

Intrinsic Mortality Governs Evolvability

Frank Veenstra¹, Pablo González de Prado Salas¹, Josh Bongard², Kasper Stoy¹ and Sebastian Risi¹

¹ITU, Copenhagen, 2300

²UVM, Burlington, VT 05405

frve@itu.dk

Abstract

An ongoing discussion in biology concerns whether intrinsic mortality, or senescence, is programmed or not. The death (i.e. removal) of an individual solution is an inherent feature in evolutionary algorithms that can potentially explain how intrinsic mortality can be beneficial in natural systems. This paper investigates the relationship between mutation rate and mortality rate with a steady state genetic algorithm that has a specific intrinsic mortality rate. Experiments were performed on a predefined deceptive fitness landscape, the hierarchical if-and-only-if function (H-IFF). To test whether the relationship between mutation and mortality rate holds for more complex systems, an agent-based spatial grid model based on the H-IFF function was also investigated. This paper shows that there is a direct correlation between the *evolvability* of a population and an indiscriminate intrinsic mortality rate to mutation rate ratio. Increased intrinsic mortality or increased mutation rate can cause a random drift that can allow a population to find a global optimum. Thus, mortality in evolutionary algorithms does not only explain evolvability, but might also improve existing algorithms for deceptive/rugged landscapes. Since an intrinsic mortality rate increases the evolvability of our spatial model, we bolster the claim that intrinsic mortality can be beneficial for the evolvability of a population.

Introduction

Mortality is a fundamental component of natural systems that is caused by intrinsic factors (senescence, or deterioration with age) or extrinsic factors such as predation, disease and accident. However, the explanation and origins of senescence are still debated. Despite the plurality of reasons that cause mortality, its rate in a given population can shape selective pressures and thereby overall evolvability. Considering recent publications by Kowald and 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.

Proponents of non-programmed senescence support theories such as mutation accumulation (Medawar, 1952), antagonistic pleiotropy (Williams, 1957), and the disposable-soma theory (Kirkwood, 1977). Proponents of programmed aging largely support theories on evolvability (Weismann,

1889; Goldsmith, 2016) and altruistic aging (Yang, 2013; Werfel et al., 2017; Herrera et al., 2017), while not necessarily excluding the aforementioned theories on non-programmed senescence. The evolvability theory of senescence states that species vary in their capacity of evolution (Goldsmith, 2016). Since a population where individuals don't die would reside in a zero-evolvability state (Goldsmith, 2008), this theory hypothesizes that programmed aging and other life span limiting features can aid the evolvability of a species.

Evolvability is the *population's* ability to traverse the fitness landscape without passing through non-functional regions (Smith, 1970; Haubold and Wiehe, 2006). This definition is different from some existing measures of evolvability (Altenberg, 1994; Wagner, 1996; Lehman, 2012) since it is not directly the ability of a population or individual to produce fitter or more varied individuals when compared to their parents. The ability of the individual or a population to produce better adapted offspring or more diverse individuals is less important than the ability of the population to find a better solution across generational time. A population that can produce better fit individuals in one generation might be unable to cross 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.

Since senescence is prevalent in nature—salmon and octopi dying after spawning of their offspring or elephants dying after running out of their last set of teeth—it is interesting to investigate whether the phenomenon of senescence can be beneficial in simulated models. Intrinsic mortality leading to evolvability has been mostly discussed hypothetically Goldsmith (2014). Notably, Herrera et al. (2017) investigated 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 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 could enable progeny to find a solution in the search space that is more distant, and perhaps ultimately more efficient, than the ancestor's solutions. As a testbed for this potential leap, we used a deceptive fitness function—an adjusted version of the H-IFF function (Watson et al., 1998)—and a genetic algorithm (GA) to simulate a population of individuals.

A common issue with GAs is that optimal hyperparameters highly depend on the given domain. Generational GAs 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 steady state algorithms has also been investigated for e.g. dynamic environments and shown to perform similar to generational EAs (Vavak and Fogarty, 1996). The application of a mortality rate in GAS can therefore also elucidate if such mechanisms should be implemented in an existing GA to better traverse the fitness landscape.

Before discussing the methods illustrating the effect of mortality, it is useful to clarify some concepts from evolutionary dynamics that lead to the premise of the 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 et al. (2006), when considering a sequence space of a specific gene, there can be several optima in this space (Schuster and Swetina, 1988). Depicted in Figure 1, 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 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 lower fit state, how can it traverse the sequence space to end up in the narrow peak that is the better fit solution? Traversing this state space would either require an individual to drastically mutate into that region, or a population could gradually move to the region through genetic drift. Nowak's mutation rate threshold values are however 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 have a chance to also occupy the narrow region no matter how high the mu-




Figure 1: Stable populations on a sequence landscape under changing mutation rates. Blue areas represent the sequence space the population occupies under different mutation rates u . u_1 and u_2 are mutation rate thresholds which make the population stable in narrow high peak (top), broader peak (middle) or unstable (bottom). Height of the black bars represents the fitness value of specific sequences. The distance between bars indicates the genetic distance between specific sequences. Adjusted from Nowak et al. (2006).

tation rate is. Therefore, if we consider δ to be the mortality rate, we claim that there exists a mortality rate threshold δ_1 and δ_2 similar to the mutation rate thresholds.

Accepting this relationship between the mortality rate and the mutation rate, we test whether the theory can be experimentally verified. Using a difficult to solve deceptive state space landscape (H-IFF) on both a steady state GA and a spatial model can elucidate how this relationship influences the evolvability of a population. The steady state GA is used as an abstract model to see the general effects of mortality on the evolutionary progression on this deceptive fitness landscape. Whereas, as described by Werfel et al. (2017), spatial models can elucidate aspects of mortality that equate to natural systems. The spatial model, which contains an inherent extrinsic mortality rate emerging from local competition, is used to isolate the influence of intrinsic mortality to see whether the concept of mortality is relevant for evolvability in natural systems.

Methodology

The experiments are divided in a benchmark optimization implementation using a steady state GA and an agent-based grid model¹. In both simulations, the fitness value of an individual is calculated based on the H-IFF function. The selection/deletion operators in the spatial model are inher-

¹The source code for the steady state GA and the agent-based grid model can be found here: <https://github.com/FrankVeenstra/ALife2018>

ent properties between the interactions of the individuals and their environment while the steady state GA 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 et al. (2006) (Figure 1), and (2) investigate if an additional intrinsic mortality rate influences the evolvability. The experiments investigate if 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.

H-IFF

The H-IFF function generates a deceptive fitness landscape for binary genomes. The fitness landscape is fractal that makes it difficult to solve with any kind of optimization strategy (Watson et al., 1998). Considering a binary genome composed of 64 bits, the H-IFF landscapes where the fitness is given over a function of the number of '1' bits in the genome is depicted in Figure 2. The gray area illustrates the possible fitness values an individual can achieve with a certain number of zeros in its genome. For example, the maximum fitness of an individual with either only zeros or only ones is 192 and this value is the global maxima for this problem. However, when half of the genome is composed of zeros and half ones, the fitness value of that particular individual ranges somewhere between 4 and 160 depending on the specific order of bits. In between each of the peaks lies another local peak which contributes to the fractal nature of the H-IFF function. This implementation is slightly different from the original implementations where the genes may be NULL, 0 or 1 Watson et al. (1998). The NULL possibility has been omitted in order to make the algorithm simpler (e.g. less computational requirements) and easier to visualize. Omitting the NULL possibility reduces the self-similarity calculations by one layer. The maximum fitness of 64-bit H-IFF is 192 instead of 448 as in the original implementation. However, the landscape is unchanged, with as many local optima as the original implementation.

Steady State GA

The steady state GA 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 a half the contingency in mutation events. Thus, a mutation rate of 0.1 means a gene is mutated with 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.




Figure 2: **The H-IFF fitness space landscape.** Potential fitness value (y) values and the number of ones present in the genome(x). All possible solutions can be plotted within the gray area.




Figure 3: **Illustration of the spatial model.** Green represents plant biomass, blue rabbit biomass. Snapshot taken after the first few cycles of the spatial model

Each steady state GA 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. Deleted individuals were marked with a fitness value of -1 but kept in the population and unable to reproduce. The population was logged after each generation. No crossover was implemented to isolate the effect of the solely the mutation rate.

We ran 20 individual 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 since these values lead to early extinc-

ure 4). The evolutionary progressions differ across various mutation rates where mortality rates or mutation rates that are too high lead to too much variation and less fit individuals. If the mortality rate or mutation rate is too low, the population quickly stagnates in a local optimum. The proper ratio of mutation rate and mortality rate leads to a population residing in an unstable local optimum, but still fit enough to traverse the *top* of the fitness landscape and explore multiple peaks. The results of 20 runs using the steady state GA are shown in Figure 5 (top). Figure 5 (middle and bottom) depicts the fitness and diversity on the H-IFF landscape over generational time, which can help to determine how single runs are able to traverse the fitness landscape. Using the optimal mutation rate to mortality rate ratio, the ability of a population to produce adaptive diversity over generational time can create diversity while still hugging the top of the landscape. Hence, informally, we refer to this phenomenon as *hill-hugging* since the GA crosses valleys but does not move to a low region in the search space compared to higher mutation/mortality rates.

The spatial model produces a similar effect to the steady state GA. Additionally, the optimal mortality rate to mutation rate in the spatial model is less prone to lose the global optima once it has been found (Figure 6). This might be because a better fit population is able to sustain more individuals than a lower fit population (Figure 7). A specific ratio of mutation rate to terminal age that is optimal for the population to traverse the state space landscape can also be observed (Figure 8). The individuals of the spatial model, in contrast to the steady state GA, were initialized with individuals residing in the middle of the state space landscape. The results illustrate that despite being in a local optimum furthest away from the global optima, the solution can still be quickly found with the right parameters.

Table 2: 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 4: 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 $R^2 = 0.89$

Individual runs with a terminal age of 60 (Figure 7) show a similar phenomenon as the non-spatial model; a lower mutation rate leads to premature convergence more quickly while a high mutation rate creates an unstable population. Moreover, the speed of finding the global maximum in the spatial model also differs depending on different mutation rate and mortality rate variables (Figure 6). Here speed was derived from the number of cycles the spatial model ran before finding the global maximum. The only difference to Figure 6 comparing mutation rate 0.02 and terminal age 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) while the difference in speed of the number of cycles was not significant (two sided Mann Whitney-u test p value 0.2). The median number of individuals simulated before finding the global maximum was 971,792 and 188,148 respectively. Thus, a higher terminal age in this comparison needed to simulate less individuals. Apart from this anomaly, the speed plot of using individuals as a measure looks almost identical to (Figure 6) and is therefore not included in the paper. For EAS, the number of individuals simulated should be minimized though simulation time in terms of cycles is what matters in natural populations.

Interestingly, when using no terminal age, the average age of the individuals in the population stays relatively the same across different mutation rates while the maximum age is significantly higher in high mutation rate scenarios (two sided Mann Whitney-u p value of $7 \cdot 10^{-8}$). These results could suggest that the best fit individual is unable to produce many functional offspring due to the high mutation rates, meaning that it has in turn a higher chance to outcompete the other individuals in the populations since the competitors are less fit. The elites thus become older in scenarios

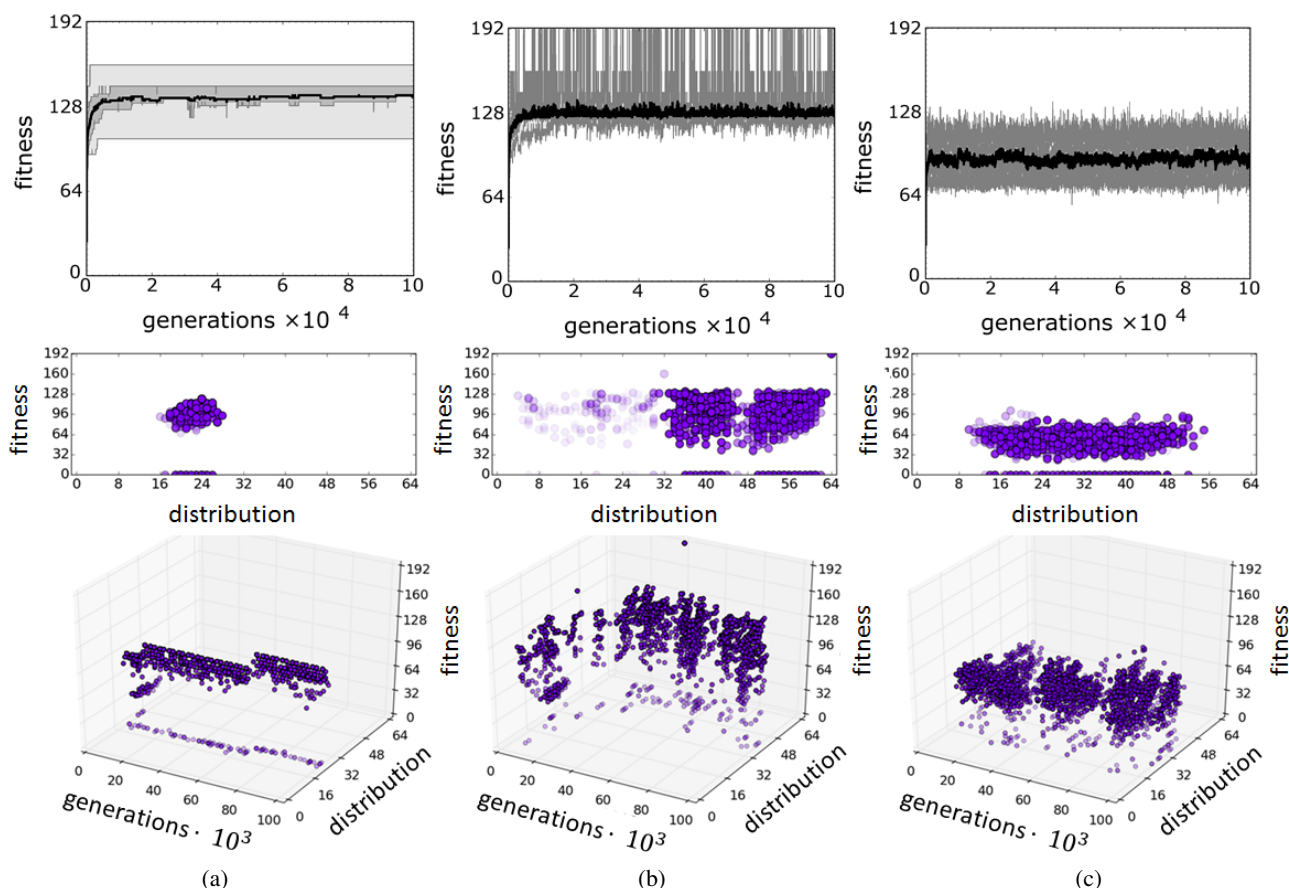


Figure 5: **Evolutionary Progress for different Mortality Rates.** (top) The average fitness and percentiles (25-75 dark grey; 0-100 light grey) of 20 runs using a mutation rate of 0.08 and a mortality rate of 0.04 (a), 0.08 (b) and 0.16 (c). Distribution of the population across the H-IFF landscape 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.

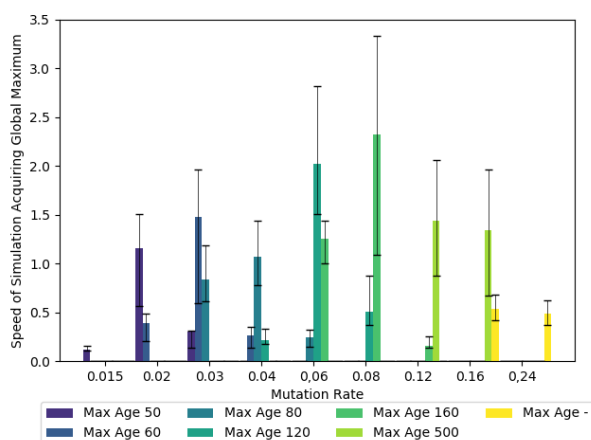


Figure 6: **Speed of solving H-IFF for the spatial model.** Times the global maximum found divided by the average number of iterations $\cdot 10,000$ the spatial model ran varying the mutation rate (x) and the maximum age.

with a high mutation rate; with no terminal age, the maximum age of the population under high mutation rates can grow as high as 20,000 cycles. As shown in Table 2, as high mutation rates also lead to more less fit individuals in the population, the mutation rate is necessarily low, otherwise the population will go extinct as denoted by ϵ in Table 2. This is what happens during an error catastrophe as seen in nature.

Discussion

This paper demonstrated that evolvability is significantly influenced by the mutation and mortality rate ratio in a genetic algorithm and a spatial model using the H-IFF fitness landscape. In particular, the H-IFF function, despite its deceptiveness, can be traversed by a steady state GA by simply including an indiscriminate mortality rate. Since a fitness landscape in nature is likely highly convoluted, we speculate that programmed aging could be, as Goldsmith (2016) men-

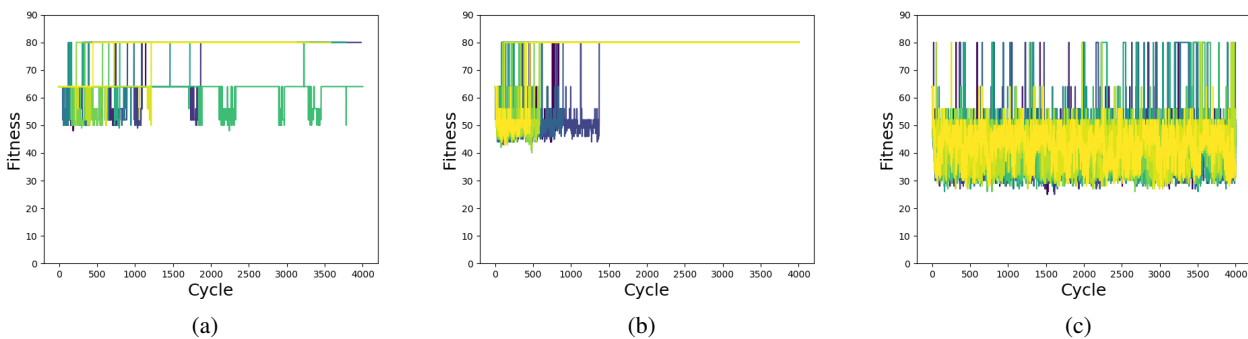


Figure 7: Individual runs showing maximum fitness values across generational time when implementing a terminal age of 60 and a mutation rate of 0.02 (a), 0.03 (b), 0.04 (c). Each cycle represents 100 iterations of the spatial model.




Figure 8: **Optimal mutation rate as a function of terminal age.** Note the logarithmic scale for terminal age. The continuous line shows an exponential fit: $0.1903 - 0.1907 \cdot e^{-0.0028x}$, with $r^2 = 0.9936$ (values closer to one indicate a better fit).

tions, 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 species finds more adaptive traits and the better it is able to cope with changing environments. In our case, the shortest path to the global maximum of H-IFF from the center of H-IFF was to mutate individuals in the appropriate lesser fit local maxima. Our results suggest that there is an optimal mortality rate for a given mutation rate, which could be considered to improve EAS, especially steady state EAS. Steady state algorithms with no chance of removing elite individuals may result in premature convergence. However, the addition of random mortality can enhance such algorithms and allow them to efficiently traverse the state space landscape. The mutation to mortality rate has an optimal ratio that depends on multiple factors. The population size and size of the functional genes are also a contributing factor though experimenting with these variables was out of the scope of this paper.

Figure 8 shows that the optimal mutation rate saturates

as the terminal age increases. The difference between a terminal age of 1000, 2000, or no terminal age, is small. 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, and this can therefore explain why individuals do not reach an older age. Although in the steady state GA the best solution is frequently lost, this is not the case in the spatial model, which indicates a potential difference in robustness of a spatial model compared to the GA but might also be attributed to the genome size used. 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. This increased stability ensures that the global maximum is not lost. Interestingly, in order to get the same amount of evolvability as before finding the global optima in spatial models, the mutation rate, intrinsic mortality rate or extrinsic mortality factors would need to increase again.

Species in natural environments suffer from both intrinsic (aging) and extrinsic death (predation, accidents). Extrinsic mortality is known to fluctuate, both in predictable ways (seasons) and depending on external factors (diseases, variable predation pressure). We have shown a clear correlation between mortality and mutation rate for optimal evolvability, which means that such fluctuations in mortality rates could have a negative impact in the evolvability of populations. Evolving an intrinsic death factor may alleviate this problem: when external pressure is high, aging is not a dominant factor. However, if external mortality is decreased, then intrinsic death prevents the death rate to mutation rate equilibrium to get too far off balance, preserving evolvability. Hence the evolutionary advantage of an intrinsic mortality rate.

Undefined domains with potential deceptive landscapes, such as robotics simulations, may also have an optimal ratio. Though we think it is likely that due to the unknown convolution of the landscape in different regions, this optimal ratio changes in different locations of the landscape and thus changes across generational time. An additional feature

of the mortality to mutation rate ratio is that this determines the broadness of the hills in the landscape that can be occupied across generational time in a stable way. A higher mutation rate or mortality rate would then determine the search space the simulator is occupying, and a broader stable region might result in more abstract general phenotypes that could be better transferable to robots. Robotic simulators, a common domain for EAs, is a promising field for future applications. It would be especially interesting to see if an optimal ratio of mortality and mutation rate can surpass existing algorithms like Age Fitness Pareto Optimization (Schmidt and Lipson, 2011), Age Layered Population Structures (Hornby, 2006) and Novelty Search (Lehman, 2012).

Conclusion

This paper demonstrated an explicit relationship between mutation rate and mortality rate for optimal evolvability on a deceptive fitness landscape in both spatial and non-spatial evolutionary models. As an alternative to proposed theories showing how intrinsic mortality is advantageous for altruistic aging, we claim that intrinsic mortality governs evolvability and that it is thereby a potentially evolvable trait, supporting theories of programmed death. Moreover, in scenarios of fluctuating extrinsic mortality rates, an intrinsic mortality rate would keep the evolvability the same which might further support why intrinsic mortality has an evolutionary benefit. The results not only increase our understanding of senescence but hold potential benefit for deceptive/rugged landscapes in evolutionary algorithms.

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*, pages 47–74.
- 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*. Azinet Press, Crownsville, 3 edition.
- Goldsmith, T. C. (2016). Aging is programmed! (A response to Kowald-Kirkwood Can aging be programmed? A critical literature review). *Aging Cell*, page 7.
- Haubold, B. and Wiehe, T. (2006). *Introduction to Computational Biology: An Evolutionary Approach*. Birkhäuser Basel, 1 edition.
- Herrera, M., Miller, A., and Nishimura, J. (2017). Altruistic aging: The evolutionary dynamics balancing longevity and evolvability. *Mathematical Biosciences and Engineering*, 14(2):455–465.
- Hornby, G. S. (2006). ALPS: The Age-Layered Population Structure for Reducing the Problem of Premature Convergence. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 815–822.
- Kirkwood, T. B. L. (1977). Evolution of aging. *Nature*, 270:301–304.
- Kowald, A. and Kirkwood, T. B. (2016). Can aging be programmed? A critical literature review. *Aging Cell*, 15(6):986–998.
- Lehman, J. (2012). *Evolution through the Search for Novelty*. PhD thesis, University of Central Florida.
- Lehman, J., Miikkulainen, R., and Sun, G. (2015). Extinction events can accelerate evolution. *PLoS ONE*, 10(8).
- Medawar, P. B. (1952). An unsolved problem of biology.
- Nowak, Martin, A., Nowak, M. A., and Nowak, Martin, A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press.
- Schmidt, M. D. and Lipson, H. (2011). Age-Fitness Pareto Optimization. In *Genetic Programming Theory and Practice VIII*, volume 8, pages 129–146.
- Schuster, P. and Swetina, J. (1988). Stationary mutant distributions and evolutionary optimization. *Bulletin of Mathematical Biology*, 50(6):635–660.
- Smith, J. M. (1970). Natural Selection and the Concept of a Protein Space. *Nature*, 225.
- Vavak, F. and Fogarty, T. (1996). Comparison of steady state and generational genetic algorithms for nuse in nonstationary environments. In *Proceedings of IEEE International Conference on Evolutionary Computation*, pages 192–195.
- Wagner, G. P. (1996). Homologues, Natural Kinds and the Evolution of Modularity. *American Zoologist*, 36:36–43.
- Watson, R., Hornby, G., and Pollack, J. (1998). Modeling building-block interdependency. *Parallel Problem Solving from Nature PPSN V*, 1498:97–106.
- Weismann, A. (1889). *Essays upon heredity and kindred biological problems*. Oxford, Clarendon Press.
- Werfel, J., Ingber, D. E., and 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.
- Yang, J. N. (2013). Viscous populations evolve Altruistic programmed ageing in ability conflict in a changing environment. *Evolutionary Ecology Research*, 15(5):527–543.