

How Much Should You Select for Evolvability?

Andrew M. Webb¹, Julia Handl² and Joshua Knowles¹

¹School of Computer Science, The University of Manchester, UK

²Decision and Cognitive Sciences Group, MBS, The University of Manchester, UK
andrew.webb@manchester.ac.uk

Abstract

We consider whether selection for evolvability leads to greater adaptive progress than selection for adaptedness alone. Our treatment bears on longstanding discussions of selection for evolvability in the literature, which have been largely limited to conceptual and qualitative arguments to date. We study a simple mathematical model of a population of individuals whose adaptedness and evolvability (here modelled as the standard deviation of mutations affecting adaptedness) are both under selective forces. In the special case of a population of size two, we show that the optimal amount of selection for evolvability depends on the ratio between the initial evolvability and the amount that evolvability can increase in the time given. Our result shows that to maximize the amount of adaptation it never pays off to select for evolvability more than to select for adaptedness itself. We have not answered the question of to what degree evolvability is selected for in nature, however we have made a small step in quantitative modelling of the evolution of evolvability and proved the existence of conditions under which selection for evolvability has a demonstrably positive effect.

Introduction

The definition of evolvability has been hard to pin down (Kirschner and Gerhart, 1998). One popular definition is that it is a property of an individual, or an aggregate property of a lineage, that increases the expected rate of adaptation of the lineage. For example, for Hansen (2006) evolvability is “the ability of the genetic system to produce and maintain potentially adaptive genetic variants.” All measures of this type of evolvability depend in some way on the probability distribution of fitness effects of mutations (Altenberg, 1995). Proposed measures include the likelihood of a mutation being beneficial (Smith et al., 2002) and the variance of the fitness of offspring prior to selection (Gallagher, 2009).

The evolution of evolvability has become a popular research topic for biologists (Dawkins, 2003; Pigliucci, 2008), researchers of evolutionary computation (Altenberg, 1994; Reisinger and Miikkulainen, 2006), and in the artificial life community (McMullin, 2012; Webb and Knowles, 2014). The questions being asked fall into two broad categories:

1. Has evolvability increased through evolution in life on Earth? If so, by what mechanisms?
2. How can we encourage the evolution of evolvability in evolutionary algorithms and artificial life simulations?

In this paper, we ask *to what extent* we should select for evolvability in simulated evolution in order to maximize the degree of adaptation overall. To our knowledge, this has not been done before. We have in mind a scenario in which adaptedness and evolvability are simultaneously under selective forces, and where the two are inextricably linked such that there is a trade-off; an individual can’t inherit its adaptedness from one parent and its evolvability from another. This would be the case, for example, if an individual’s evolvability were determined by the way in which its traits are encoded in the genetic material that undergoes variation. If one individual is well adapted, but another, less well adapted, individual is more evolvable due to having a more suitable encoding, in general it is not straightforward to re-encode the more adapted individual using the better encoding.

An answer to the question of to what extent we should select for evolvability will primarily be useful in evolutionary computation and artificial life simulations, where often the goal is to maximize the degree of adaptation. It may be less useful in answering questions about natural evolution, where there is no such forward planning, though it still might provide a useful piece in a larger puzzle.

Our contributions are as follows.

- We introduce a simple model of a population in which adaptedness and evolvability are simultaneously under selective forces, and in which we control the relative importance of adaptedness and evolvability during selection.
- For the special case that the population is of size two, we answer exactly to what extent evolvability should be selected for (to the exclusion of selecting for adaptedness) to maximize the total amount of adaptation.
- We discuss the difficulties of using the same method to analyze a larger population.

- We list the ways in which the model might be extended in order to better reflect realistic scenarios.

The Model

We have a population of N individuals, each with two traits A and B . The A value of an individual represents adaptiveness to some environment, and as such it is a value to be maximized. The B value of an individual represents the “evolvability with respect to trait A ”; it is the key factor in determining the rate of increase of A in the course of evolution. Here, that means that the B s determine the standard deviations of mutations affecting the A s.

In each generation, we rank the population by the value $\gamma A + (1 - \gamma)B$, where the A s and B s have first been normalized by dividing by the standard deviations of those traits in the population. The weighting parameter γ is under our control, and takes values in the range $[0, 1]$. This parameter represents a trade-off between selecting for adaptedness and evolvability. The top proportion p become the parents of the next generation; we select with replacement from the set of parents to form the next generation.

The reason we include the normalization step is that, without it, as the variance of one of the traits in the population becomes large, evolution stops acting on the other trait, regardless of the value of γ . This is illustrated in Figs. 1 and 2. These show, for a large population with normally distributed A and B traits, the expected increase in the population mean values of A and B if we select the top half by value $A + B$, where the standard deviation of the A values is 1 and the standard deviation of the B values is parameterized by β . Without normalization, as Fig. 1 shows, the expected increase of each trait is a nonlinear function of both of the trait standard deviations. As the standard deviation of the B s, β , tends to infinity, the expected increase in trait A due to selection tends to zero. With normalization, as Fig. 2 shows, the expected increase of each trait is proportional to the standard deviation of just that trait.

After the selection step, we mutate the A and B values as follows. We add Gaussian noise to each individual’s B value with a constant standard deviation, β . We add Gaussian noise to each individual’s A value with standard deviation αB (i.e., a constant times that individual’s B value). Since standard deviations must be positive, we prevent the B s from taking negative values; when a mutation makes a B value negative, we set it to a small positive value ϵ .

Algorithm 1 shows the process of evolution in this model. Table 1 lists the parameters and their roles.

Our question, stated in terms of the model, is as follows. What value of the parameter γ maximizes, at some particular future time t_{end} , the expected mean value of A ? We answer this question exactly in the special case that the population size $N = 2$ and the proportion selected as the parents of the next generation $p = 1/2$.




Figure 1: Without the normalization step, when we select for $A + B$, the expected increase in the mean value of each trait (ΔA and ΔB) is a nonlinear function of the standard deviations of both traits. As the standard deviation of the B s, β , tends to infinity, trait A stops being selected for.




Figure 2: With the normalization step, when we select for $A + B$, the expected increase in the mean value of each trait (ΔA and ΔB) is a linear function of the standard deviation of just that trait.

Derivation of the Result

We restrict ourselves to the special case that the population size $N = 2$ and the proportion kept during selection $p = 1/2$. We can restate the problem so that we only have to keep track of one A value and one B value in each generation; let $A(t), B(t)$ be the A and B values of the *parent* for generation t , with $A(0) = A_0, B(0) = B_0$.

In each generation, we duplicate the parent, mutate both copies, and then, after normalizing by dividing the A s and B s by their population standard deviations, we select as the parent for the next generation the individual with the maximum value of $\gamma A + (1 - \gamma)B$. Since the parents are identical before mutation, we are essentially selecting between mutation events. The normalization step means that, as far as the selection operator is concerned, all mutations have a standard deviation of 1.

We can achieve the same result by drawing four ‘normalized’ mutations from the standard normal distribution $\mathcal{N}(0, 1)$. M_{A_1} and M_{A_2} are the normalized mutations af-

Algorithm 1 Our model of simultaneous evolution of adaptiveness and evolvability.

```

1: Initialize vector  $A(t=0)$  with  $N$  elements of value  $A_0$ 
2: Initialize vector  $B(t=0)$  with  $N$  elements of value  $B_0$ 
3: for each  $t \leftarrow 0..t_{end}$  do
4:    $A' \leftarrow A(t)/\text{std\_dev}(A(t))$ 
5:    $B' \leftarrow B(t)/\text{std\_dev}(A(t))$ 
6:   Sort  $A(t), B(t)$  by the corresponding value  $\gamma A' + (1 - \gamma)B'$ 
7:   for each  $n \leftarrow 1..N$  do
8:      $i \leftarrow$  random variate drawn from the discrete uniform distribution  $[1, pN]$ 
9:      $M_A \sim \mathcal{N}(0, \alpha^2 B(t)[i]^2)$ 
10:     $M_B \sim \mathcal{N}(0, \beta^2)$ 
11:     $A(t+1)_n \leftarrow A(t)_i + M_A$ 
12:     $B(t+1)_n \leftarrow B(t)_i + M_B$ 
13:    if  $B(t+1)_n < 0$  then
14:       $B(t+1)_n \leftarrow \epsilon$ 

```

fecting the A s, and M_{B_1} and M_{B_2} are the normalized mutations affecting the B s. We will then select the pair of mutations with the largest value of $\gamma M_A + (1 - \gamma)M_B$, and multiply each by the desired mutational standard deviation, undoing the normalization step. We then apply these mutations to the parent to get the A and B values of the parent of the next generation.

Let $\langle M_A^+, M_B^+ \rangle$ be the $\langle M_A, M_B \rangle$ pair with the maximum value of $\gamma M_A + (1 - \gamma)M_B$. That is,

$$\langle M_A^+, M_B^+ \rangle = \langle M_{A_m}, M_{B_m} \rangle, \text{ where} \quad (1)$$

$$m = \underset{i \in \{1,2\}}{\text{argmax}} (\gamma M_{A_i} + (1 - \gamma)M_{B_i}).$$

The two components in the sum in (1) are distributed as

$$\gamma M_{A_{1,2}} \sim \mathcal{N}(0, \gamma^2) \quad (2)$$

$$(1 - \gamma)M_{B_{1,2}} \sim \mathcal{N}(0, (1 - \gamma)^2). \quad (3)$$

Using the results (21), (22) from the appendix, we obtain the expected values of these components in the pair with the maximum sum, which are

$$E[\gamma M_A^+] = \frac{\gamma^2}{\sqrt{\gamma^2 + (1 - \gamma)^2} \sqrt{\pi}} \quad (4)$$

$$E[(1 - \gamma)M_B^+] = \frac{(1 - \gamma)^2}{\sqrt{\gamma^2 + (1 - \gamma)^2} \sqrt{\pi}}. \quad (5)$$

From one generation to the next, A will increase by $\alpha B(t)M_A^+$, and B will increase by βM_B^+ . These are the ‘unnormalized’ mutations, which have the expected values (taking the expectation over the possible values of M_A^+, M_B^+)

N	—	The population size. Here we set $N = 2$.
p	—	The proportion of individuals chosen by truncation selection as parents of the next generation. Here we set $p = 1/2$.
A_0	—	The initial value of the trait A in the population.
B_0	—	The initial value of the trait B in the population. Non-negative.
α	—	A parameter adjusting the standard deviation of A mutations. Non-negative.
β	—	A parameter adjusting the standard deviation of B mutations. Non-negative.
t_{end}	—	The time at which we want to maximize the expected mean value of A , with respect to the parameter γ . Positive integer.
γ	—	A parameter under our control representing a trade-off between selection for the traits A and B . In the range $[0, 1]$.

Table 1: Parameters of the model.

$$E[\alpha B(t)M_A^+] = \frac{\gamma}{\sqrt{\gamma^2 + (1 - \gamma)^2} \sqrt{\pi}} \alpha B(t) \quad (6)$$

$$E[\beta M_B^+] = \frac{1 - \gamma}{\sqrt{\gamma^2 + (1 - \gamma)^2} \sqrt{\pi}} \beta. \quad (7)$$

The nonlinear trade-off between the rates of increase of A and B , determined by the parameter γ , is shown in Fig. 3, and is given by

$$f(\gamma) = \frac{\gamma}{\sqrt{\gamma^2 + (1 - \gamma)^2}}. \quad (8)$$

Since $\alpha B(t)M_A^+$ and βM_B^+ represent the change in A and B from generation t to $t+1$, we have the recurrence relations

$$A(t+1) = A(t) + \alpha B(t)M_A^+(t) \quad (9)$$

$$B(t+1) = B(t) + M_B^+(t), \quad (10)$$

and the expected value of $A(t)$ and $B(t)$ (now taking the expectation over the mutation events in *every* generation) satisfy the recurrence relations

$$E[A(t+1)] = E[A(t)] + f(\gamma) \frac{\alpha E[B(t)]}{\sqrt{\pi}} \quad (11)$$

$$E[B(t+1)] = E[B(t)] + f(1 - \gamma) \frac{\beta}{\sqrt{\pi}}. \quad (12)$$

Solving the second recurrence relation gives the expected value of B at time t , which is




Figure 3: The trade-off between selecting for traits A and B in a population of size 2. The function $f(\gamma)$ gives the expected increase in trait A due to selection, in units of the standard deviation of A mutations. The function $f(1 - \gamma)$ plays the same role for the expected increase in trait B .

$$E[B(t)] = B_0 + f(1 - \gamma) \frac{\beta t}{\sqrt{\pi}}, \quad (13)$$

and solving the first recurrence relation gives the expected value of A at time t , which is

$$E[A(t)] = A_0 + \frac{\alpha t}{4\pi} \left(\frac{4\sqrt{\pi} B_0 \gamma}{\sqrt{\gamma^2 + (1 - \gamma)^2}} - \beta(t - 1) \left(1 + \frac{1}{\gamma^2 + (1 - \gamma)^2} \right) \right), \quad (14)$$

To find a value of γ that maximizes the expected value of A at time t_{end} , we set the derivative of the above expression with respect to γ equal to zero. The solution γ^* in the range $[0, 1]$ that satisfies this equation is

$$\gamma^* = \frac{3}{4} + \frac{1}{4} (z - \sqrt{z + 3\sqrt{z - 1}}), \quad \text{where} \quad (15)$$

$$z = \sqrt{1 + \frac{2}{\pi} \left(\frac{\beta(t_{end} - 1)}{B_0} \right)^2}.$$

The optimal value of the trade-off parameter, γ^* , depends only on B_0, β , and t_{end} , and is an increasing function of $B_0/(\beta(t_{end} - 1))$. Figure 4 shows γ^* as a function of $B_0/(\beta(t_{end} - 1))$, while Fig. 5 shows γ^* as a function of the reciprocal $\beta(t_{end} - 1)/B_0$. Both are shown so that both asymptotes are clear.

That γ^* depends on this quantity makes sense; it is the ratio between the initial evolvability B_0 and βt_{end} , which is related to the amount that evolvability can increase in the course of evolution in the time given. If the initial evolvability is large compared to the capacity to increase evolvability, then it pays off to focus more on increasing the trait A . As




Figure 4: The optimal value of the trade-off parameter γ as an increasing function of $B_0/(\beta(t_{end} - 1))$. As B_0 becomes large, the optimal value asymptotically approaches 1.




Figure 5: The optimal value of the trade-off parameter γ as a decreasing function of $\beta(t_{end} - 1)/B_0$. As β or t_{end} become large, the optimal value asymptotically approaches 1/2.

we look further to the future and t_{end} becomes large, the initial evolvability value has less of an effect and γ^* tends towards 1/2. The optimal value γ^* is never less than 1/2; it never pays off to select more for evolvability than for adapt- edness.

Figure 6 shows the optimal value of γ found by numerical methods for a range of values of B_0, β , and t_{end} . For each setting of the parameters, we plot the value of γ with the highest mean value of A at time t_{end} measured over one hundred thousand trials (the low population size make the outcome noisy). The numerical results closely agree with the answer obtained here, verifying the result¹.

Figure 7 shows, for a particular setting of the parameters, the expected value of A over time for three strategies; setting $\gamma = 1$ (so that only A is selected for), setting $\gamma = 1/2$ (so that we select equally for A and B), and setting $\gamma = \gamma^*$ (the optimal value). It can be seen that the γ^* strategy dominates. Figure 8 shows the same with a different setting of the parameter β . Note that for the γ^* strategy, the plots

¹There is a small discrepancy, because our result does not account for the fact that, after a mutation, we set negative B values to small positive values. This manifests when B_0 is small enough that B is small compared with β in the initial generations.




Figure 6: A comparison between the optimal value of γ obtained by numerical methods and the exact result.

do not show A over time for a particular value of γ ; for each time t , the plot shows the expected value of A when using the (constant) value of γ that maximizes $A(t)$.

Larger Population Sizes

With a larger population size, the model is harder to analyze for two reasons. The first is that the trait variances in the population depend on mutations accumulated over multiple generations; because more than one parent is selected in each generation there will be residual variation from the previous generation. Moreover, this residual (post-selection) variation will not be normally distributed. The result is that the traits A and B will no longer be normally distributed, but will be skewed by an amount depending on the trade-off parameter γ and the proportion kept during selection p . The trait distributions will change over time, approaching an equilibrium shape.

The second problem is that, because more than one parent is selected in each generation, correlation builds up between the A and B values in the population; individuals selected for having high A values are likely to have inherited large B values. The result of this correlation is that there is indirect selection for trait B when selecting for trait A .

Figure 9 shows (with $N = 100, p = 1/2$), as functions of γ , the measured mean increase in traits A and B during selection in generation 10, in units of the mutational standard deviation of A and B , respectively. Figure 10 shows the same in generation 50. The asymmetry is due to indirect selection for trait B , and the mean increase of each trait due to selection changes over time because both the trait distributions and the correlation between the traits are changing over time. Compare these with the stationary (in time) and symmetric functions giving the expected per-generation increase of traits A and B for a population of size two, shown in Fig. 3. Without an exact expression for the expected increase of traits A and B due to selection in a larger population, we cannot deduce the optimal value of the trade-off parameter γ .




Figure 7: The expected value of A over time for three strategies for setting γ . $A_0 = 0, \alpha = 1, B_0 = 1, \beta = 0.1$.




Figure 8: The expected value of A over time for three strategies for setting γ . $A_0 = 0, \alpha = 1, B_0 = 1, \beta = 1$.

Figure 11 shows the optimal value of γ found by numerical methods, with a population size $N = 100$ and $p = 1/2$. The exact result from the population size $N = 2$ case is shown for comparison.

Related Work

Here we review research related to selection for evolvability from the research fields of evolutionary biology, evolutionary computation, and artificial life. There has been much debate about how to define evolvability, and about whether it is a property of individuals or populations. In our work we have followed Conrad (1972), Altenberg (1994), Kirschner and Gerhart (1998), and others in defining evolvability to be a property of individuals, related to their amenability to adaptive evolution, or capacity to produce potentially more well-adapted offspring.

Even amongst those who agree with this definition, there is no consensus about exactly how to measure evolvability. Proposed measures include the likelihood of a beneficial mutation, the expected fitness of offspring after selection, and the expected variance in fitness of offspring prior to selection. Gallagher (2009) gives a summary of these and




Figure 9: The function $f_1(\gamma)$ shows the measured mean increase in trait A in generation 10 in units of the standard deviation of A mutations. The function $g_1(\gamma)$ plays the same role for trait B . The functions are asymmetric; there is indirect selection for B when selecting for A . $N = 100$.




Figure 10: The function $f_2(\gamma)$ shows the measured mean increase in trait A in generation 50 in units of the standard deviation of A mutations. Function $g_2(\gamma)$ plays the same role for trait B . The functions are not the same as those for generation 10, and the functions are asymmetric. Compare with Fig. 3. $N = 100$.

other measures. It is this last measure, the pre-selection variance of offspring, which is closest to ours; in our model, the evolvability of an individual is the standard deviation of mutations affecting the individual's offspring.

In biology, questions have been asked about whether evolvability has evolved in life on Earth, and if it has, whether it evolved by natural selection or by some additional mechanisms (Pigliucci, 2008). The first question is motivated by the fact that evolvability confers a future, rather than present, advantage, and it's not obvious that evolution has or is able to select for evolvability even if doing so would be beneficial in the long term (Kirschner and Gerhart, 1998). Amongst those who believe evolvability has evolved, proposed mechanisms include indirect selection due to correlation between fitness and evolvability (Altenberg, 1994), direct selection for evolvability as an adaptation or as a byproduct of selection for environmental robustness (Hansen, 2006; Visser et al., 2003), higher level




Figure 11: A comparison between the optimal value of γ obtained by numerical methods (with population size $N = 100$) and the exact result (with $N = 2$).

selection between clades and groups (Alberch, 1991), and a process of repeated extinctions and radiations (Dawkins, 2003).

In recent experiments in artificial life, the underlying encodings of self-replicators (i.e., the way in which the replicators are encoded in their heritable genetic information) has been allowed to evolve. The hope is that more evolvable encodings (with respect to the environment) will emerge. Many of these simulations were implemented in Avida and its variants. Avida is an artificial life platform in which assembly-like computer programs self replicate (Ofria and Wilke, 2004).

Baugh and McMullin (2013) and Hasegawa and McMullin (2013) have, respectively, designed replicators for the Tierra and Avida self-replication platforms in which part of the self-replicating program is interpreted as genetic information, and another part is interpreted as a decoding mechanism, that decodes the genetic information. By allowing both to evolve, the way that the replicators are encoded in the genetic portion can change over time.

Egri-Nagy and Nehaniv (2003) have implemented a variant of Avida in which each replicating program has its own, different, instruction set, which itself can evolve. The goal is similar; the way that a replicator's behaviour is encoded can change over time, and more evolvable encodings might be discovered.

Webb and Knowles (2014) aimed to study the differing capacities to evolve evolvability between 'non-self-encoding' and 'self-encoding' replicators. In both cases, each replicator implement a decoder that interprets its genetic information, and the decoder itself can evolve over time. In 'self-encoders', the decoder determines the way in which it *itself* is encoded in the genetic information. The authors concluded that there may have been insufficient selection for evolvability in their simulations to distinguish between the two types of replicator.

In the evolutionary computation field, there is a general concern with evolvability in the sense of mechanisms that

might improve the capacity for adaptive evolution. Placing the degree of mutational variation under evolutionary control has been studied by Eiben et al. (1999), mostly from an empirical perspective, though important theoretical work in this area has also been done (Rudolph, 2001).

Reisinger and Miikkulainen (2006) list some ways in which evolvability has been allowed to evolve in evolutionary algorithms. For example, Ebner et al. (2002) study the evolution of evolvability in neutral networks, in which neighbouring genotypes can encode the same phenotype. Neutral mutations, whose relevance to evolvability was first outlined by Maynard Smith (1970), are those that leave the phenotype unchanged, while possibly changing the phenotypic neighbourhood. Reisinger and Miikkulainen give as other examples evolutionary algorithms with indirect encodings (Stanley and Miikkulainen, 2003) and Estimation-of-Distribution algorithms (Pelikan et al., 2002).

Altenberg (1994) shows that in genetic programming, evolvability can evolve by the implicit selection of blocks of code for what he calls their ‘constructional selection’, or their ability to improve programs in the population when inserted into them.

Conclusion and Future Work

We have introduced a model of a population simultaneously evolving an adaptive trait A and a trait B that is the “evolvability with respect to A ”, with the aim of answering to what extent we should select for evolvability in order to maximize the total amount of adaptation over a given time period.

We have answered this question exactly in the special case that the population size is two, with one individual selected as the parent of the next generation. We find that the optimal weighting parameter γ is never less than $1/2$ and is an increasing function of $B_0/(\beta(t_{end} - 1))$, asymptotically approaching 1, where B_0 is the initial evolvability, β is the standard deviation of mutations affecting B , and t_{end} is the time at which we want to maximize A with respect to γ .

The model is straightforward to analyze with a population size of two, because the two individuals only ever differ in the mutations that happen within the current generation, and the mutations are independent and normally distributed; there are no residual differences between individuals from earlier mutations and there is no correlation between the A and B values within the population. As a result, the expected increase in A due to selection is a function of γ times the standard deviation of A mutations, and the expected increase in B is the same function of $1 - \gamma$ times the standard deviation of B mutations (see Fig. 3). For larger populations, the trait distributions are not normal, change over time, and become correlated, making analysis more difficult.

The applicability of our result is limited by the assumptions of the model, which are as follows.

- “Evolvability” is determined wholly by the B value,

and is independent of time, environmental variables, and where each individual is on the A -landscape.

- The A s and the B s can increase without limit.
- We have perfect knowledge of the B s. In practice, we would likely have to rely on estimates of the evolvability of an individual or lineage, derived from observations of the effects of past mutations.

In future work we aim to extend the model to overcome one or more of these limitations.

Appendix

If we have two random variables A and B both distributed as $\mathcal{N}(0, \sigma^2)$, then the maximum of A and B has the expected value

$$\begin{aligned} E[\max(A, B)] &= \int_{-\infty}^{\infty} a\phi(a) \int_{-\infty}^a \phi(b) db da \quad (16) \\ &+ \int_{-\infty}^{\infty} b\phi(b) \int_{-\infty}^b \phi(a) da db \\ &= 2 \int_{-\infty}^{\infty} a\phi(a) \int_{-\infty}^a \phi(b) db da, \end{aligned}$$

where $\phi(x)$ is the pdf of the distribution. This can be understood as follows. We integrate over the possible values of A , multiplying the probability of getting that value by the probability that the B value is less than it (i.e., the probability that the A is the maximum of the pair). For each possibility we multiply by the value of A to get the expected value. We then do the same thing for the case where the B value is the greater of the pair. Because A and B have the same distributions, these integrals are equal, so we evaluate it once and double the result. Evaluating the integral gives the result

$$E[\max(A, B)] = \frac{\sigma}{\sqrt{\pi}}. \quad (17)$$

In this paper we make use of the following result giving the expected values of the pair of numbers (out of two pairs) that has the maximum sum. Suppose we have four normal random variables distributed as

$$\begin{aligned} A_1 &\sim \mathcal{N}(0, \sigma_A^2), & B_1 &\sim \mathcal{N}(0, \sigma_B^2) \\ A_2 &\sim \mathcal{N}(0, \sigma_A^2), & B_2 &\sim \mathcal{N}(0, \sigma_B^2). \end{aligned} \quad (18)$$

Let $\langle A_m, B_m \rangle$ be the $\langle A, B \rangle$ pair with the maximum sum. That is,

$$m = \operatorname{argmax}_{i \in \{1, 2\}} (A_i + B_i). \quad (19)$$

The expected value of A_m is given by

$$E[A_m] = 2 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} a \phi_A(a) \phi_B(b) \Phi(a+b) da db, \quad (20)$$

$$\text{where } \Phi(a+b) = \int_{-\infty}^{a+b} \phi_{A+B}(c) dc,$$

and $\phi_A(x)$ is the pdf of each of $A_{1,2}$, $\phi_B(x)$ is the pdf of each of $B_{1,2}$, and ϕ_{A+B} is the pdf of each of $A_1 + B_1$ and $A_2 + B_2$.

In words, we integrate over the possible values of A_1 and B_1 , multiplying the joint probability of getting those values by the probability that the sum of the *other* $\langle A, B \rangle$ pair takes a value less than $A_1 + B_1$, and we multiply by the A_1 value to get its expected value. We then integrate over the possible values of A_2 and B_2 (for the case where the sum of the second pair is greater than the sum of the first), which gives the same integral again. Adding the two integrals together gives the expected value of A_m .

Evaluating the above integral gives the value

$$E[A_m] = \frac{\sigma_A^2}{\sqrt{\sigma_A^2 + \sigma_B^2} \sqrt{\pi}}, \quad (21)$$

and by symmetry the expected value of B_m is

$$E[B_m] = \frac{\sigma_B^2}{\sqrt{\sigma_A^2 + \sigma_B^2} \sqrt{\pi}}. \quad (22)$$

Acknowledgement

Thanks to Joe Mellor and Henry Reeve for their help in obtaining the result in the appendix. Thanks also to George Kourtis, Andrew Leeming, and Ubai Sandouk.

References

- Alberch, P. (1991). From genes to phenotype: dynamical systems and evolvability. *Genetica*, 84(1):5–11.
- Altenberg, L. (1994). The evolution of evolvability in genetic programming. In Kinnear, K. E., editor, *Advances in Genetic Programming*, Complex Adaptive Systems, pages 47–74, Cambridge. MIT Press.
- Altenberg, L. (1995). Genome growth and the evolution of the genotype-phenotype map. In *Evolution and Biocomputation*, pages 205–259. Springer.
- Baugh, D. and McMullin, B. (2013). Evolution of G-P mapping in a von Neumann self-reproducer within Tierra. In *Advances in Artificial Life, ECAL*, volume 12, pages 210–217.
- Conrad, M. (1972). The importance of molecular hierarchy in information processing. *Towards a Theoretical Biology*, 4:222.
- Dawkins, R. (2003). The evolution of evolvability. *On Growth, Form and Computers*, pages 239–255.
- Ebner, M., Shackleton, M., and Shipman, R. (2002). How neutral networks influence evolvability. *Complexity*, 7(2):19–33.
- Egri-Nagy, A. and Nehaniv, C. L. (2003). Evolvability of the genotype-phenotype relation in populations of self-replicating digital organisms in a Tierra-like system. In *Advances in Artificial Life*, pages 238–247. Springer.
- Eiben, Á. E., Hinterding, R., and Michalewicz, Z. (1999). Parameter control in evolutionary algorithms. *Evolutionary Computation, IEEE Transactions on*, 3(2):124–141.
- Gallagher, A. (2009). *Evolvability: a formal approach*. PhD thesis, University of Oxford.
- Hansen, T. F. (2006). The evolution of genetic architecture. *Annual Review of Ecology, Evolution, and Systematics*, pages 123–157.
- Hasegawa, T. and McMullin, B. (2013). Exploring the point-mutation space of a von Neumann self-reproducer within the Avida world. In *Advances in Artificial Life, ECAL*, volume 12, pages 316–323.
- Kirschner, M. and Gerhart, J. (1998). Evolvability. *Proceedings of the National Academy of Sciences*, 95(15):8420–8427.
- Maynard Smith, J. (1970). Natural selection and the concept of a protein space. *Nature*, 225:563–64.
- McMullin, B. (2012). Architectures for self-reproduction: Abstractions, realisations and a research program. In *Artificial Life*, volume 13, pages 83–90.
- Ofria, C. and Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial life*, 10(2):191–229.
- Pelikan, M., Goldberg, D. E., and Lobo, F. G. (2002). A survey of optimization by building and using probabilistic models. *Computational Optimization and Applications*, 21(1):5–20.
- Pigliucci, M. (2008). Is evolvability evolvable? *Nature Reviews Genetics*, 9(1):75–82.
- Reisinger, J. and Miikkulainen, R. (2006). Selecting for evolvable representations. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pages 1297–1304. ACM.
- Rudolph, G. (2001). Self-adaptive mutations may lead to premature convergence. *Evolutionary Computation, IEEE Transactions on*, 5(4):410–414.
- Smith, T., Husbands, P., Layzell, P., and O’Shea, M. (2002). Fitness landscapes and evolvability. *Evolutionary Computation*, 10(1):1–34.
- Stanley, K. O. and Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130.
- Visser, J., Hermisson, J., Wagner, G. P., Meyers, L. A., Bagheri-Chaichian, H., Blanchard, J. L., Chao, L., Cheverud, J. M., Elena, S. F., Fontana, W., et al. (2003). Perspective: evolution and detection of genetic robustness. *Evolution*, 57(9):1959–1972.
- Webb, A. M. and Knowles, J. (2014). Studying the evolvability of self-encoding genotype-phenotype maps. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 79–86. MIT Press.