape without passing through non-functional regions (Smith, 1970; Haubold & Wiehe, 2006), and we define it as follows:

Evolvability The ability of a *population* to create *adaptive* genetic diversity across *generations*.

Hereby, we consider: (1) the entire population (or gene pool) instead of the individual, (2) genetic diversity to be adaptive (leading to a better-fit solution, in contrast to non-adaptive genetic diversity), and (3) the measurement of evolvability over many generations instead of one. The definition of evolvability used here is different from some existing measures of evolvability (Altenberg, 1994; Wagner, 1996; Lehman, 2012). The main difference of the definition we gave is that evolvability should be measured over several generations. The ability of the individual or a population to produce better-adapted offspring or more diverse individuals in one generation is less important than the ability of the population to keep producing better offspring across generational time. A population that can produce fitter individuals in one generation might be unable to traverse a local *valley* in the fitness landscape that needs to be crossed to find a better solution. Therefore, we analyze the efficiency at which a population is able to traverse the fitness landscape over generational time as a proxy of evolvability ¹.

Intrinsic mortality leading to evolvability has been mostly discussed hypothetically (Goldsmith, 2014). Notably, Herrera et al. (2016) investigated the evolvability of a population of agents in a rapidly changing environment. They show that a terminal age allows the population to better continuously adapt to its environment. In addition, Lehman et al. (2015) showed that extinction events can lead to a better evolvable EA. However, in contrast to the work presented

¹See Veenstra (2018) for a discussion on evolvability

here, extinction events were discriminative and kept certain *elites* in the population. Although evolution is usually seen as including incremental improvements over generational time, some solutions might require evolutionary steps that make individuals *worse* than their ancestors. This less elitist approach could enable progeny to find a solution in the search space that is more distant, and perhaps, ultimately more efficient, than the ancestral solutions. As a testbed for this potential leap, we used a deceptive fitness function—an adjusted version of the hierarchical if-and-only-if (H-IFF) function (Watson et al., 1998)—and a genetic algorithm (GA) to simulate an evolving population. We show that on this testbed, intrinsic mortality alters the evolvability of the population. In addition, a spatial model that inherently contains extrinsic mortality factors, is used to show that the added intrinsic mortality rate still changes the evolvability of the population. Though too much mortality is detrimental, depending on the mutation rate of the population, a certain mortality rate actually increases the evolvability of the evolving populations.

We have described in Veenstra et al. (2018) how this intrinsic mortality parameter can be beneficial for finding solutions on the H-IFF function. This paper extends the previous by including additional results on the H-IFF landscape and a comparison about how a population evolves using our method and age-fitness pareto optimization (AFPO) (Schmidt & Lipson, 2011). It furthermore highlights the relevance of these results to nature, supporting the theory of programmed death inducing evolvability. The next sections of the introduction briefly give an overview on the existing theories of senescence and whether mortality can be programmed (subsection 1.1 & 1.2). It follows up with a brief illustration about how intrinsic mortality could promote evolvability (subsection 1.3). Subsequently, the experimental setups of the evolutionary algorithms and the spatial model (section 2), and their results are subsequently presented (section 3). Afterwards, the results, and their impact, are discussed in section 4. Overall, this paper demonstrates not only the evolutionary advantage of senescence in nature, but also how this concept could potentially be beneficial in evolutionary computation.

1.1 Summary of Theories on Senescence

Mortality is a fundamental component of natural systems that is caused either by intrinsic factors (senescence) or extrinsic factors such as predation, disease, and accidents. It initially seemed that aging is an evolutionary disadvantage for individuals since their personal fitness is lowered when an individual dies from internal mechanisms. There are, however, several theories explaining the cause and function of this biological phenomenon as an alternative to being a direct disadvantage. Darwin already mentioned in the Sixth Edition of *On the Origin of Species* (not in older editions) that longevity is related to the scale of organization, expenditure, and general activity of organisms, which has likely been determined by natural selection (Darwin, 1872). Although having an evolutionary disadvantage for the individual, it may have several advantages for the maintenance of a species. Furthermore, Weismann. (1889) has claimed that aging is determined by the “needs of the species”, which is subject to the same mechanical process of regulation as other structures and functions of organisms. Weismann also proposed that this intrinsic mortality makes a population more evolvable since there is a higher turnover rate of individuals.

An alternative theory by Medawar (1952) proposes that aging could just be a phenomenon that arises due to the simple neglect of selective pressure on older organisms, and older organisms are, by chance, more likely to have encountered mortality-inducing factors that limit their lifespan. Considering a steady probability of death for each individual in a specific population, a survivor curve can be created displaying how many individuals in a specific age group are alive (Figure 1). Consider a population of 100 individuals wherein every month there is a specific chance for individuals to die—a fixed probability of death, or mortality rate. This probability shapes the age range that evolutionary selection can act on. In the case of a 20% mortality chance every month, the range of selection is quite low, whereas this range is quadrupled when the chance is only 5% (Figure 1). Depending on this rate, the number of older individuals in a population differs. Since there are less older individual as displayed by the survivor curve, the selection pressure on old age is therefore reduced. Due to the lower selection pressure on old individuals, an accumulation of mutations can therefore creep in, which results in the senescence of the individuals. An *accumulation of mutations* leading to senescence could be the result of

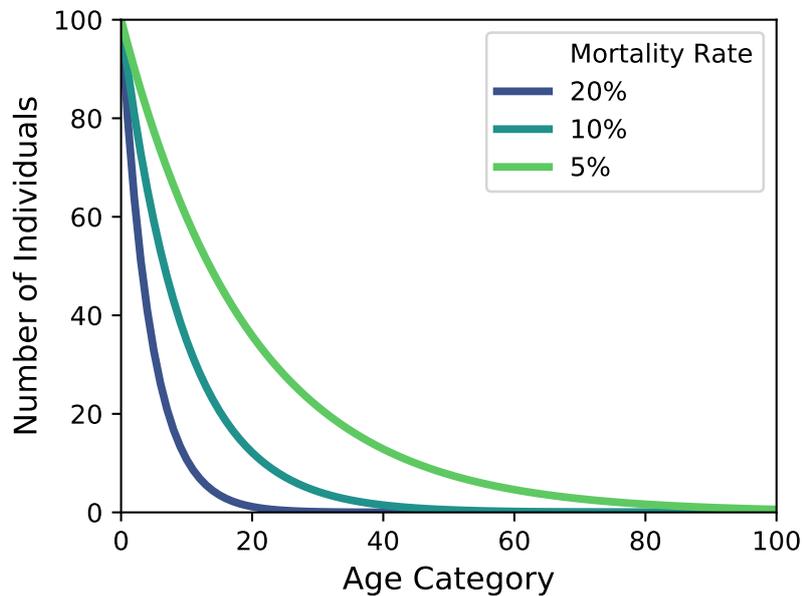


Figure 1: **Survivor curves.** Considering a fixed number of individuals entering the population every iteration, and accepting that there is a 5%, 10% or 20% probability of an individual dying by chance, it can be seen that after a certain number of iterations there will simply not be any survivors left in a specific age category. Based on Medawar (1952).

a decline in selection pressure on older individuals. In this scenario, older individuals are more likely to die due to mere chance or wear-and-tear processes. Genes beneficial in early life would therefore have a higher selective advantage and thereby a higher chance to propagate themselves into the next generation. However, mere accumulation of mutations as an explanation for senescence is difficult to hold for most species since it has been shown that single genes that cause aging have been conserved throughout different species over evolutionary time (Guarente & Kenyon, 2000).

Genes could also have evolved to be beneficial early in life while being potentially deleterious later in life. This is a theory known as *antagonistic pleiotropy* (Williams, 1957). In this theory, a gene can have a pleiotropic effect by promoting reproductive success and survival early in life, while being detrimental later in life. Through this process, evolutionary biologists have argued that this inherent trade-off makes it difficult for natural selection to evolve old age in the first place. Although the scientific literature contains an ubiquity of examples of pleiotropic genes,

it is more difficult to see how mutations in genes can improve the personal fitness early in life, while having a deleterious effect later in life. These genes are sometimes referred to as “putative” disease alleles (Carter & Nguyen, 2011) since evidence that these types of alleles really have a benefit early in life has yet to be acquired.

Another alternative theory to the accumulation of mutations and antagonistic pleiotropy theories is the *disposable soma hypothesis* proposed by Kirkwood (1977), which is one of the dominating theories today (Shefferson et al., 2017). In the disposable soma theory, the body of an individual organism can allocate limited resources to various cellular processes and needs to make compromises between its metabolism, reproduction, repair and maintenance functions. For example, a population only focusing on repair can be outcompeted by a population that instead spends more energy on growth. Combined with Medawar’s survival curve, wherein death can also be caused by extrinsic factors such as predation, this would suggest that maintenance and repair are also of lesser importance later in life, since the probability of an individual reaching old age by chance is already low. Not allocating any resources to the repair of an organism with increasing age would thus lead to an organism’s deterioration with age as a side effect.

Other more recent theories consider the potential altruistic effect of senescence in which aging can be beneficial for coping with a changing environment (Yang, 2013; Mitteldorf & Martins, 2014; Herrera et al., 2016). In this case, it has been artificially shown that a terminal age is beneficial for a population in rapidly changing environments that necessitate adaptive changes in the genome. Similarly, a resulting benefit from senescence, or intrinsic mortality, is the reduction of over-consumption of environmental resources that gives a selective incentive for intrinsic mortality (Werfel et al., 2017). In addition, Lehman et al. (2015) showed that extinction events could lead to a better evolvable evolutionary algorithm, though in this case the extinction events were discriminative and kept certain *elites* in the population. Considering the previous results, from an optimization perspective, it seems that senescence, or simply mortality, can be beneficial for a population regarding evolvability and altruistic aging.

The evolvability theory of senescence claims that senescence increases the evolvability of a population (Goldsmith, 2014). For this case, we think mortality may aid evolvability in two ways:

1. When individuals die, a higher turnover rate of new individuals arises in the population (proposed by Weismann. (1889))
2. Mortality reduces selective pressure on the best individuals in the population and thereby decreases convergence, promoting diversification

For the first reason, a greater number of individuals that can live during a specific period leads to a larger proportion of acquired genetic adaptations by a population. In other words, more individuals that have been “evaluated” in the environment yield more individual fitness results from potentially differing phenotypes. Mortal populations thus contain a higher turnover rate of individuals as compared to immortal populations. The second argument implies that if older, fit individuals outcompete younger, slightly less fit individuals, the population can get stuck in a local optimum, or be in a state that is considered less evolvable.

The actual mechanisms of senescence would most likely be a combination of the theories of senescence that have been described. However, by considering the population (or gene pool) as a whole instead of thinking about the benefits for the individual, there is no reason for individually detrimental phenotypic traits not to pose an evolutionary advantage. For instance, if mutations are the main factor driving senescence, mutations also drive evolution due to the introduction of new variations of genes in the population. In contrast, a non-mutating population with an immortal life would reside in a zero-evolvability state (Goldsmith, 2008). The evolution of complex organisms can thus be a compromise between evolvability of the species and personal benefit to the individual (Goldsmith, 2008).

The antagonistic pleiotropy theory might surely be an explanation for senescence, though if senescence turns out to be beneficial for a population, this antagonistic effect would actually be an altruistic effect. If the personal fitness of an individual could be prolonged by adjusting the self-repair energy expenditure as mentioned by the disposable soma theory, there are evolutionary pressures toward better self-repair mechanisms. However, the lack of self-repair might also simply lead to more mutations, making a population more evolvable. Alternatively, the lack of self-repair mechanisms could lead to senescence, which could be advantageous as well. Essentially, if there is a selective pressure toward senescence, we can speak of programmed death.

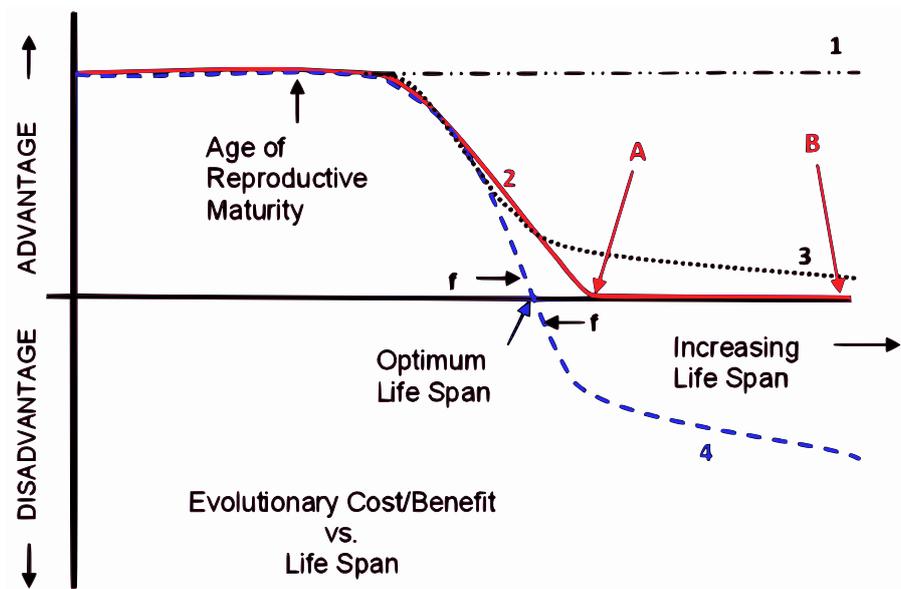


Figure 2: **Theories of aging plotted as an advantage and a disadvantage over time.** Line 1 denotes the advantage of longevity of an individual if longevity would not lead to a decrease in personal or inclusive fitness. Line 2 (solid red line) represents Medawar (1952), where the advantage of longevity would decrease with age but would not have a negative impact. Antagonistic pleiotropy and disposable-soma theories are represented by line 3 (dotted). In this case an increased lifespan does have an advantage to the inclusive fitness of a population but decreases with age. Weismann and Goldsmith support concept 4 where there exists an optimum life-span. Individuals that grow older than this lifespan will have a negative impact on the population. The figure is taken from Goldsmith (2014).

1.2 Programmed vs. Non-programmed Senescence

The theories on senescence can be further categorized into non-programmed and programmed theories. In Kowald et al. (2016), non-programmed senescence supports theories such as mutation accumulation, (Medawar, 1952), antagonistic pleiotropy, (Williams, 1957), and disposable-soma theory (Kirkwood, 1977). In contrast, in Goldsmith (2016) programmed aging theories on evolvability (Weismann., 1889; Goldsmith, 2016; Mitteldorf et al., 2014) and altruistic aging (Yang, 2013; Werfel et al., 2017; Herrera et al., 2016) theories are supported, and this has been corroborated by biological examples. It seems that supporters of the non-programmed theories generally exclude programmed theories (Kowald et al., 2016; Shefferson et al., 2017) whereas the programmed senescence supporters do not specifically exclude non-programmed theories. The impact of the theories considering lifespan as an advantage or disadvantage are summa-

rized in Figure 2 taken from Goldsmith (2014), where the increased age is only considered to be a disadvantage in line 4 (supporting programmed senescence), but not in lines 1 (classical no negative effect), 2 (mutation accumulation) and 3 (antagonistic pleiotropy and disposable soma). In Kowald et al. (2016), Figure 2 has been critiqued due to line 3 not representing the antagonistic pleiotropy and disposable soma theory well. It is critiqued since the parameters of these theories lead to a specific average lifespan that is optimal; hence, mortality is an emergent factor. However, if there were a gene that, at no cost, would improve the lifespan of the organism, it would have an advantage for both the antagonistic pleiotropy and disposable soma theories, but not for the programmed aging theories. Age itself is not the disadvantage, but rather the disadvantage is the product of the pleiotropic gene, or the trade-off between soma, respectively. Therefore, the distinction is still valid but should be taken lightly. The main useful distinction is the evolutionary advantage of the process of senescence—or the evolutionary disadvantage of long life posed on the gene pool—which line 4 displays.

1.3 Why Intrinsic Mortality can Promote Evolvability

In order to understand how mortality can induce evolvability, it is useful to clarify some concepts from evolutionary dynamics that led to the experiments of this paper. Considering any population of individuals at carrying capacity in an environment and stating that the mortality rate is fixed in this population, the mutation rate greatly influences the types of genes in the population and the resulting stable attractor space in a quasispecies equilibrium. As explained by Nowak, Martin (2006), when considering a sequence space of a specific gene, there can be several optima in this space (Schuster & Swetina, 1988). Depicted in Figure 3, if the average mutation rate u is below a specific critical value u_1 , the stable (robust) state of the gene in the population will end up in a narrow peak. When the mutation rate is at a value between u_1 and u_2 , the narrow peak becomes an unstable region in the sequence space for the population and the population will in turn converge to the broader sequence space with a lower maximum fitness value. If the mutation rate is in turn increased to be higher than u_2 , there will be no stable state and the sequence space in which the population resides is random. However, if genes in a population of individuals already reside in the broader less-fit state, how can they traverse the sequence space to end up

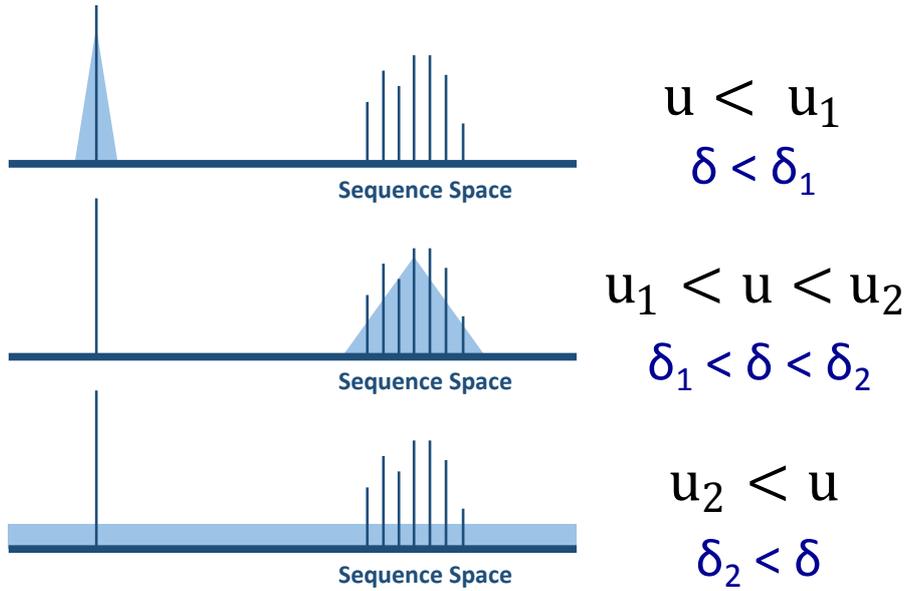


Figure 3: **Fitness landscape with one peak and a hill based on the sequence space.** u represents a mutation rate and δ represents the mortality rate in a given population. The two threshold values for u and δ are believed to have a similar effect on the stable region the population occupies on the fitness landscape.

in the narrow peak that is the better-fit solution? Traversing this fitness landscape would either require an individual to drastically mutate into that region, or a population to gradually move to the region through genetic drift. Nowak’s mutation-rate threshold values are, moreover, only valid for a population of mortals. If immortality could occur, the immortal individual residing in the narrow peak will always stay there (since it cannot be outcompeted) and eventually, its offspring also have a chance to occupy the narrow region, no matter how high the mutation rate is. Mortality in Nowak’s model is thus a requirement. Therefore, if we consider δ to be the mortality rate, we claimed in Veenstra et al. (2018) and Veenstra (2018) that there exists a mortality rate threshold δ_1 and δ_2 similar to the mutation rate thresholds (Figure 3). This is the initial hypothesis that has formed the premise of why mortality promotes evolvability.

Using H-IFF as the difficult-to-solve deceptive fitness-landscape on both an steady-state genetic algorithm SSGA and a spatial model can help us understand how this relationship influences the evolvability of a population. The SSGA is used as an abstract model to view the general effects of mortality on the evolutionary progression on this deceptive fitness landscape. Additionally, as described by Werfel et al. (2017), spatial models can elucidate aspects of mortality

that equate to natural systems. The spatial model—that contains an inherent extrinsic mortality rate emerging from local competition—is used to isolate the influence of intrinsic mortality to see whether it affects evolvability in natural systems.

2 Methodology

The experiments are divided in a benchmark optimization implementation using a SSGA and an agent-based grid model ². In both approaches, the fitness value of an individual is calculated based on the H-IFF function. The selection/deletion operators in the spatial model are inherent properties between the interactions of the individuals and their environment while the SSGA uses a random selection operator. Including an extrinsic mortality mechanic in the spatial model allows us to (1) test whether the mutation rate can alter the stable region in the sequence space of the genomes as explained by Nowak, Martin (2006) (Figure 3), and (2) investigate if an additional intrinsic mortality rate influences the evolvability. The SSGA uses both 64-bit as 128-bit genomes. In addition, the SSGA is compared to a state of the art evolutionary algorithm called Age-Fitness Pareto Optimization (AFPO)(Schmidt & Lipson, 2011) and we discuss how each implementation traverses the fitness landscape differently. The experiments investigate whether mortality alters the evolutionary progression of a population of individuals containing binary genomes and whether this enables the population to traverse the state space landscape more efficiently.

2.1 Hierarchical if-and-only-if

There are many binary fitness-landscapes that have been designed to be convoluted and deceptive, such as the Chuang fl (Chuang & Hsu, 2010), the Royal Road (Mitchell et al., 1991), or the hierarchical if-and-only-if (H-IFF) function (Watson et al., 1998). H-IFF is the function that we use throughout the paper to represent our deceptive fitness landscape. The H-IFF function creates a fractal deceptive fitness landscape and can be used to evaluate the performance of implementations of evolutionary algorithms. In H-IFF, a binary genome is evaluated based on

²The source code for the SSGA and the agent-based grid model can be found here: <https://github.com/FrankVeenstra/ALife2018>

Table 1: **Example of the score of a 16-bit H-IFF genome.** The table shows the fitness derived from the bit string 0001-1011-1111-1111. In each layer, whenever two sequences of bits above a cell are similar, a value is added to the total score of that layer. After summing all the scores in each layer, the scores are multiplied depending on the layer number. For example, the score of layer two is multiplied by 2, the score of layer three is multiplied by 4 and the score of layer four is multiplied by 8. This genome has a fitness value of 14 (out of a maximum of 32). Adapted from Watson et al. (1998).

	0	0	0	1	1	0	1	1	1	1	1	1	1	1	1	sum
1	1		0		0		1		1		1		1		1	$6 \cdot 1$
2	0			0			1			1			$2 \cdot 2$			
3	0						1						$1 \cdot 4$			
4	0															$0 \cdot 8$

self-similarity. One can check for self-similarity in the genome in multiple layers by initially checking the similarity of a pair of bits across the genome, continuing in the next layer by checking a pair of two bits, followed by checking a pair of four bits, etc. In each layer, a fitness value can be ascribed to the self-similarity score of the genome. This score is usually derived from the number of self-similar parts in the genome and the layer depth that is being checked.

As an example, Table 1 illustrates how one can derive a fitness value from a 16-bit genome that results in four layers on which to check for self-similarity. Note that in the original implementation a *null* bit was possible in the genome, which resulted in an additional fifth layer. This *null* possibility has been omitted in this paper to increase the computational efficiency and ease the visualization of genetic change over time. The omission of this *null* possibility makes the evolutionary progression easier to visualize by plotting the fitness value over the number of ones in the genome (Figure 4). The gray area in Figure 4 illustrates the possible fitness values an individual can achieve when having a certain number of ones in its genome. The landscape is, however, unchanged, with as many local optima as the original implementation. The maximum fitness of an individual with either only zeros or only ones is 192, and this is the value of the global maxima on H-IFF. However, when half of the genome is composed of zeros and the other half of ones, the fitness value of that particular individual ranges somewhere between 4 and 160 depending on the specific order of the bits.

For H-IFF, there are two potential global maximums regardless of the length of the genome. One global maximum contains a bit string of only ones while the other contains only zeros. In between these extremes, there are many local optima and one can generally state that there is

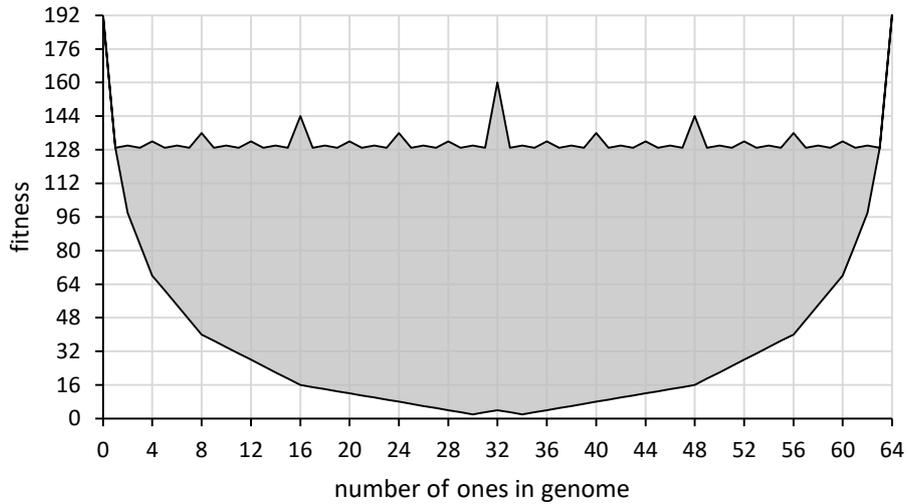


Figure 4: **Fitness landscape of the adjusted H-IFF function.** The global optima in this landscape are at the edges of the distribution and there are local optima fractally dispersed between every other two local optima.

a local optimum between each two high optima. This makes the landscape inherently fractal and deceptive. To understand how different sets of genomes correspond to fitness values based on the landscape produced by H-IFF, Figure 5 illustrates how four genomes are located on the fitness landscape of Figure 4 using the explanation in Table 1. A score for self-similarity in this illustration is simply denoted by a red color. The area of the fields of the table being colored red directly translates into the fitness value. In this example, the best fit genome is a bit string of only zeros.

Considering a binary genome that consists of 64 bits, there are a total of 2^{64} total possible configurations of a genome. This large search space in turn makes it difficult for algorithms to solve this binary function. Moreover, the 128-bit example that is used as well thus contains an immense search space of 2^{128} total possible configurations.

2.2 Steady State Genetic Algorithm

The steady state genetic algorithm (SSGA) with a mortality rate implements a 64-bit genome composed of either ones or zeros which are randomly initialized. Genes in the genome are mutated with a probability given by the mutation rate. Note that mutating a gene will randomly assign a bit of 1 or 0, so the gene swaps a bit with half the probability of the mutation rate in mutation

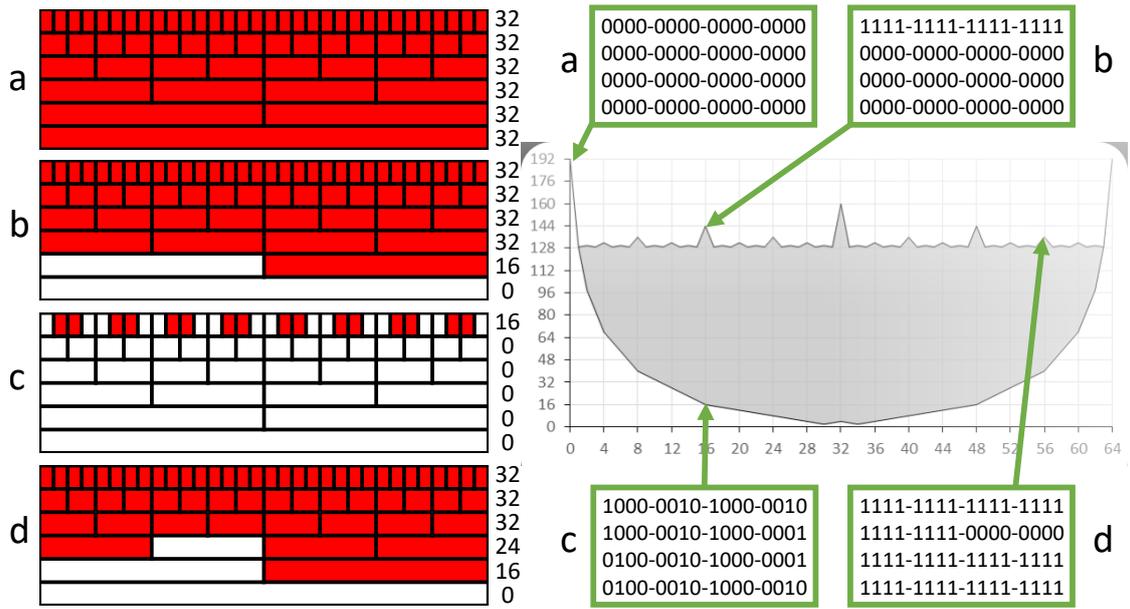


Figure 5: **Explanation of the H-IFF function.** Four genomes of length 64 are shown (a,b,c,d) with their corresponding fitness value. Left shows the scoring tables where red indicates a reward for self-similarity, as shown in Table 1. The red area directly translates into a fitness value of the individuals.

events. Thus, a mutation rate of 0.1 means a gene is mutated with a 10% probability but changed with only 5% probability. This implementation ensures that a mutation rate of 1.0 does not produce offspring with the complementary bit string of their parent’s genome, but rather an entirely random set of bits.

Each SSGA iteration is as follows: (1) a random individual is chosen, (2) the chosen genome is copied, mutated and evaluated, (3) the new genome is compared to a random individual in the existing population, and (4) the new genome replaces this second individual if the fitness for the new genome is higher. For a population size n , a generation consists of n iterations. After each generation, individuals are independently checked for deletion with a probability given by the mortality rate. The population was logged after each generation. No crossover was implemented to isolate the effect of the mutation rate.

We performed three experiments using the SSGA. The first determined the correlation between the mortality rate and the mutation rate on 64-bit H-IFF, the second on 128 bit H-IFF, and the third compared the SSGA with mortality to age-fitness pareto optimization (AFPO) Schmidt &

Lipson (2011). For the first experiment, 20 simulations of 100,000 generations and a population size of 50 individuals with different values for the mortality rate and mutation rate. A mutation rate sweep from 0.0 to 1.0 was done changing the mutation rate approximately exponentially. A similar sweep was done for the mortality rate, although the 0.64 and 1.0 mortality rates have been excluded from the results since these values led to early extinction of the population. For the second experiment, again, 20 simulations of 100,000 generations and a population size of 50 was used with varying mutation rates and mortality rates. In the last experiment we compared the optimal mutation/mortality rate found by doing this sweep to the optimal mutation rate of AFPO on 64-bit H-IFF again using a population size of 50.

To briefly illustrate, AFPO starts by initializing a population of random individuals. Each individual has been assigned an age of 0, and these individuals and their asexually produced offspring increment their age value by one after each generation. After evaluating all of the individuals in this initial population, all fittest individuals in all possible age categories are kept while the rest of the individuals are discarded. Afterwards, all ages of all individuals are incremented, and a random individual is inserted into the population. This random individual has an age of 0, similar to the initial population. The new population, including the random individual, is used to generate an offspring population that is the same size as the maximum population size. Afterwards, tournament selection (with tournament size 2) determines which individuals from the offspring population and the existing population will form the population in the next generation. This tournament is, however, also steady state since offspring are set up to compete with an individual in the original parent population sequentially. An individual can only outcompete and replace another individual, if its fitness is higher and its age is lower, thus dominating on both fronts.

For this last experiment using AFPO, we ran 200 simulations using these optimal parameters and kept track on how many evolutionary runs found the global maximum on 64-bit H-IFF. The optimal mutation rate was around 0.1 and was subsequently set to 0.1 in the experiments. The optimal mortality rate was tuned using the 0.1 optimal mutation rate of AFPO, which resulted in using a mortality rate of 0.05.

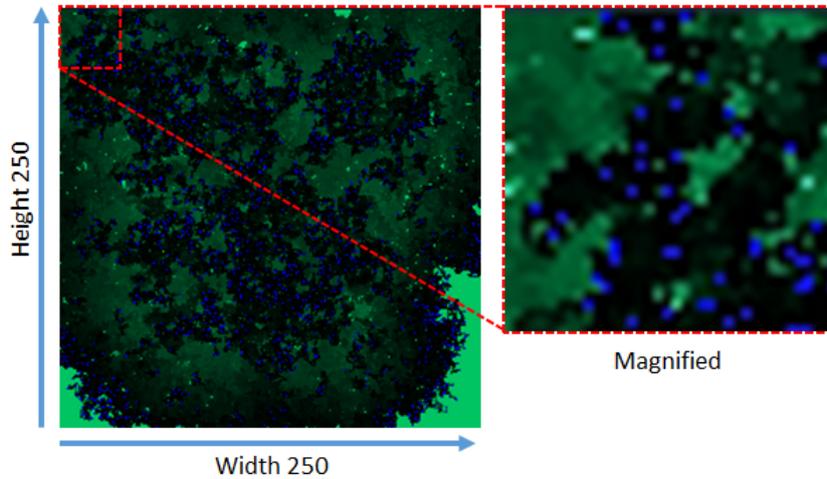


Figure 6: **Illustration of the spatial model.** Green represents plant biomass, blue rabbit biomass. Snapshot taken after the first few cycles of the spatial model

2.3 Spatial Model

The spatial model is an agent-based grid model. Like the SSGA, the spatial model implements the H-IFF fitness landscape, though it only implements 32-bit H-IFF to reduce computational requirements, which were considerably higher compared to the SSGA. In contrast to the steady state implementation, the population was initialized with genomes in the middle local optima of the H-IFF fitness landscape (0000-0000-0000-0000-1111-1111-1111-1111) with a corresponding fitness value of 64. From this starting configuration it is particularly challenging to find the global optima since no random individuals are close to any of the global optima which can be the case in random initialization. Moreover, when randomizing genomes, the fitness of random individuals can by chance be so low that the population is never able to survive. Hence, the local optimum was chosen as the initial genome of all individuals in the population.

The spatial model is similar to a predator-prey model and consisted of a 250×250 grid where cells are either type 0 (prey) or 1 (predator). One can imagine the prey and predators to be plants and rabbits respectively, where predators were subject to evolution and each predator cell contained a binary genome. The fitness value derived from a predator's genome translates into food consumption efficiency (a metabolic efficiency). The ability to efficiently acquire food from the environment enables predators to grow faster, thereby producing more offspring. At

each iteration, a given amount of biomass is added to plant (prey) cells, according to a biomass production rate, which was an absolute value of +0.0016, with a maximum biomass value of 1.0. Predator cells attempted to move to a neighboring cell with 4/5 chance (moving up, down, left, right or stay put). If the target position was occupied by another predator, or if the position is out of the grid, the predator did not move. For computational efficiency, the grid is sequentially updated from left to right and top to bottom, ensuring that predators move only once per iteration. Predator cells reproduce with a 1/10 chance if their biomass is above twice the reproduction cost (reproduction cost was 0.4). Offspring start with a biomass equal to the reproduction cost times 0.8 (0.32). The reproduction cost was subtracted from the biomass of the parent predator. All parameters used were chosen based on prior experimentation.

When a predator cell moves on a prey cell it consumes the prey's biomass with an efficiency rate of ($fitness/maxfit$). The predator will not increase its biomass over the 1.0 limit; any unused plant biomass is left in the cell. At every iteration, predator cells lose 0.02 biomass as a maintenance cost. Predators with biomass below 0.01 are removed from the population (extrinsic mortality from starvation).

In contrast to the SSGA, the spatial model implements both intrinsic and extrinsic mortality. For intrinsic mortality we implemented a terminal age, and the extrinsic mortality results from local competition, which is similar to the implementation of Werfel et al. (2017). To explore the relationship between intrinsic mortality and mutation rate different mutation rates and terminal ages are compared. The main results indicate how often, and how quickly, the global maximum was found on 32-bit H-IFF.

3 Results

In both the SSGA and the spatial model, a correlation was found between the mortality rate and mutation rate, and the performance of evolving populations on H-IFF. This correlation is clearly found on 64-bit H-IFF. However, 128-bit H-IFF would be too expensive computationally to produce such clear results. We do include some experiments on 128-bit H-IFF to demonstrate that it is indeed solvable using the SSGA and a mortality rate. The clear correlation on 64-bit H-IFF and the spatial model suggest that intrinsic mortality could lead to more evolvable

Table 2: **Times the optimal solution is found in the ssga.** Varying the mutation rate (u) and terminal age (δ). Results are taken from 20 runs of each set of parameters on 64-bit H-IFF. The lower-case value represents the average number of generations (thousands) that had to be simulated before finding the global optimum. Mutation rates above 0.32 and below 0.01 have been omitted since the global optima is never found in these scenarios

$u \backslash \delta$	0.0	0.005	0.01	0.02	0.03	0.04	0.06	0.08	0.12	0.16	0.24	0.32
0.01	0	0	0	0	0	0	0	0	0	0	0	2
0.02	0	0	0	0	0	0	0	0	0	0	2	2
0.03	0	0	0	0	0	0	0	0	0	0	20 ₂₆	0
0.04	0	0	0	0	0	1	0	0	0	14 ₄₃	3	0
0.06	0	0	0	0	0	0	0	3	20 ₁₇	11 ₅₀	0	0
0.08	0	0	0	0	0	1	18 ₂₂	19 ₁₉	3	0	0	0
0.12	0	1	0	15 ₁₈	19 ₃₄	7	0	0	0	0	0	0
0.16	2	11 ₃₅	17 ₃₅	3	0	0	0	0	0	0	0	0
0.24	4	0	0	0	0	0	0	0	0	0	0	0
0.32	4	0	0	0	0	0	0	0	0	0	0	0

populations, at least as an alternative to a high mutation rate.

3.1 Steady State Genetic Algorithm

The number of times the global maximum was found in each of the 20 evolutionary runs is presented in Table 2. In addition, the subscript values in the table represent the average number of generations ($\cdot 10^3$) it took the runs to find the global maximum, on average. As can be seen in the table, the relationship between the mortality rate and the mutation rate in the SSGA is very specific for finding the global maximum on 64-bit H-IFF. Moreover, the mutation rate and mortality rate explain 89% of the variation seen in the ability to reach to the global optimum in H-IFF (Figure 7). The results of the 20 runs using the SSGA are displayed in Figure 8 (top). To see how single runs are able to traverse the fitness landscape, Figure 8 (middle and bottom) depicts the fitness and diversity on the H-IFF landscape over generational time. Depicted by blue dots on the graphs are the individuals of a population at specific intervals.

Though the optimal ratio between the mutation rate and the mortality rate changes when varying the size of the genome, different optimal rates still exist for the larger genome. On 128-bit H-IFF, which contains an immense number of possible configurations almost 20 orders of magnitude larger than 64-bit H-IFF, evolving populations are still able to find the global maximum within 100,000 generations when a specific mutation rate to mortality rate ratio is

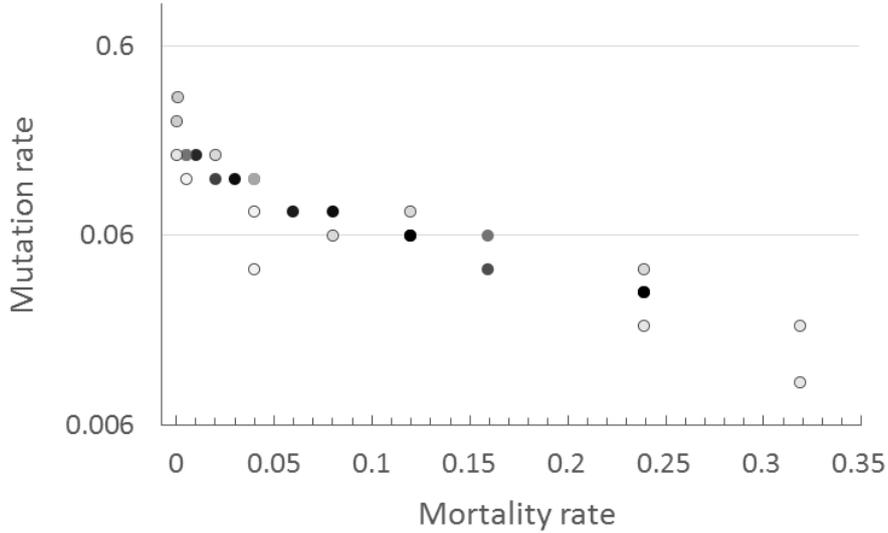


Figure 7: **Relationship between the mutation rate and mortality rate.** Mutation rate is shown in logarithmic scale. Symbols represent number of optimal solutions found for 64-bit H-IFF. Darker colors represent more solutions for those parameters (up to 100% success). Exponential fit for the data: $y = 0.1538 \times e^{-7.28x}$, with $R^2 = 0.89$

used (Figure 9). On 128-bit H-IFF, the maximum achievable fitness value is 448, rather than 192. Out of 20 evolutionary runs, the global maximum on H-IFF, though highly unstable, was found 3 times in different runs when using a mutation rate of 0.03 and a mortality rate of 0.12. It was also found 3 times using a mutation rate of 0.06 and a mortality rate of 0.02, and it was found only once when using a mutation rate of 0.02 and a mortality rate 0.16. The other combinations of mutation and mortality rates that did not find the global maxima used similar values as done for the sweep in Table 2.

Figure 10 displays how single evolutionary runs of the SSGA and AFPO, traverse the 64-bit H-IFF landscape over time. As can be observed in a single evolutionary run of 100,000 generations, the mortality rate implementation seems to exhibit more diversity while remaining at the top of the landscape. In contrast, many individuals in AFPO reside in less-fit (lower) areas of the fitness landscape, likely due to too frequent insertions of random individuals at each generation. The evolutionary progression of the SSGA implementing mortality should convey the informal principle of hill-hugging that was introduced in the previous paper (Veenstra et al., 2018) since, especially compared to AFPO, the steady state implementation does not seem to

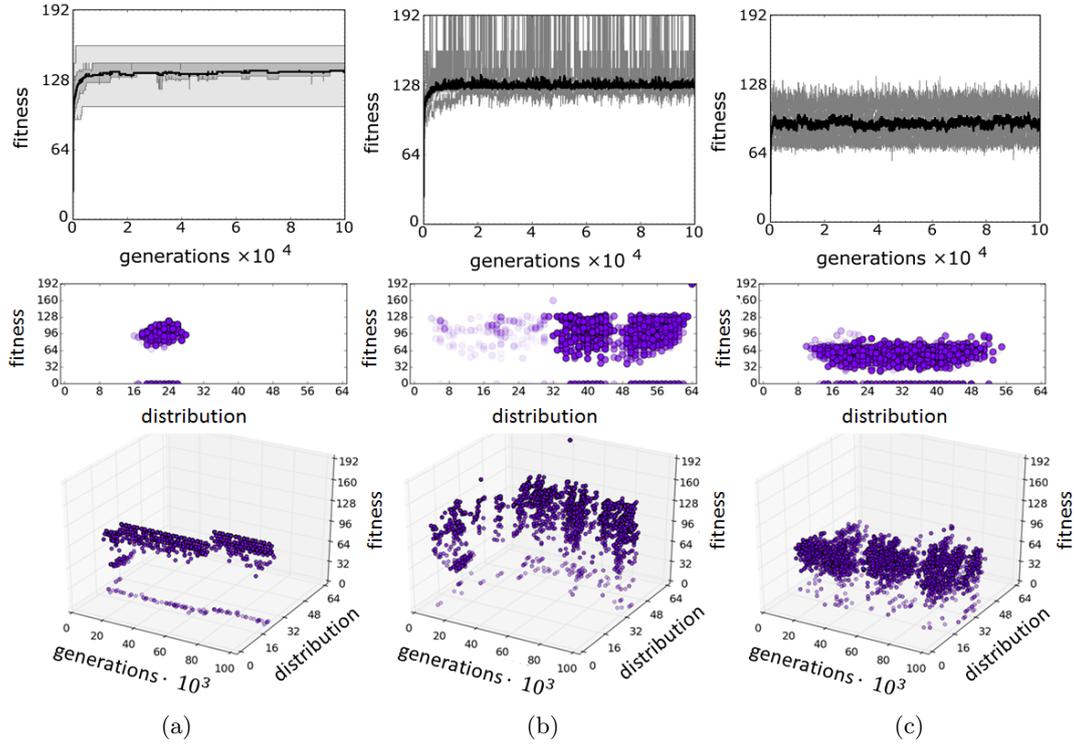


Figure 8: **Evolutionary progress for different mortality rates.** (top) The average fitness and percentiles (25-75 dark grey; 0-100 light grey) of 20 runs using a mutation rate of 0.08 and a mortality rate of 0.04 (a), 0.08 (b) and 0.16 (c). Distribution of the population across the H-IFF landscape of a single run in comparing the distribution and fitness of individuals across the landscape (middle) and plotting the distribution and fitness over generational time (bottom). The blue dots represent individual genomes and the area on H-IFF they occupy at fixed intervals.

occupy low-fit regions of the fitness landscape. The two approaches were furthermore compared over 200 separate evolutionary runs that lasted 100,000 generations (Figure 11). Out of the 200 runs, only one run of AFPO found the global maximum, whereas this number was 190 out of 200 runs in the SSGA with the optimal mortality rate. Age-Fitness Pareto Optimization performs better than SSGAs when no mortality is implemented, but the inclusion of the optimal mortality rate made the SSGA find the global maximum more often than AFPO.

3.2 Spatial Model

Similar to the steady state implementation, the spatial model showed a tight correlation between the mutation rate and mortality rate for the populations that were more efficient at finding the

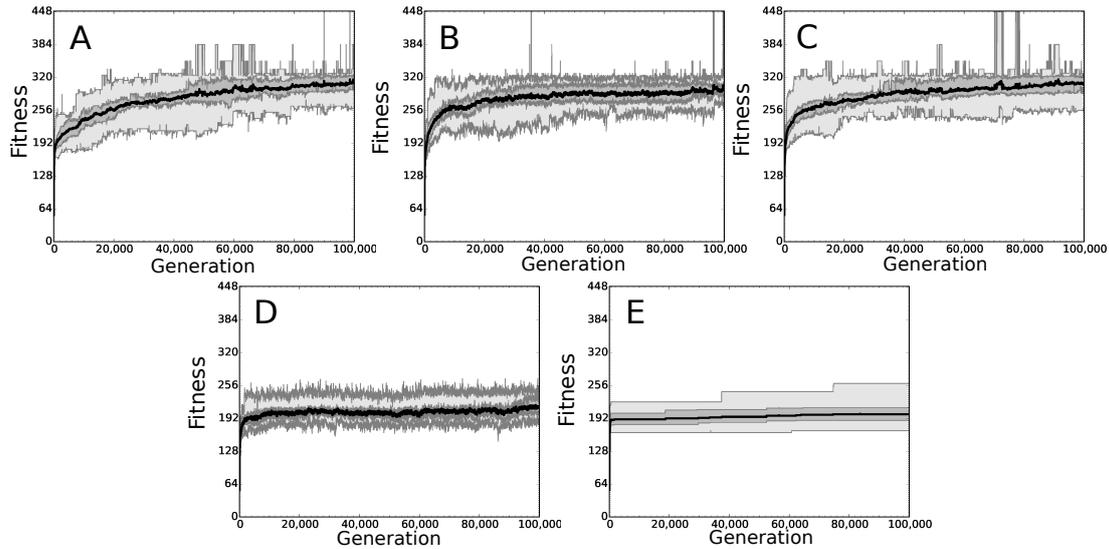


Figure 9: **Relationship between mutation rate and mortality rate on 128-bit H-IFF.** This figure illustrates the evolutionary progression of 20 runs with a mutation rate and mortality rate of 0.02 and 0.16 (A), 0.03 and 0.12 (B), and 0.06 and 0.02 (C), respectively. A relationship with a mutation rate or mortality rate that was too high led to the evolutionary progression depicted in D, while a too low mutation and mortality rate led to the progression depicted in E.

maximum fitness on 32-bit H-IFF, as seen in Table 3. Looking at the individual runs when using a terminal age of 60 (Figure 12), we observe a phenomenon similar as to the SSGA, in which a lower mutation rate leads to premature convergence while a high mutation rate creates unstable populations. Moreover, the speed of finding the global maximum in the spatial model also differs when using different mutation-rate and mortality-rate variables (Figure 13). In this case, the speed was derived from the number of cycles the spatial model ran before finding the global maximum. A different measure of speed would be to count the number of individuals that had been simulated before the maximum had been found. For example, the only real difference to Figure 13 when comparing a mutation rate of 0.02 and a terminal age of 50 with mutation rate 0.16 and terminal age 500, is that the latter needs to simulate significantly less individuals before the maximum is found (two-sided Mann Whitney-u test p value 0.008). The difference in speed in terms of the number of cycles was not significant in this comparison (two-sided Mann Whitney-u test; p value 0.2); i.e., the median number of individuals simulated before finding the global maximum was 971,792 in the low terminal age scenario and 188,148 in the high terminal

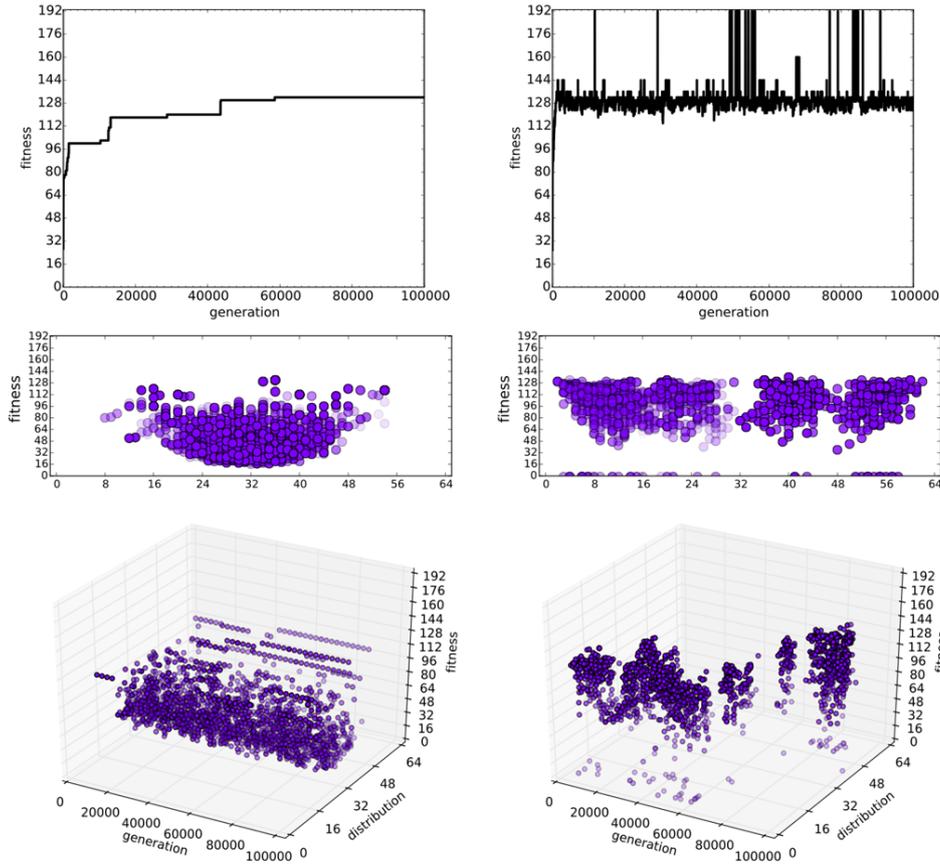


Figure 10: **Single evolutionary runs comparing AFPO (left) and H-IFF (right).** The graphs display the maximum fitness (top), distribution of the population in 2d (middle), and the distribution of the population in 3d (bottom) of a single run. The blue dots represent individuals on the population plotted at intervals.

age scenario. Thus, a higher terminal age in this comparison needed to simulate less individuals. The speed plot of using the number individuals as a measure looks almost identical to Figure 13. For evolutionary algorithms in general, the number of individuals simulated should be minimized, though producing more individuals within a short time might be more valuable for real world populations.

Another interesting result when not using a terminal age is that although the average age of the individuals in the population remains relatively fixed across different mutation rates, the maximum age is significantly higher in high mutation-rate scenarios (two-sided Mann Whitney-u p value of $7 \cdot 10^{-8}$). The difference between the maximum age of the low mutation rate scenario

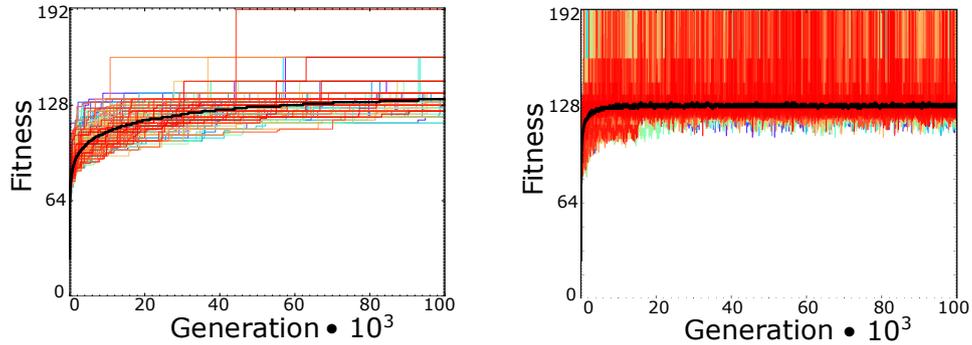


Figure 11: **200 evolutionary runs are shown for both AFPO and mortality rate.** The solid black line represents the average value of the maximum fitness of each evolutionary run. Though it can be seen that the average maximum fitness in AFPO is ever-increasing, the average maximum fitness using the mortality rate stays around the same value. The graph’s color is reddish due to the order of the curves of the evolutionary runs being plotted in the graph.

Table 3: **Number of optimal solutions for 32-bit H-IFF on a spatial model.** Results are taken from 20 runs for each combination of mutation rate (u) and terminal age (TA). ϵ marks combinations where the population went extinct in all runs. Data from terminal age 1000 and 2000 not shown.

$u \setminus TA$	40	50	60	80	120	160	500	-
0.01	5	0	0	0	0	0	0	0
0.015	19	8	0	0	0	0	0	0
0.02	ϵ	20	15	0	0	0	0	0
0.03	ϵ	10	20	20	0	0	0	0
0.04	ϵ	1	11	20	12	0	0	0
0.06	ϵ	ϵ	2	6	20	20	0	0
0.08	ϵ	ϵ	1	1	20	20	0	0
0.12	ϵ	ϵ	0	0	2	9	20	0
0.16	ϵ	ϵ	ϵ	0	0	1	20	17
0.24	ϵ	ϵ	ϵ	0	0	0	1	14
0.32	ϵ	ϵ	ϵ	0	0	0	1	2

(mutation rate 0.03) and the high mutation rate (mutation rate 0.24) is displayed in Figure 15. These results were interpreted as the fittest individuals being unable to produce many functional offspring due to the high mutation rates. This meant that the older, fitter individuals had a higher chance to outcompete the other, likely less fit individuals, in the populations. The elite thus became much older in scenarios where there was a high mutation rate—when looking at no terminal age, the maximum age of the population under high mutation rates could grow as high as 20,000 cycles.

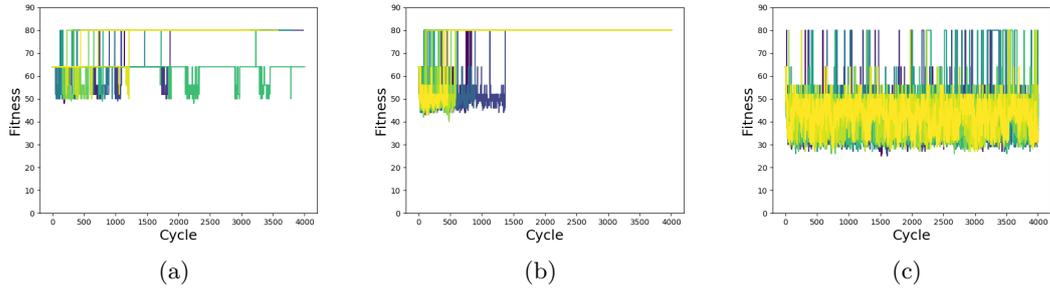


Figure 12: Individual runs showing maximum fitness values across generational time when implementing a terminal age of 60 and a mutation rate of 0.02 (a), 0.03 (b), 0.04 (c). Each cycle represents 100 iterations of the spatial model.

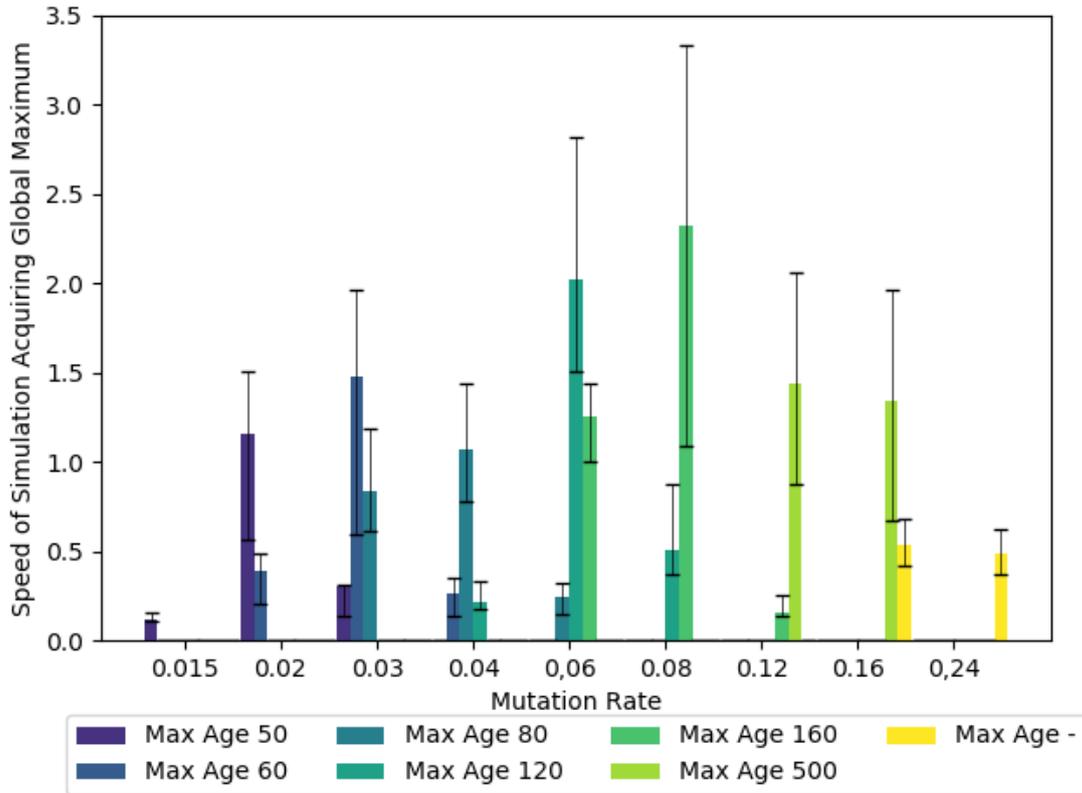


Figure 13: **Speed of solving H-IFF for the spatial model.** Times the global maximum found divided by the average number of iterations $\cdot 10,000$ the spatial model ran varying the mutation rate (x) and the maximum age.

As displayed in Table 3, as high mutation rates also lead to a greater number of less-fit individuals in the population, the mutation rate is necessarily low, otherwise the population

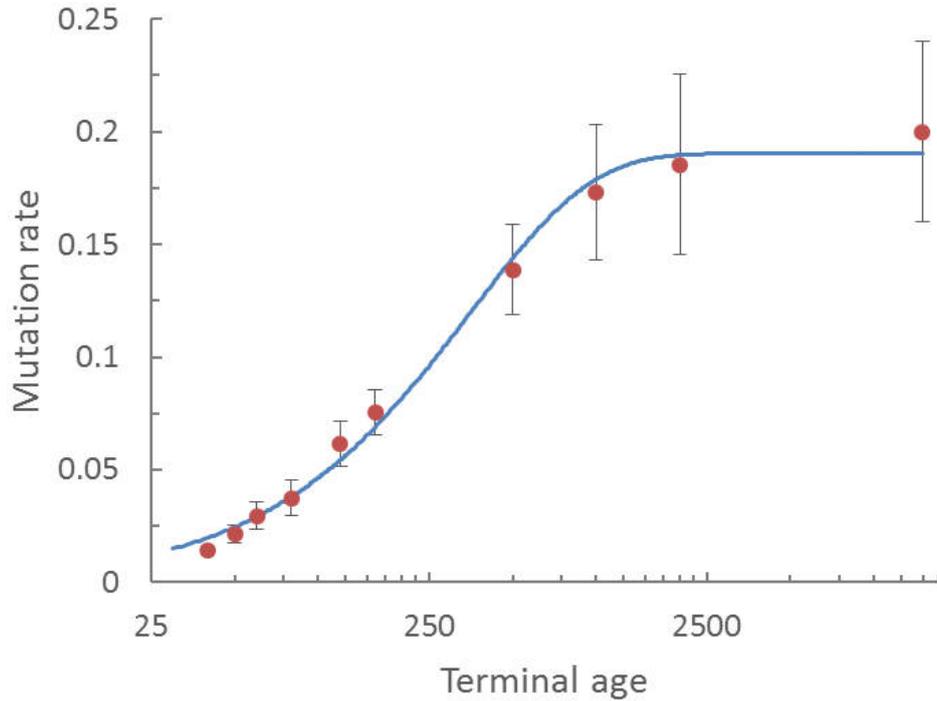


Figure 14: **Optimal mutation rate as a function of terminal age.** Note the logarithmic scale for terminal age. The continuous line shows an exponential fit: $0.1903 - 0.1907 \cdot e^{0.0028}$, with $r^2 = 0.9936$ (values closer to one indicate a better fit).

would go extinct as denoted by ϵ in Table 3. Both a mortality rate and a mutation rate could thus lead to an *error catastrophe*, where a species goes extinct due to excessive mutations. In conclusion, the mortality rate and mutation rate shape the evolvability of the population simulated in the spatial model.

4 Discussion

Our results show that evolvability is greatly influenced by the mutation and mortality rate ratio in both evolutionary algorithms and a spatial model. In particular, the H-IFF function, despite its deceptiveness, can be traversed by an SSGA through simply including an indiscriminate mortality rate. This allows it to find the global optima bit 64-bit H-IFF as well as on 128-bit H-IFF. SSGAs with no chance of removing elite individuals are prone to premature convergence, as demonstrated by the few times the global maximum has been found when no mortality rate

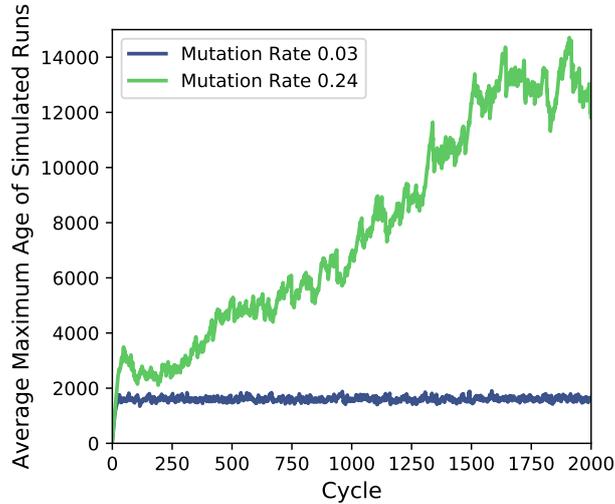


Figure 15: **Average maximum age difference in populations.** The graph shows the average of the maximum ages of 10 individuals in the population in the spatial model over time. The lines display the oldest individuals in the 0.03 mutation rate and the 0.24 mutation rate runs.

was implemented (e.g. Figure 8.a). Likewise, the search approaches random search when both the mortality rate and mutation rate are too high. As in the SSGA, the intrinsic mortality rate also influences the rate at which the global maximum on H-IFF is found in the spatial model. Therefore, evolvability is also enhanced in models where mortality already arises from external mortality factors as a result of local competition. Since a fitness landscape in nature is likely highly convoluted—and possibly deceptive—we speculate that programmed aging could be, as Goldsmith (2014) has mentioned, beneficial for the evolvability of a population. The better a species can traverse the fitness landscape without going through low fitness regions, the more plausible it is that the population will find more adaptive traits. This, in turn, makes the population better and potentially increases its ability to cope with changing environments.

It is relevant to note that the mortality factor we are considering in this paper is very easy to implement—simply remove individuals from the population at random to increase the evolvability of the population. In Lehman (2015), extinction events have been used to increase the evolvability of a population, and it has been shown in Veenstra (2018) that the mortality rate or the extinction events both lead to an algorithm that can have increased evolvability. As long as some sort of

mortality mechanism is implemented, it is likely that the populations evolvability is increased on deceptive fitness landscapes.

A common issue with EAs is that optimal hyperparameters highly depend on the given domain. Generational EAs inherently implement a mortality mechanism, since the entire population is replaced by a new population of offspring every generation (when no elitism is implemented). Moreover, deletion in a steady-state genetic algorithm has also been investigated (e.g., for dynamic environments) and shown to perform similarly to generational EAs (Vavak & Fogarty, 1996). The application of a mortality rate in EAs can therefore also elucidate if such mechanisms should be implemented in an existing EA to better traverse the fitness landscape.

AFPO, developed to combat premature convergence, relies on a new random individual being inserted in a specific region close the maximum fitness of the landscape by accident (Schmidt & Lipson, 2011). From this starting point, the new individuals on the pareto front in this age category need to quickly find a good solution before they are outcompeted by a younger strain climbing another local hill. A potential improvement to AFPO is thus also to insert a new random individual at intervals, not every generation, so the individuals in the new age category have time to climb a local hill, preventing them from being outcompeted by chance. In the SSGA, despite the occasional loss of the best individual in a population, the entire population of individuals remains close to the top of the fitness landscape.

Figure 14 shows that the optimal mutation rate saturates as the terminal age increases. The increment in mutation rate as terminal age increases over 1000 (up to infinity, or no terminal age) is negligible. As described in the mutation accumulation theory by Medawar (1952), this terminal age is so high that natural selection would not be influenced by it, since the average individual probably dies before reaching this age. The 64-bit SSGA simulations frequently lost an optimal solution previously found (Figure 8.b). This happens less frequently in the spatial model (Figure 12), which may be the result of increased robustness of the spatial model as well as an effect of the genome size (which is smaller in the spatial model). In addition, the spatial model has a higher probability of keeping the best solution in the population since a fitter population can sustain more individuals than a less fit population. Their metabolic efficiency is, on average, higher. This increased stability ensures that the global maximum is not lost in the spatial model.

There are a few ways through which we see how an intrinsic mortality rate could arise in

natural populations. A continuous rate of evolvability can enable one population to outcompete other, less evolvable populations. In addition, individuals in a population with an intrinsic mortality rate might keep their optimal mortality rate steady when external factors fluctuate. Finally, a higher mortality rate and lower mutation rate might prevent scenarios such as error catastrophes commonly seen in nature.

The mutation rate and mortality rate have an optimal ratio that depends on multiple factors, such as reproduction and development, crossover, selection pressure, etc. However, we think that the few essential factors that always shape evolvability are mortality, reproduction, and variation. These factors can be altered through different mechanisms:

- **Variation:** Development (Kriegman et al., 2017; Veenstra, 2018), crossover and mutation rates influence the location of the search space of the progeny.
- **Reproduction:** The population size and density, foraging efficiency, reproduction rate, etc. might shape how many new individuals are born.
- **Mortality:** Different factors, extrinsic or intrinsic, dictate the amount of mortality and thereby available space for offspring.

For creating an optimally evolvable population, we think there is a specific amount for each of these factors for a given domain. This relationship between these factors is depicted in the ternary graph of Figure 16. To clarify, without reproduction or variation, populations reside in zero-evolvability states since no new individuals or variation arises, respectively. In addition, without the inclusion of any kind of mortality, the EA likely ends up prematurely converging. A prematurely converged population would then be susceptible to being outcompeted by another population that hasn't prematurely converged. Hence, a selective pressure on intrinsic mortality could emerge in populations to increase evolvability.

In addition, species in natural environments suffer from both intrinsic (aging) and extrinsic (predation, accidents and parasitism) mortality. Extrinsic mortality is known to fluctuate, both in predictable ways (seasons) and depending on external factors (diseases, variable predator pressure). As there is a clear correlation between the mortality and mutation rate for optimal evolvability, this means that such fluctuations in mortality rates could affect the evolvability

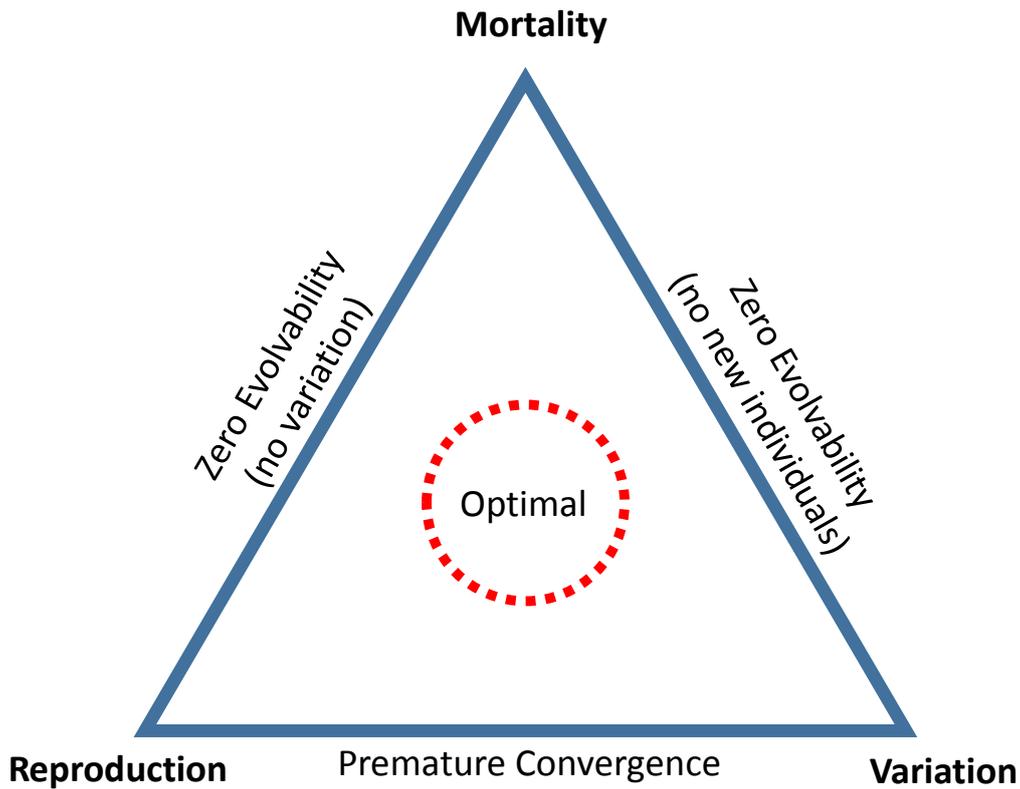


Figure 16: Ternary graph depicting the relationship between reproduction rate/efficiency, mortality, and variation rate. Reproduction, mortality and variation shape evolvability. For a given sequence space, there exists an optimal relationship between these factors (denoted by the red circle). The exact place of the position of the circle likely depends on the domain the evolutionary algorithm is applied to. No mortality usually leads to premature convergence in SSGAs, and no reproduction and no mutations lead to a zero-evolvability state since no new individuals or variation is being introduced in the population.

of populations. Evolving an intrinsic mortality factor may alleviate the problem having an fluctuating evolvability rate. When an external mortality rate is high, aging is not a dominant factor. If external mortality is decreased, then intrinsic mortality prevents the death-rate-to-mutation-rate equilibrium from being out of balance, potentially preserving evolvability in a population. A possible benefit to senescence could thus be to keep the evolvability of a population steady under conditions of fluctuating extrinsic mortality rates. As discussed by Herrera et al. (2016), intrinsic mortality seems to be beneficial in changing environments, and therefore a steady optimal evolvability rate might be required.

An evolving population might benefit from low mutation rates and high mortality rates since

it gradually changes the genetic information in the population's progeny instead of producing drastic changes. In a low mutation rate and a high mortality rate population, the next generation of individuals would be closer to the previous generation, but can be far away in the long run. A high mutation rate and low mortality rate implementation, instead creates a next generation that consists of individuals that are already more distant from their parents compared to the previous population. It is thereby possible that a higher mortality rate prevents the occurrence of error catastrophes that can lead to the extinction of a species.

The results of this paper, in conjunction with the work shown by Mitteldorf & Martins (2014); Herrera et al. (2016), suggest that mortality, or aging, is beneficial to the evolvability of a population, and could therefore be programmed. In addition, we have argued how the increased evolvability leads to a selective pressure that might result in an increase of the rate of mortality. As mentioned in the introduction, mortality reduces the personal fitness of an individual drastically though the inclusive fitness might be increased. Therefore, the death of an individual can altruistically help the evolution of the population.

5 Conclusion

An explicit relationship between the mutation rate and mortality rate for optimal evolvability on a deceptive fitness landscape in both spatial and non-spatial evolutionary models has been presented. These results not only increase our understanding of senescence, but hold potential benefit in applications to evolutionary algorithms and robotics. As an alternative to proposed theories on aging showing how intrinsic mortality is advantageous for altruistic aging, we claim that intrinsic mortality governs evolvability and that it is thereby a potentially favorable trait, ultimately supporting the theories on programmed death. Moreover, in scenarios of fluctuating extrinsic mortality rates, an intrinsic mortality rate would keep the evolvability stable, a potentially relevant evolutionary benefit of intrinsic mortality. Also, considering potentially unrealistically high mutation rates to otherwise grant the same level of evolvability, these high mutation rates in nature would likely lead to an error catastrophe, leading to the extinction of the species altogether. Senescence might therefore be a tool to enhance evolvability in evolving populations.

6 Acknowledgements

This project has received funding from "Flora Robotica", a European Union's Horizon 2020 research and innovation program under the FET grant agreement, no. 640959. Computation/simulation for this paper was supported by the Vermont Advanced Computer Core, University of Vermont.

References

- Altenberg, L. (1994). The Evolution of Evolvability in Genetic Programming. *Advance in Genetic Programming*, (pp. 47–74).
URL <http://dynamics.org/Altenberg/PAPERS/EEGP/>
- Anderson, R. C., Wood, J. B., & Byrne, R. A. (2002). Octopus senescence: The beginning of the end. *Journal of Applied Animal Welfare Science*, 5(4), 275–283.
- Aristotle (1910). Book IV. In *The History of Animals*.
- Carter, A. J. R., & Nguyen, A. Q. (2011). Antagonistic pleiotropy as a widespread mechanism for the maintenance of polymorphic disease alleles. *BMC Medical Genetics*, 12.
- Chuang, C.-Y., & Hsu, W.-L. (2010). Multivariate Multi-model Approach for Globally Multimodal Problems. In *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation*, (pp. 311–318). Portland: ACM.
URL <http://doi.acm.org/10.1145/1830483.1830544>
- Darwin, C. R. (1872). *On the Origin of Species by means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life*. John Murray, 6 ed.
- Goldsmith, T. C. (2008). Aging, evolvability, and the individual benefit requirement; medical implications of aging theory controversies. *Journal of Theoretical Biology*, 252(4), 764–768.
- Goldsmith, T. C. (2014). *The Evolution of Aging*. Crownsville: Azinet Press, 3 ed.
- Goldsmith, T. C. (2016). Aging is programmed! (A response to Kowald-Kirkwood Can aging be programmed? A critical literature review). *Aging Cell*, (p. 7).
- Guarente, L., & Kenyon, C. (2000). Genetic pathways that regulate ageing in model organisms. *Nature*, 408(6809), 255–262.
URL <http://www.nature.com/doi/10.1038/35041700>

- Haubold, B., & Wiehe, T. (2006). *Introduction to Computational Biology: An Evolutionary Approach*. Birkhäuser Basel, 1 ed.
- Herrera, M., Miller, A., & Nishimura, J. (2016). Altruistic aging: The evolutionary dynamics balancing longevity and evolvability. *Mathematical Biosciences and Engineering*, *13*, 455–465.
- Kirkwood, T. B. L. (1977). Evolution of aging. *Nature*, *270*, 301–304.
- Kowald, A., & Kirkwood, T. B. L. (2016). Can aging be programmed? A critical literature review. *Aging Cell*, *15*(6), 986–998.
- Kriegman, S., Cheney, N., Corucci, F., & Bongard, J. C. (2017). A Minimal Developmental Model Can Increase Evolvability in Soft Robots. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO'17*, (pp. 131–138).
- Kiryakov, P., Gomez-Perez, A., Glebov, A., Asbah, N., Bruno, L., Meunier, C., Iouk, T., & Titorenko, V. I. (2016). Empirical verification of evolutionary theories of aging. *Aging*, *8*(10), 2568–2589.
- Lehman, J. (2012). *Evolution through the Search for Novelty*. Ph.D. thesis, University of Central Florida.
URL <http://joellehman.com/lehman-dissertation.pdf>
- Lehman, J. (2015). Enhancing Divergent Search through Extinction Events. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO'15*, (pp. 951–958).
- Lehman, J., Miikkulainen, R., & Sun, G. (2015). Extinction events can accelerate evolution. *PLoS ONE*, *10*(8).
- Lin, K., Hsin, H., Libina, N., & Kenyon, C. (2001). Regulation of the *Caenorhabditis elegans* longevity protein DAF-16 by insulin/IGF-1 and germline signaling. *Nature Genetics*, *28*(2), 139–145.
- Medawar, P. B. (1952). An unsolved problem of biology.
- Mitchell, M., Forrest, S., & Holland, J. H. (1991). The Royal Road for Genetic Algorithms: Fitness Landscapes and GA Performance. *Proceedings of the First European Conference on Artificial Life*, (1), 245–254.
- Mitteldorf, J., & Martins, A. C. R. (2014). Programmed Life Span in the Context of Evolvability. *The American Naturalist*, *184*(3), 289–302.
URL <http://www.journals.uchicago.edu/doi/10.1086/677387>
- Nowak, Martin, A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press.
- Nusbaum, T. J., & Rose, M. R. (1994). Aging in *Drosophila*. *Comparative Biochemistry and Physiology Part A: Physiology*, *109*(1), 33–38.
URL <http://www.sciencedirect.com/science/article/pii/0300962994903093>

- O'Dor, R. K., & Webber, D. M. (1986). The constraints on cephalopods: why squid aren't fish. *Canadian Journal of Zoology*, 64(8), 1591–1605.
 URL <http://www.nrcresearchpress.com/doi/abs/10.1139/z86-241-23.UkR1TyqF9D1>
- Schmidt, M. D., & Lipson, H. (2011). Age-Fitness Pareto Optimization. In *Genetic Programming Theory and Practice VIII*, vol. 8, (pp. 129–146).
- Schuster, P., & Swetina, J. (1988). Stationary mutant distributions and evolutionary optimization. *Bulletin of Mathematical Biology*, 50(6), 635–660.
- Shefferson, R. P., Jones, O. R., & Salguero-Gómez, R. (2017). *The Evolution of Senescence in the Tree of Life*. Cambridge University Press.
- Smith, J. M. (1970). Natural Selection and the Concept of a Protein Space. *Nature*, 225.
 URL <http://dx.doi.org/10.1038/225563a0>
- Vavak, F., & Fogarty, T. (1996). A comparative study of steady state and generational genetic algorithms for use in nonstationary environments. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 1143, (pp. 297–304).
- Veenstra, F. (2018). *The Watchmaker's Guide to Artificial Life: On the Role of Death, Modularity and Physicality in Evolutionary Robotics*. Ph.D. thesis, IT University of Copenhagen.
 URL <https://en.itu.dk/~media/en/research/phd-programme/phd-defences/2018/phd-thesis-final-version-frank-veenstra-pdf.pdf?la=en>
- Veenstra, F., Salaslez, P. G. d. P., Bongard, J., Stoy, K., & Risi, S. (2018). Intrinsic Mortality Governs Evolvability. In *The 2018 Conference on Artificial Life*, (pp. 242–249).
- Wagner, G. P. (1996). Homologues, Natural Kinds and the Evolution of Modularity. *American Zoologist*, 36, 36–43.
- Watson, R., Hornby, G., & Pollack, J. (1998). Modeling building-block interdependency. *Parallel Problem Solving from Nature PPSN V, 1498*, 97–106.
 URL <http://www.springerlink.com/content/t594k1x114971150/>
- Weismann, A. (1889). *Essays upon heredity and kindred biological problems*. Oxford, Clarendon Press.
 URL <http://www.biodiversitylibrary.org/item/23551>
- Werfel, J., Ingber, D. E., & Bar-Yam, Y. (2017). Theory and associated phenomenology for intrinsic mortality arising from natural selection. *PLoS ONE*, 12(3).
- Williams, G. C. (1957). Pleiotropy, natural-selection, and the evolution of senescence. *Evolution*, 11(4), 398–411.

Wodinsky, J. (1977). Hormonal inhibition of feeding and death in octopus: control by optic gland secretion. *Science (New York, N.Y.)*, 198, 948–951.

Yang, J. N. (2013). Viscous populations evolve Altruistic programmed ageing in ability conflict in a changing environment. *Evolutionary Ecology Research*, 15(5), 527–543.

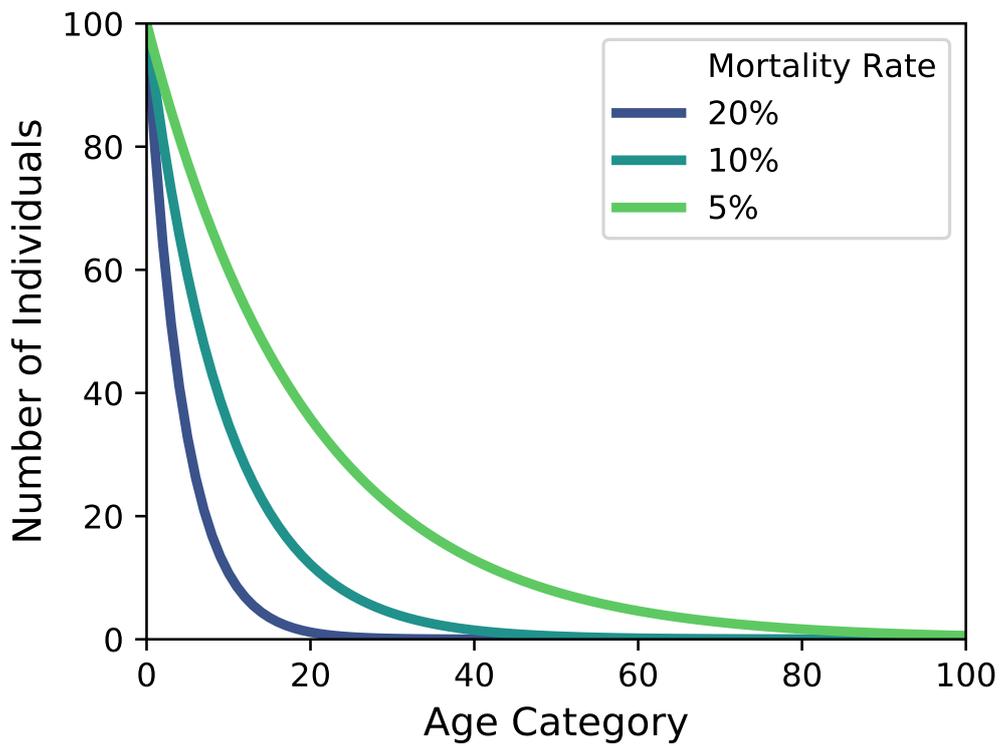


Figure 1: **Survivor curves.** Considering a fixed number of individuals entering the population every iteration, and accepting that there is a 5%, 10% or 20% probability of an individual dying by chance, it can be seen that after a certain number of iterations there will simply not be any survivors left in a specific age category. Based on Medawar (1952).

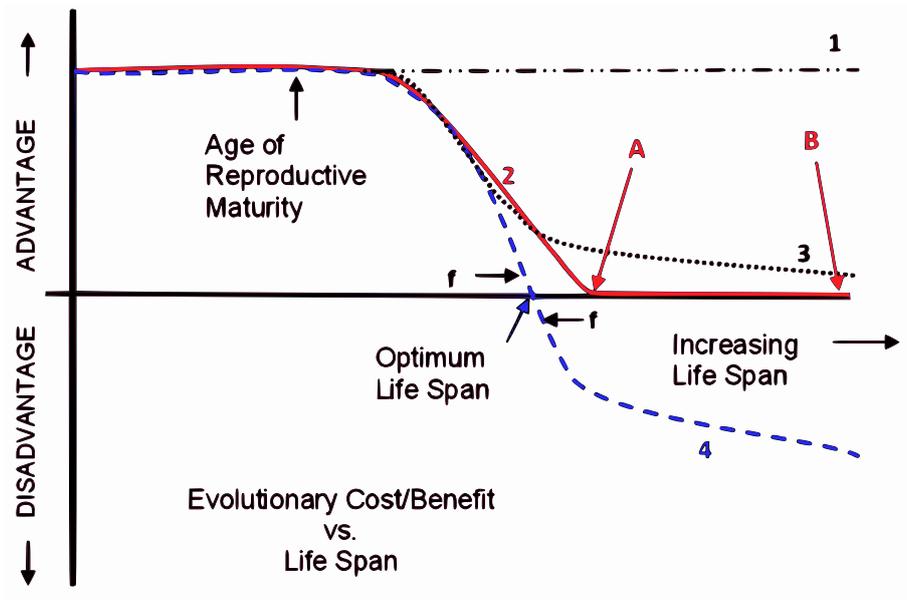


Figure 2: **Theories of aging plotted as an advantage and a disadvantage over time.** Line 1 denotes the advantage of longevity of an individual if longevity would not lead to a decrease in personal or inclusive fitness. Line 2 (solid red line) represents Medawar (1952), where the advantage of longevity would decrease with age but would not have a negative impact. Antagonistic pleiotropy and disposable-soma theories are represented by line 3 (dotted). In this case an increased lifespan does have an advantage to the inclusive fitness of a population but decreases with age. Weismann and Goldsmith support concept 4 where there exists an optimum life-span. Individuals that grow older than this lifespan will have a negative impact on the population. The figure is taken from Goldsmith (2014).

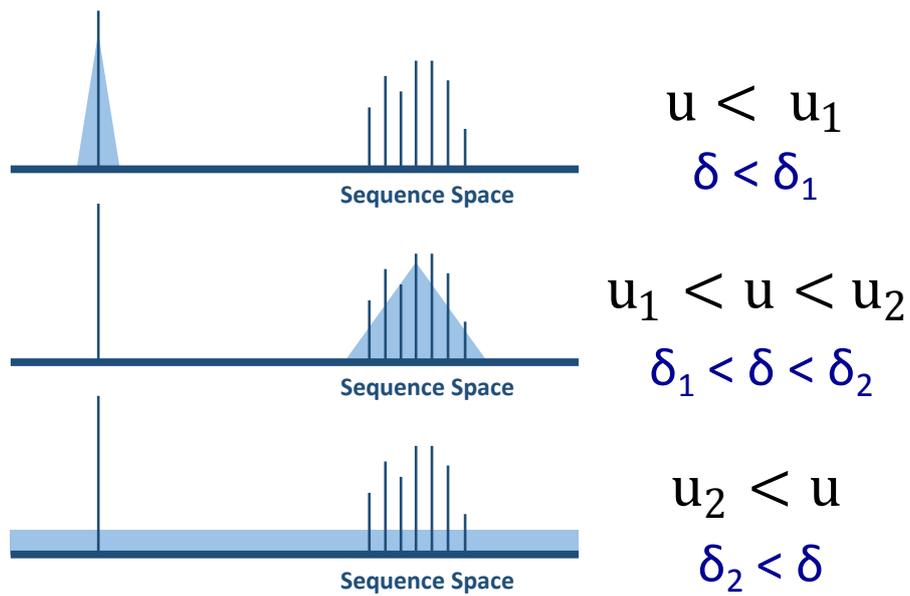


Figure 3: **Fitness landscape with one peak and a hill based on the sequence space.** u represents a mutation rate and δ represents the mortality rate in a given population. The two threshold values for u and δ are believed to have a similar effect on the stable region the population occupies on the fitness landscape.

Table 1: **Example of the score of a 16-bit H-IFF genome.** The table shows the fitness derived from the bit string 0001-1011-1111-1111. In each layer, whenever two sequences of bits above a cell are similar, a value is added to the total score of that layer. After summing all the scores in each layer, the scores are multiplied depending on the layer number. For example, the score of layer two is multiplied by 2, the score of layer three is multiplied by 4 and the score of layer four is multiplied by 8. This genome has a fitness value of 14 (out of a maximum of 32). Adapted from Watson et al. (1998).

	0	0	0	1	1	0	1	1	1	1	1	1	1	1	1	1	sum
1	1		0		0		1		1		1		1		1		$6 \cdot 1$
2	0			0			1			1						$2 \cdot 2$	
3	0						1										$1 \cdot 4$
4	0																$0 \cdot 8$

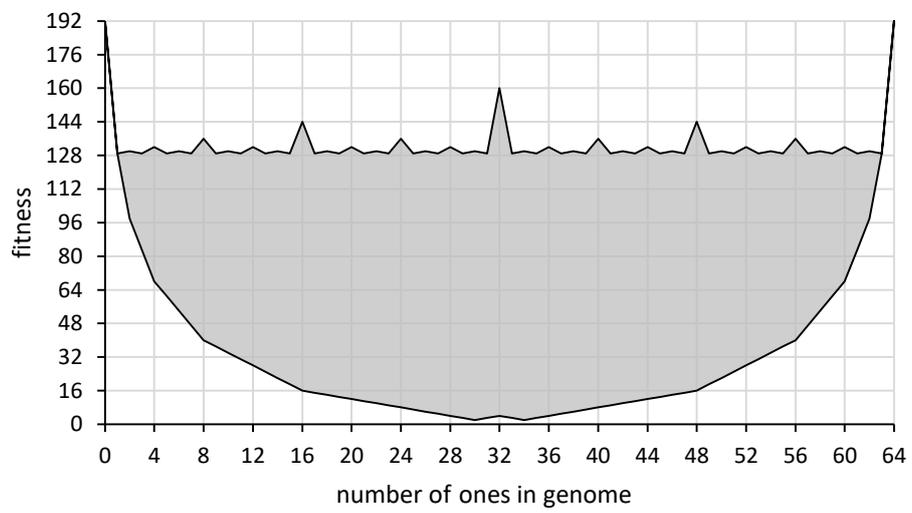


Figure 4: **Fitness landscape of the adjusted H-IFF function.** The global optima in this landscape are at the edges of the distribution and in there are local optima fractally dispersed between every other two local optima.

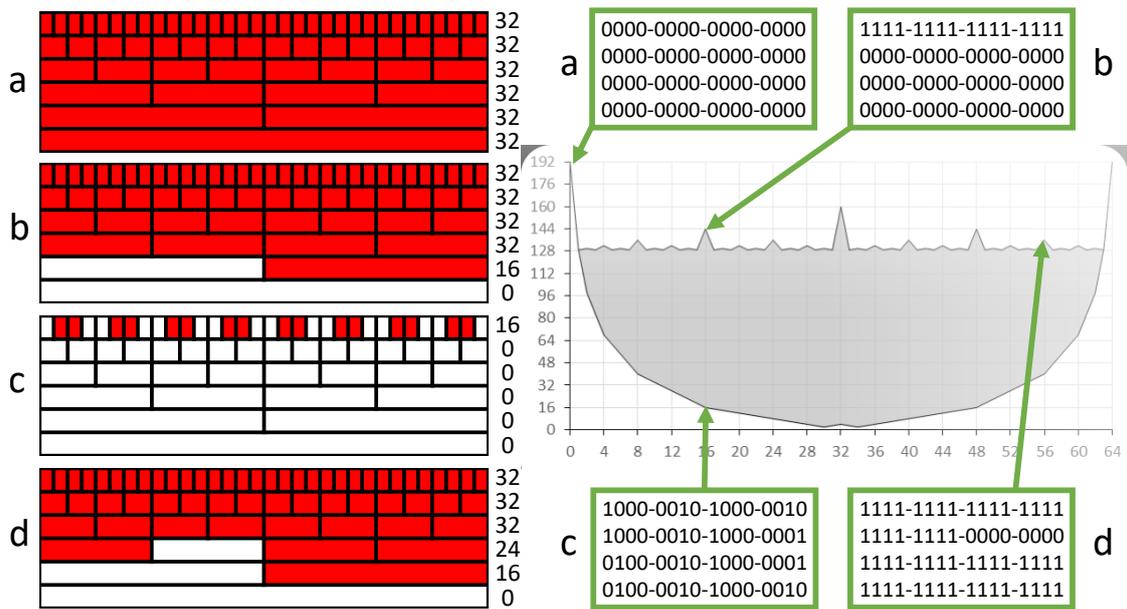


Figure 5: **Explanation of the H-IFF function.** Four genomes of length 64 are shown (a,b,c,d) with their corresponding fitness value. Left shows the scoring tables where red indicates a reward for self-similarity, as shown in Table 1. The red area directly translates into a fitness value of the individuals.

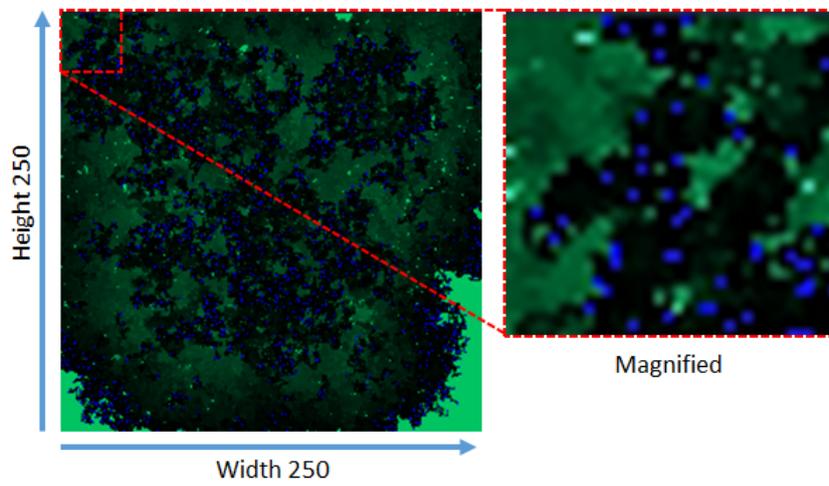


Figure 6: **Illustration of the spatial model.** Green represents plant biomass, blue rabbit biomass. Snapshot taken after the first few cycles of the spatial model

Table 2: **Times the optimal solution is found in the ssga.** Varying the mutation rate (u) and terminal age (δ). Results are taken from 20 runs of each set of parameters on 64-bit H-IFF. The lower-case value represents the average number of generations (thousands) that had to be simulated before finding the global optimum. Mutation rates above 0.32 and below 0.01 have been omitted since the global optima is never found in these scenarios

$u \backslash \delta$	0.0	0.005	0.01	0.02	0.03	0.04	0.06	0.08	0.12	0.16	0.24	0.32
0.01	0	0	0	0	0	0	0	0	0	0	0	2
0.02	0	0	0	0	0	0	0	0	0	0	2	2
0.03	0	0	0	0	0	0	0	0	0	0	20 ₂₆	0
0.04	0	0	0	0	0	1	0	0	0	14 ₄₃	3	0
0.06	0	0	0	0	0	0	0	3	20 ₁₇	11 ₅₀	0	0
0.08	0	0	0	0	0	1	18 ₂₂	19 ₁₉	3	0	0	0
0.12	0	1	0	15 ₁₈	19 ₃₄	7	0	0	0	0	0	0
0.16	2	11 ₃₅	17 ₃₅	3	0	0	0	0	0	0	0	0
0.24	4	0	0	0	0	0	0	0	0	0	0	0
0.32	4	0	0	0	0	0	0	0	0	0	0	0

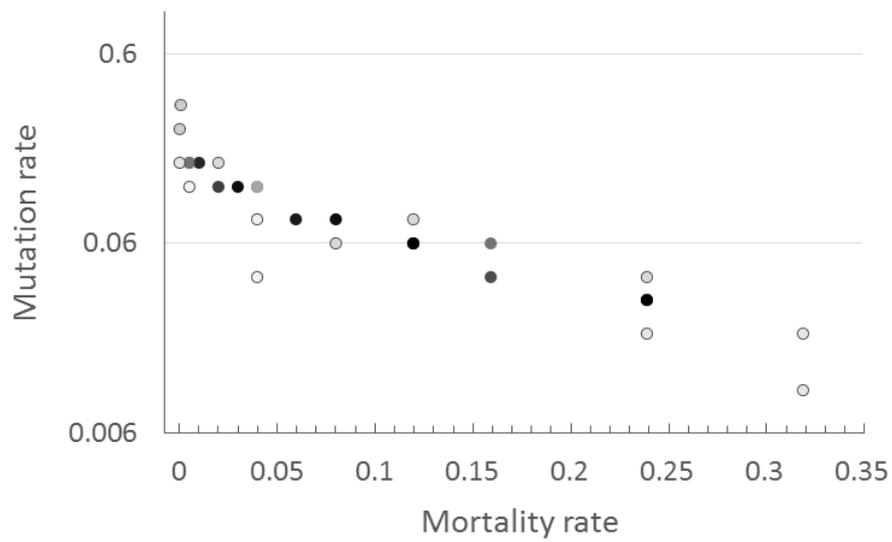


Figure 7: **Relationship between the mutation rate and mortality rate.** Mutation rate is shown in logarithmic scale. Symbols represent number of optimal solutions found for 64-bit H-IFF. Darker colors represent more solutions for those parameters (up to 100% success). Exponential fit for the data: $y = 0.1538 \times e^{-7.28x}$, with $R^2 = 0.89$

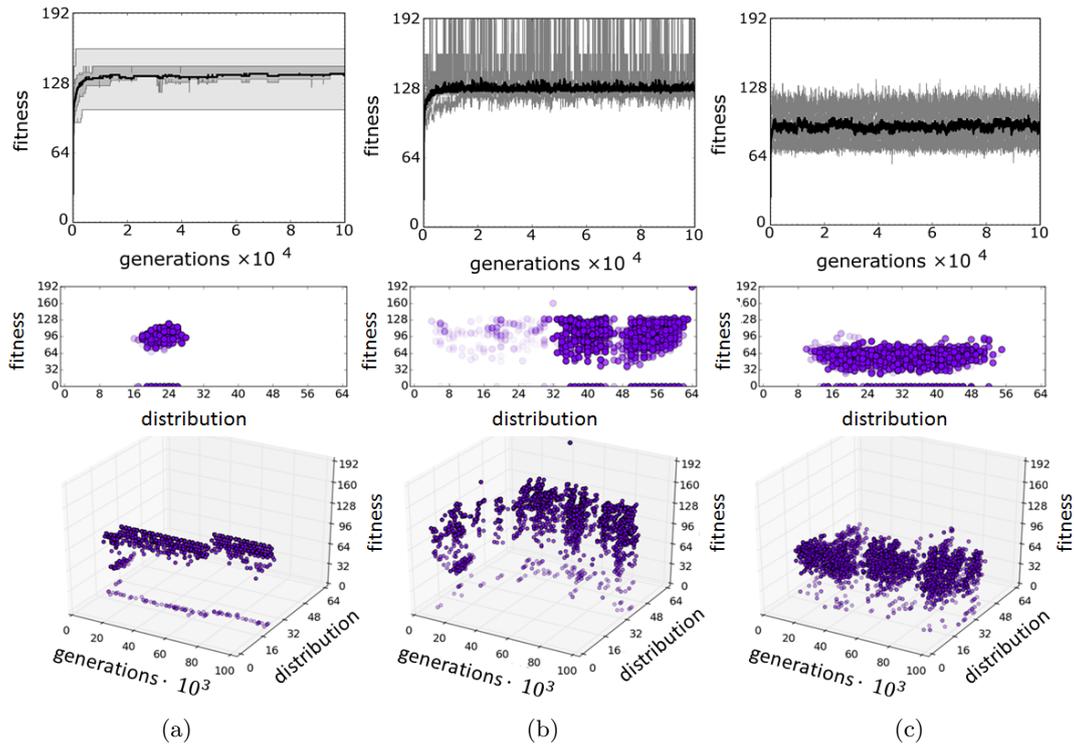


Figure 8: **Evolutionary progress for different mortality rates.** (top) The average fitness and percentiles (25-75 dark grey; 0-100 light grey) of 20 runs using a mutation rate of 0.08 and a mortality rate of 0.04 (a), 0.08 (b) and 0.16 (c). Distribution of the population across the H-IFF landscape of a single run in comparing the distribution and fitness of individuals across the landscape (middle) and plotting the distribution and fitness over generational time (bottom). The blue dots represent individual genomes and the area on H-IFF they occupy at fixed intervals.

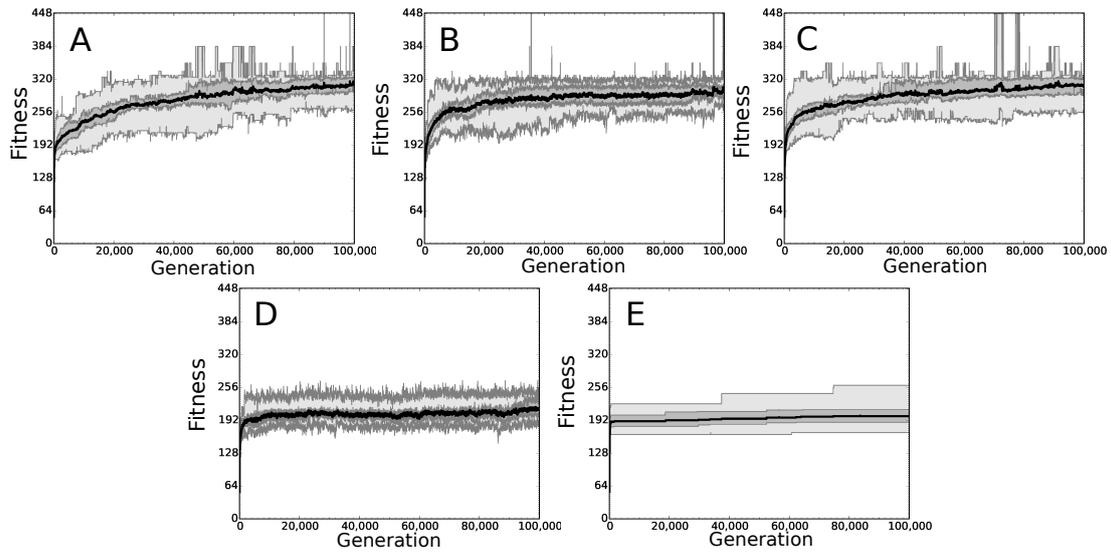


Figure 9: **Relationship between mutation rate and mortality rate on 128-bit H-IFF.** This figure illustrates the evolutionary progression of 20 runs with a mutation rate and mortality rate of 0.02 and 0.16 (A), 0.03 and 0.12 (B), and 0.06 and 0.02 (C), respectively. A relationship with a mutation rate or mortality rate that was too high led to the evolutionary progression depicted in D, while a too low mutation and mortality rate led to the progression depicted in E.

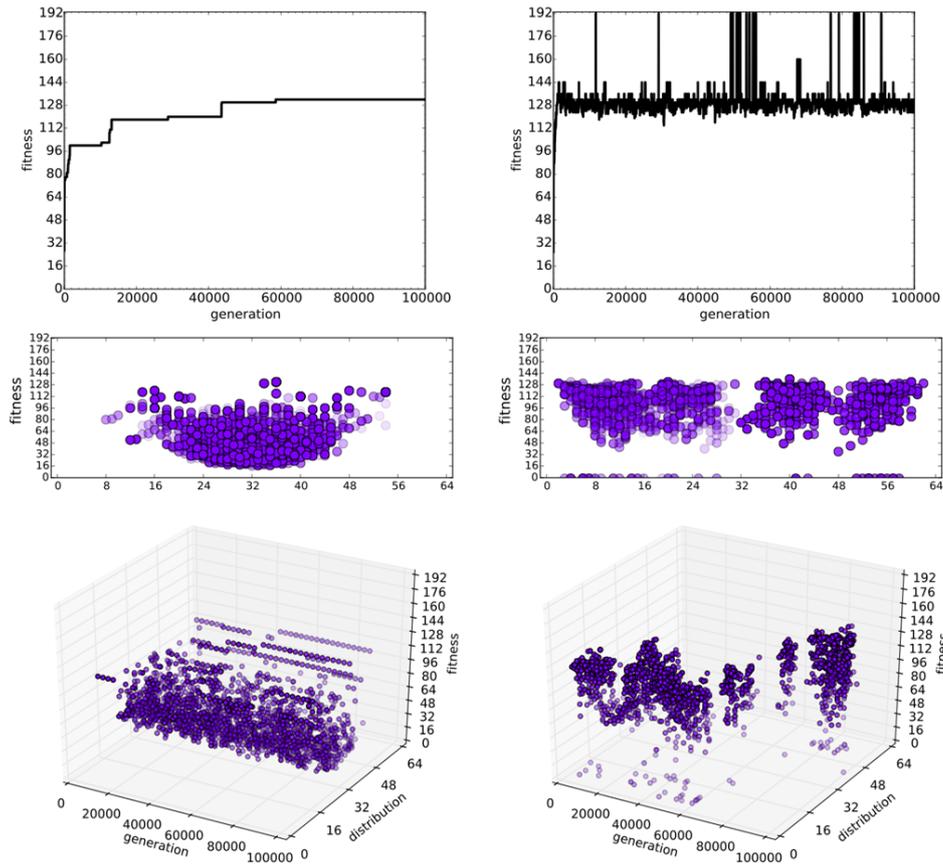


Figure 10: **Single evolutionary runs comparing AFPO (left) and H-IFF (right).** The graphs display the maximum fitness (top), distribution of the population in 2d (middle), and the distribution of the population in 3d (bottom) of a single run. The blue dots represent individuals on the population plotted at intervals.

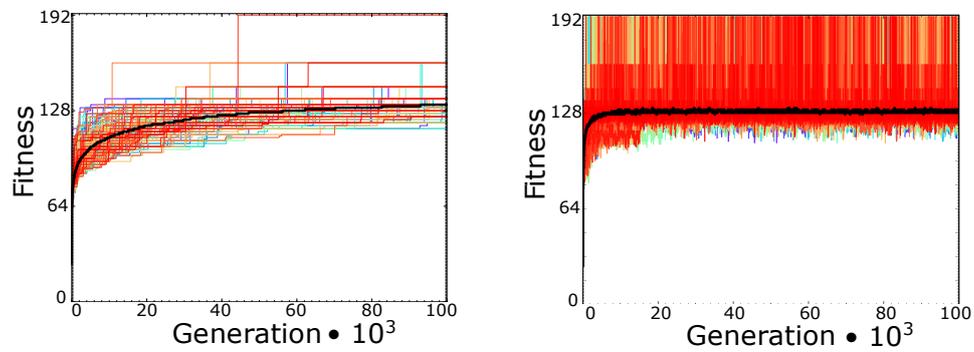


Figure 11: **200 evolutionary runs are shown for both AFPO and mortality rate.** The solid black line represents the average value of the maximum fitness of each evolutionary run. Though it can be seen that the average maximum fitness in AFPO is ever-increasing, the average maximum fitness using the mortality rate stays around the same value. The graph's color is reddish due to the order of the curves of the evolutionary runs being plotted in the graph.

Table 3: **Number of optimal solutions for 32-bit H-IFF on a spatial model.** Results are taken from 20 runs for each combination of mutation rate (u) and terminal age (TA). ϵ marks combinations where the population went extinct in all runs. Data from terminal age 1000 and 2000 not shown.

$u \setminus TA$	40	50	60	80	120	160	500	-
0.01	5	0	0	0	0	0	0	0
0.015	19	8	0	0	0	0	0	0
0.02	ϵ	20	15	0	0	0	0	0
0.03	ϵ	10	20	20	0	0	0	0
0.04	ϵ	1	11	20	12	0	0	0
0.06	ϵ	ϵ	2	6	20	20	0	0
0.08	ϵ	ϵ	1	1	20	20	0	0
0.12	ϵ	ϵ	0	0	2	9	20	0
0.16	ϵ	ϵ	ϵ	0	0	1	20	17
0.24	ϵ	ϵ	ϵ	0	0	0	1	14
0.32	ϵ	ϵ	ϵ	0	0	0	1	2

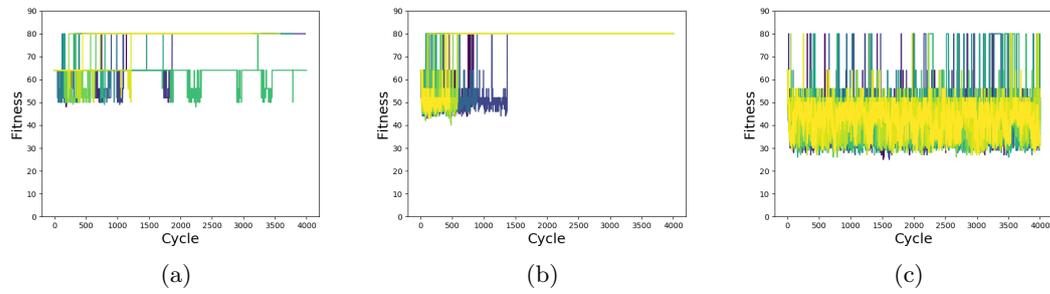


Figure 12: Individual runs showing maximum fitness values across generational time when implementing a terminal age of 60 and a mutation rate of 0.02 (a), 0.03 (b), 0.04 (c). Each cycle represents 100 iterations of the spatial model.

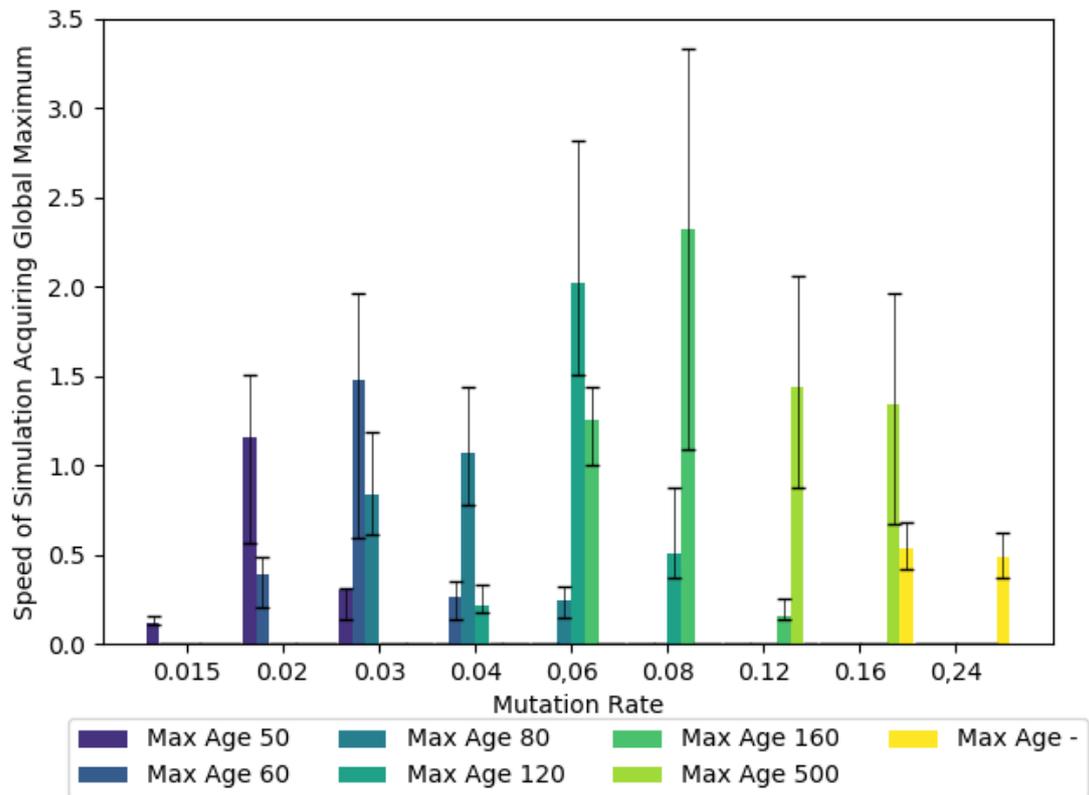


Figure 13: **Speed of solving H-IFF for the spatial model.** Times the global maximum found divided by the average number of iterations $\cdot 10,000$ the spatial model ran varying the mutation rate (x) and the maximum age.

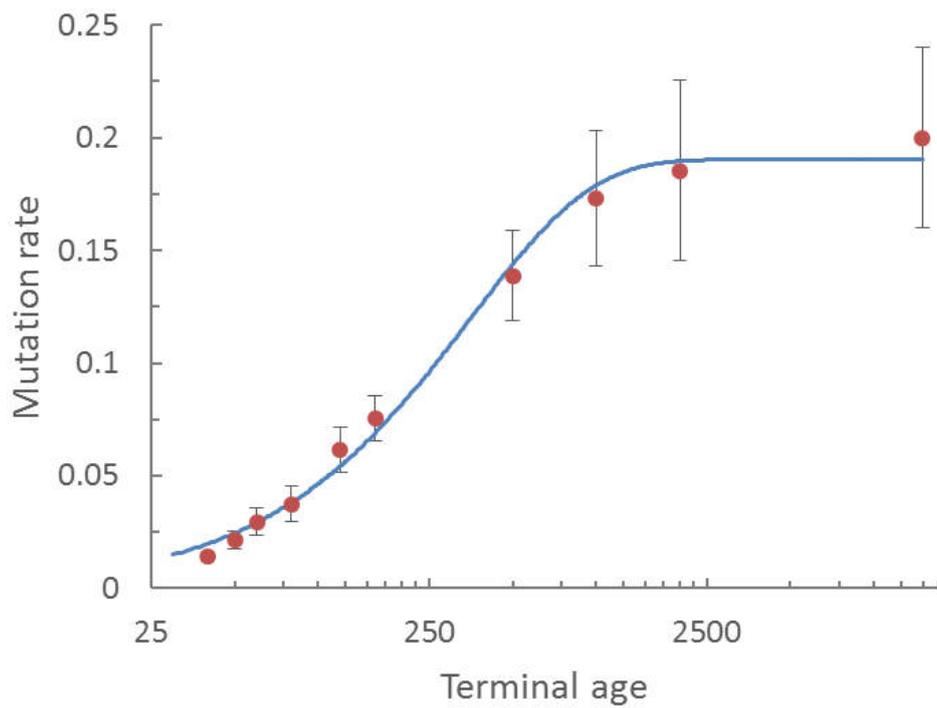


Figure 14: **Optimal mutation rate as a function of terminal age.** Note the logarithmic scale for terminal age. The continuous line shows an exponential fit: $0.1903 - 0.1907 \cdot e^{0.0028}$, with $r^2 = 0.9936$ (values closer to one indicate a better fit).

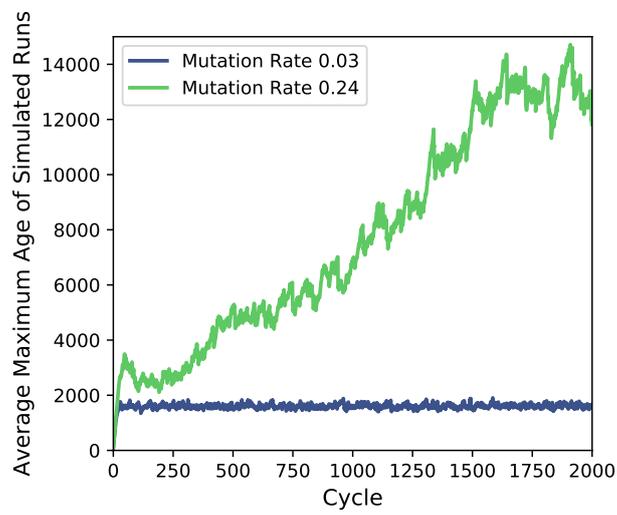


Figure 15: **Average maximum age difference in populations.** The graph shows the average of the maximum ages of 10 individuals in the population in the spatial model over time. The lines display the oldest individuals in the 0.03 mutation rate and the 0.24 mutation rate runs.

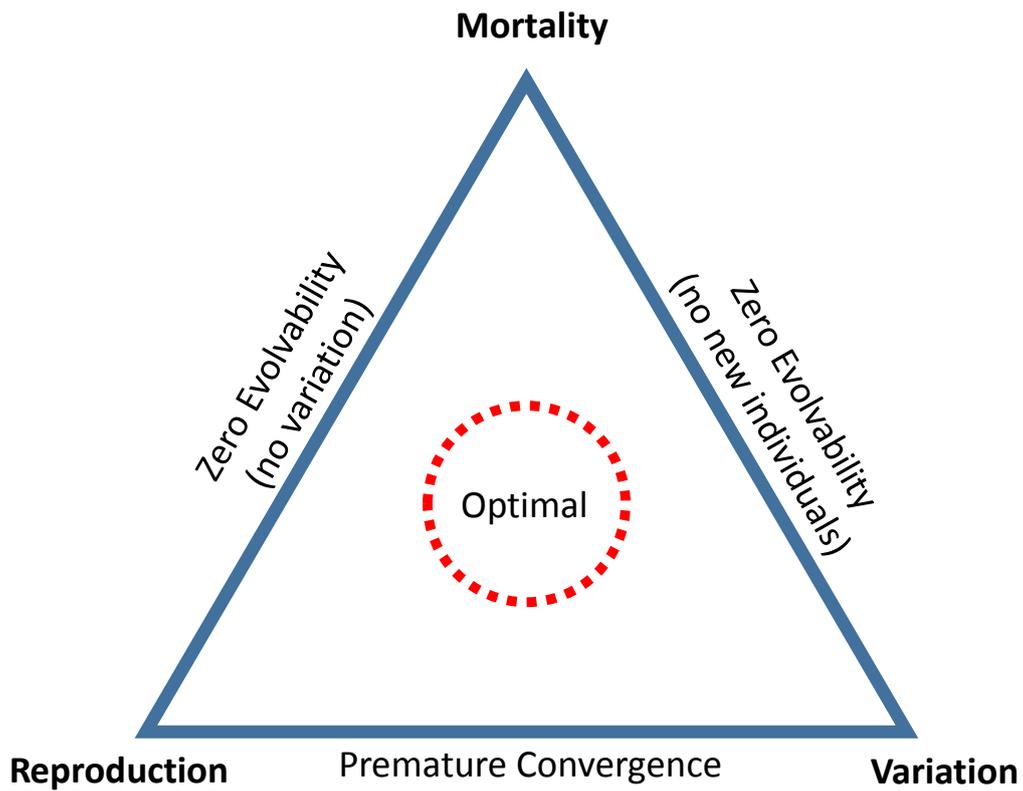


Figure 16: Ternary graph depicting the relationship between reproduction rate/efficiency, mortality, and variation rate. Reproduction, mortality and variation shape evolvability. For a given sequence space, there exists an optimal relationship between these factors (denoted by the red circle). The exact place of the position of the circle likely depends on the domain the evolutionary algorithm is applied to. No mortality usually leads to premature convergence in SSGAs, and no reproduction and no mutations lead to a zero-evolvability state since no new individuals or variation is being introduced in the population.