

Lineage selection leads to evolvability at large population sizes

Nathaniel Virgo^{1*}, Eran Agmon² and Chrisantha Fernando³

¹Earth-Life Science Institute (ELSI), Tokyo, Japan

²Department of Biological Sciences, Columbia University, New York City, NY

³Google DeepMind, London, UK

*nathanielvirgo+ecal@gmail.com

Abstract

The biological machinery of evolution can itself be subject to natural selection. Several mechanisms have been proposed through which this can happen. Here we argue that one of these — lineage selection — becomes a strong selective force when the time scale of fixation in the population is comparable to the time scale of adaptation. This implies that lineage selection will be enhanced by anything that slows down fixation; in particular, we expect its effects to be strong when populations are very large and when spatial diffusion is limited. To demonstrate this we construct a simple model of a spatially structured population evolving on a fixed, but infinite, fitness landscape. This landscape consists of a smooth, evolvable path surrounded by rugged local peaks. Our model exhibits an extremely strong dependence on population size: as its size is increased the population evolves along the smooth path, avoiding local peaks, for exponentially longer times. These results suggest that selection for evolvability might become an increasingly important force as we consider larger spatiotemporal scales, and in particular that it might help to explain the evolution of the modern cellular architecture from some previous, less evolvable state.

Introduction

Biological cells implement evolution through some of the most complex molecular machines in existence, including the ribosome. Clearly these machines are the product of evolution, and clearly they are effective in facilitating evolution. Our strong intuition is that they evolved for this very purpose, i.e. that they arose through selection precisely because of their effects on the evolvability of life's ancestors. To evaluate this we must understand not only how selection can act on evolvability, but also under what circumstances selection for evolvability can become a dominant evolutionary force, outweighing other competing selection pressures.

Several mechanisms have been proposed through which the machinery of evolution can be subject to evolutionary forces. Although it's likely that multiple mechanisms play important roles (in combination with each other as well as individually), we focus here on lineage selection, or selection between descendants of ancestors with different evolutionary traits, as illustrated in Box 1.

We make two contributions in this paper: first we argue that spatial structure should greatly increase the amount of time for which lineages can coexist, and thus lineage selection should be much stronger in spatial populations than well mixed ones; secondly we use a simple model to show how larger population sizes can result in greater selection for evolvability.


Our model consists of a spatially embedded asexual population evolving on a fixed, but infinite, fitness landscape. The ruggedness of this landscape varies as a function of a genetic parameter, and hence lineages can vary in the rate at which their fitness increases. We find that in large populations, faster-evolving lineages can out-compete slower-evolving ones, with the result that selection maintains the evolvability parameter close to its optimum value. The effect of population size is very pronounced, with evolvability being maintained for exponentially longer times as population size increases.

These results are important because they suggest that in biology, lineage selection may be an increasingly important force at larger spatial and temporal scales. In particular, at the scale of the evolution of all life on Earth we expect it to be a strong and consistently maintained evolutionary force, perhaps strong enough to explain how the genetic machinery of a modern cell could have evolved by natural selection from earlier, much less sophisticated mechanisms.

Lineage selection and the evolution of evolvability

The idea of lineage selection has been proposed independently multiple times. It was first formulated by Eshel (1973), and then popularised by Dawkins (1989) in the proceedings of the first Artificial Life conference. Dawkins conjectures that the first animal with a segmented body might not have been a particularly fit organism, but that its descendants survived because their flexible body plan made them able to evolve to fit many new environments; he goes on to suggest that selection for evolvability in this sense might be an important force in the biological world. The term “lineage selection” was coined by Aboitiz (1991).

Box 1: What is lineage selection?



The above schematic shows two individuals, A and B, initially coexisting in a population. (The rest of the population is not shown.) At first B and its descendants are less fit (produce fewer offspring on average) than descendants of A. However, the lineage descending from B increases in fitness over time relative to A's lineage and eventually becomes fixed in the population, along with any traits possessed by B that facilitated this change in fitness. Through this mechanism traits can be selected that affect the rate of change of fitness over evolutionary time. In a smaller population, A might have out-competed B before it had a chance to catch up; in order to select for evolvability the lineages must coexist for long enough for the effects of evolution to be felt.

The logic of lineage selection, as illustrated in Box 1, is simple: if one individual has a greater capacity for evolution than another then the fitness of its descendants will tend to increase faster over evolutionary time. Provided that neither lineage goes extinct then after a sufficiently long time there should tend to be more descendants of the more evolvable individual than of the less evolvable one.

Lineage selection occurs whenever the fitness value of a gene changes over time, and the concept can be used to understand kin selection (Akçay and Van Cleve, 2016), evolutionary bet hedging (Cooper and Kaplan, 1982) and the error threshold phenomena in quasispecies theory (Demetrius, 1987) in addition to evolution of evolvability. Related concepts can be used to measure the strength of selection in populations (Leibler and Kussell, 2010). Graves and Weinreich (2017) review further applications of the concept.

Evolution of evolvability has become an important research topic, and many other mechanisms have been discovered through which evolvability can be affected or enhanced by selection. In particular, Altenberg (1995); Wagner and Altenberg (1996) pointed out that much of the issue in evolvability is in the properties of the genotype-phenotype map, i.e. how genetic variation affects variation in pheno-

type. They proposed several mechanisms (not including lineage selection) through which selection could cause the genotype-phenotype map to become more modular, and thus more evolvable. A great deal of subsequent research has focused on this important issue, particularly in the context of genetic regulatory systems and development.

However, this leaves open questions about evolution of evolvability in prokaryotes, or even simpler organisms prior to the evolution of the ribosome. Complex genetic regulatory systems are largely a feature of eukaryotes, and the existence of a genetically modifiable genotype-phenotype map is in itself a complex feature that must have evolved somehow. Our central claim here is that lineage selection provides a simple and robust mechanism through which evolvability can be selected for in *any* evolving population, especially when some simple ecological conditions are met.

As an explanation for evolution of evolvability, lineage selection has received surprisingly little attention from the modelling community. An exception is (Brought, 2005), who investigates its effect on mutation rates in genetic algorithms and its interaction with the Baldwin effect.

The notion of evolvability is a complex one that resists a single definition; see Pigliucci (2008) for a review. Following Wagner and Altenberg (1996) the particular notion of evolvability that we are interested in is in the ability to produce the right kinds of variation to increase in fitness over time. We consider one lineage more evolvable than another if it increases more rapidly in fitness over evolutionary time.

This, we claim, is exactly what lineage selection selects for. As per Box 1, selection ultimately depends not only on the number of offspring over a single generation but on the total number of descendants that an individual leaves many generations later. If we define evolvability as the (expected) rate of increase of fitness then it can only correlate positively to the eventual number of offspring, and hence the selection pressure from lineage selection is always in the direction of increasing evolvability in this sense. (This does not preclude the possibility of other competing pressures that would lead evolvability to decrease.)

In order for evolvability in this sense to be important, there should be some capacity for fitness to increase. If the population is close to a fitness optimum lineage selection still occurs, but it can only select for robustness (in the sense of a decreased probability of mutations that decrease fitness) rather than evolvability. This can lead broad fitness peaks to be selected over higher narrow ones, a phenomenon that has been dubbed "survival of the flattest" (Wilke et al., 2001). To stimulate on-going evolution rather than convergence toward a local optimum, modelling studies of evolution of evolvability typically use a fitness function that changes over time (for example, Earl and Deem, 2004; Crombach and Hogeweg, 2008; Canino-Koning et al., 2016). In our model below we take a different approach, of using an infinite fitness landscape in which there is always (in principle) capac-

ity for further innovation.

The idea of selecting directly for evolvability might seem teleological, and indeed it is sometimes claimed that evolvability cannot be subject to selection for this reason. However, lineage selection is blind to the future; it selects evolvable traits not because they will facilitate evolution in the future, but because they *did* facilitate evolution in the past. In some cases we expect that traits that facilitated evolution in the past will also facilitate it in the future, and these are the cases of greatest interest. This point is made eloquently by Watson and Szathmáry (2016) in relation to a different mechanism for evolution of evolvability: evolution, like all forms of learning, fundamentally relies on things that were true in the past continuing to be so in the future, and generalisation from past to future is not mysterious.

Throughout this section we have referred to lineages of organisms, but of course lineage selection can also apply with another unit of selection at its base level. In the case of sexual populations there is no longer a single line of descent leading to each organism, but we can still consider lineages of genes. The concept of organism lineage reappears at larger evolutionary scales due to speciation.

Lineage selection and absorption time

As Box 1 illustrates, lineage selection for evolvability requires multiple lineages to coexist for long enough for evolution to affect their relative growth rates. If one lineage becomes fixed in the population (i.e. the other lineages become extinct) before the effects of evolution have been felt within the lineages, then there will be no selection pressure for evolvability. The question, then, is under what circumstances multiple lineages will coexist in a population for sufficiently long times.

As pointed out by (Graves and Weinreich, 2017, p. 18), in a well mixed population the time for which lineages coexist scales logarithmically with population size, which puts strong limits on the effectiveness of lineage selection. In this section we argue that this limit is dramatically weakened in spatially structured populations.

To see this, let us imagine a new mutant entering a population. The mutation is neutral (i.e. it does not have an immediate effect on fitness, either positively or negatively), but it may have a positive or negative effect after some time, due to its effect on the evolution of other genes. For selection to act on this there must still be descendants of the mutant present in the population at much later times, and there must also be descendants of other, non-mutant individuals for them to compete against.

The term “absorption time” is used in population genetics for the expected time before a mutant either becomes extinct or becomes fixed in the population. (This is distinct from “fixation time”, which refers to the expected time before the mutant becomes fixed, discounting cases where it is lost from the population.) Thus the absorption time is one

way to quantify the time scale on which multiple lineages can coexist in the same system.

For a well mixed population of size N under neutral evolution, the absorption time for a type that initially occupies a fraction f of the population is

$$\bar{t}(f) = -4N(f \log f + (1 - f) \log(1 - f)),$$

measured on a time scale of generations (Ewens, 2004, p. 160). Setting $f = 1/N$, representing a new mutant that exists as a single individual, this tends towards $4 \log N$ as N increases. This logarithmic scaling means that doubling the population only increases the absorption time by a fixed amount, meaning that populations would have to be very large indeed for the absorption time to be comparable to the time scale of evolution.

Spatial structure can change this picture dramatically. To take a simple example of a spatial population, which we use in our simulation model below, let us imagine that the population exists on a line, with each member having two neighbours. (The ends meet to form a circle.) Reproduction consists of replacing one of an individual’s neighbours with its offspring, so that all interactions are local. In this case the absorption time must be at least of the order N , since the mutant strain will always occupy a contiguous span, which can only spread at a constant rate at its ends. This can be modelled as a random walk with two absorbing boundaries — one for fixation and one for extinction — with the result that the scaling is indeed proportional to N .

Thus in the case of one spatial dimension the time for which lineages coexist depends linearly rather than logarithmically on the population size. When one considers the size of bacterial populations it becomes quite plausible that lineages could coexist for long enough for substantial evolutionary change to take place. Different spatial structures, such as two-dimensional lattices or patchy environments, will give rise to different scaling laws, but in general any kind of spatial structure should improve upon the logarithmic scaling of the well mixed case.

This argument is not limited to spatial structure. Anything that slows down absorption time or otherwise promotes coexistence of lineages should enhance the effects of lineage selection and lead to greater selection pressure for evolvability. In particular, lineage selection should be enhanced by diversifying selection or the formation of multiple niches, the latter of which allows lineages to coexist indefinitely.

We note that spatial structuring, diversifying selection and niche formation are all used in genetic algorithms, usually with the stated purpose of maintaining diversity. Our argument here may provide a refined explanation for their effectiveness: by maintaining diversity not only at the phenotypic level but also at the level of genetic representation, they naturally allow the algorithm to search for more evolvable solutions, in addition to solutions that are simply fitter. Our model below can be seen as an existence proof that this can

happen, though it remains to be shown whether it also commonly occurs in practical situations.

The absorption time alone is presumably not the whole story. The strength of lineage selection must also depend on many other factors, such as the strength of selection, the rate of evolution within each lineage, the rate at which evolvability-affecting mutations arise and are lost to drift, and the variance and higher moments of the absorption time. The precise relationships between these factors must determine the strength of selection for evolvability in ways that have yet to be fully investigated. However, high absorption times are surely necessary in order for lineage selection to be a strong force, and considering the expected absorption time gives us a strong intuition about the kinds of systems in which lineage selection will be most effective. Below we demonstrate this using a simple model, showing that in a spatially structured population the effects of lineage selection depend very strongly on population size.

A simple model

In this section we demonstrate these effects using a simple model, in which a population of organisms evolves with a fixed fitness landscape. This landscape is chosen to be “open-ended” (in the weak sense that fitness can increase indefinitely) while having both evolvable and unevolvable paths to higher-fitness regions. This model is not intended to correspond to any particular biological scenario. Instead it serves two purposes: firstly, as a kind of existence proof, showing that lineage selection can lead to evolution of evolvability, even in the case of a fixed fitness landscape and fixed mutation rate; and secondly, as a testbed for our theoretical predictions regarding the effects of population size.

For the sake of easy visualisation we use a two dimensional fitness landscape. Each organism has a “genome” consisting of two integers, u and v , which can take on any value, positive or negative. Mutations will change u and/or v by one unit in either direction, as detailed below. The fitness function is given by

$$f(u, v) = \beta(v + \alpha|u|\phi(u, v)),$$

where α and β are fixed parameters and $\phi : \mathbb{Z} \times \mathbb{Z} \rightarrow [0, 1]$ is a function that sends every pair of integers to a fixed number between 0 and 1. These are chosen randomly (independently and with a uniform distribution) at the start of the simulation, but fixed thereafter.¹

The parameter α controls how rapidly the ruggedness increases with distance from the $u = 0$ line, and β controls the overall strength of selection. Throughout this paper we use $\alpha = 0.5$ and $\beta = 1.0$.

¹In practice we compute them on the fly using a hash function, as was done previously by Altenberg (1995), for example. This guarantees that each (x, y) pair will always be mapped to the same number while avoiding the need to store previously calculated values. We use the `murmurhash3` library for this purpose.




Figure 1: Part of the infinite landscape used in our model, showing the smooth evolvable path surrounded by rugged unevolvable peaks.

Figure 1 shows part of this infinite landscape. The idea is that the value of f can be increased indefinitely, either by increasing v or by moving u far from zero. However, increasing $|u|$ increases the ruggedness of the landscape, and a population that strays too far from the $u = 0$ line is likely to become stuck on a local peak. Individuals with small values of u therefore have a greater evolvability.

To study dynamics on this landscape we initialise a population of size N and iterate the following two steps. (More details are given below.)

1. Choose two random members of the population, call them i and j .
2. Let $\Delta f = f(u_j, v_j) - f(u_i, v_i)$ be the difference in fitness between the two individuals. With a probability given by $1/(1 + e^{\Delta f})$, replace j with a mutated copy of i . Otherwise, replace i with a mutated copy of j .

Generally we will measure time in generations, by which we mean N iterations of this process.

This can be seen as a simplified genetic algorithm without recombination along the lines of Harvey’s microbial GA (Harvey, 2009), or as a Moran process with fitness value e^f . The reason for using e^f rather than just f is that the probability $1/(1 + e^{\Delta f})$ does not change if both $f(u_i, v_i)$ and $f(u_j, v_j)$ are increased by the same amount, so that the evolutionary dynamics remain the same if the whole population moves along the landscape in the v direction.

To generate a mutated copy of an individual, we change u by one unit (up or down, randomly) with probability μ , and we apply the same rule to v , independently of u . Throughout the paper we set $\mu = 0.1$.

The spatial structure of the population determines the details of step 1 above. In this paper we concentrate on the spatial case, in which individuals are arranged along a line, and we always choose two adjacent neighbours. The new individual always replaces one member of this pair, so that




Figure 2: Maximum value of the fitness function obtained after 10^6 generations (i.e. $10^6 N$ death/reproduction events) as a function of population size in a spatial population, for $\alpha = 0.5$, $\beta = 1.0$, $\mu = 0.1$, showing an exponential dependence on population size. Hollow blue points indicate that the variance in u was still high and the value of f still increasing at the end of the run; red indicates that the population converged on a local peak before the millionth generation. The results of three independent runs are shown for every value of N up to 5000, in steps of 100.

both competition and reproduction are local. The line wraps around at the ends, so the individuals at positions 0 and $N - 1$ are neighbours.

The initial population consists of N individuals for which $u = v = 0$, so that the population starts on the evolvable path; the question is whether (and for how long) evolution will follow this path, rather than veering off to the side and getting stuck on a local peak.

Model results

According to our theoretical predictions above, individuals with a value of u close to 0 should have an evolutionary advantage, even if in the short term they are less fit than other members of the population. However, this advantage should be dependent on the absorption time. Therefore larger populations should perform better on this landscape, especially in the case of a spatially structured population.

Figure 2 shows results from the spatial version of the simulation, showing that the value of f attained after 1,000,000 generations increases exponentially with N . This curve would still be exponential if time were measured in reproduction events rather than generations, since the number of reproductions per generation scales only linearly with N . The exponential scaling in Figure 2 is our main result.

Figure 3 shows that the population size does not affect the rate of increase of f (when time is measured in generations), but it instead affects the probability that the population will get stuck on a local peak. f increases linearly over time, but




Figure 3: Time series for ten independent runs, with population sizes of 1000 (red, dashed) and 2000 (blue), showing that the fitness f increases linearly over time until the population converges onto a local peak. The top plot shows the fitness value and the bottom plot shows the average value of the evolvability parameter u . In each plot the variation in the population is represented by a shaded area, which is drawn from -1 to $+1$ standard deviations around the mean. Changing the population size does not affect the rate at which fitness increases before reaching a local peak, but it does affect the probability of getting stuck on a local peak. The parameters are the same as for Figure 2.

stops increasing when the population leaves the region near $u = 0$ and the variance in u collapses. The probability of such a collapse occurring at any given time appears to depend exponentially (negatively) on the population size. This demonstrates the strong dependence on population size that we expected to see, due to the increased coexistence time of lineages within larger populations.

To understand why this happens we can visualise the evolutionary dynamics during this phase of steady increase in f , which we do in Figure 4. At any given time, the majority of the population exists in subpopulations that are concentrated on local peaks, some distance away from the evolvable path at $u = 0$. However, typically a much smaller number of individuals have small values of $|u|$ and large values of v , and it is the descendants of these that tend to become the future population.

This can be seen by marking the ancestors of the current population, as in figures 4b and 4c. Here we colour each of the past individuals according to how many of its descendants are present in the final population. (If multiple individuals have the same value of u and v , the maximum such score is plotted.) This score is zero for most individuals, since most lineages eventually go extinct; the red points show precisely the lines of descent leading to the final population, with the grey points showing other lineages that branched off from this red evolutionary core, only to

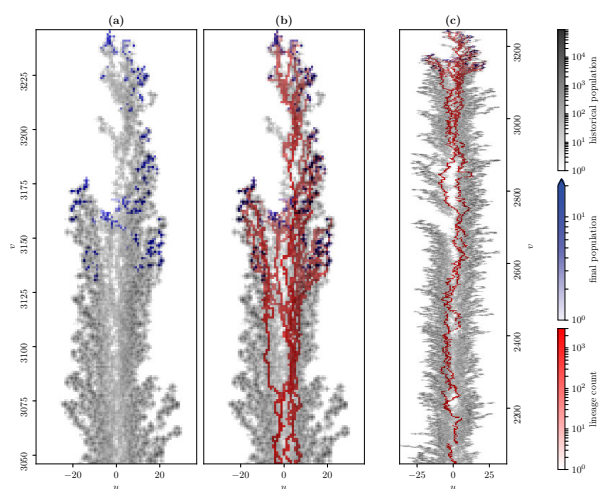


Figure 4: Visualisations of the evolutionary dynamics for a single run, with $N = 3000$. Each plot shows a portion of the fitness landscape, looking down on it from above. (a) shows the population after 5×10^5 iterations in blue. All previous generations are shown in grey, with each cell being coloured according to the total number of individuals that ever existed with those parameters. Logarithmic colourmaps are used because the populations tend to be highly concentrated on local peaks. (b) additionally shows the lineages leading to the final generation, as detailed in the text. (c) is a “zoomed out” version of (a) with a compressed v axis, showing the coalescence of the lineages to a single common ancestor with v around 2200. Note that distant ancestors tend to be much closer to $u = 0$ than typical members of the final population.

eventually go to extinction.

The red lines in Figure 4 are similar to a phylogenetic tree, and exhibit the phenomenon of coalescence (Kingman, 1982), or merging of lines of descent. This is seen in plot 4c, in which the lines of descent meet at a single last common ancestor, with v around 2200.

The difference between the historical (grey) and ancestral (red) populations is notable. At any given time a rather small fraction of the population occupies the region near $u = 0$, with the majority being concentrated on local peaks away from the evolvable path. However, the ancestors tend not to be typical members of the historical populations. Instead they are drawn from the small fraction that occupy the more evolvable region of the landscape. This phenomenon — that current individuals tend to be descended from the more evolvable members of historical populations — is a consequence of lineage selection. This history-oriented perspective provides us with a way to reason about lineage selection in biology, as discussed at the end of the paper.

The effect of lineage selection can be seen more directly in Figure 5. In population genetics, fitness is usually defined in terms of the number of offspring an individual leaves in




Figure 5: The average number of descendants an individual leaves after a given number of generations, plotted as a function of u . Individuals near the evolvable path (i.e. with u in the vicinity of 0) are less fit, in the sense of having fewer offspring and leaving fewer descendants in the short term. This manifests as a small pale area at the bottom of the figure, near $u = 0$. However, after many generations their expected number of descendants greatly outweighs that of individuals further from $u = 0$. The plot was created by forming the graph of parent-offspring relationships for a run of around 10^4 generations, with $N = 5000$ and the same parameters as in previous figures. We then take a random sample of 1.8×10^6 of the resulting individuals. The lineage descending from each individual forms a tree structure, and we calculate the number of nodes as a function of their distance from the root.

the next generation. However, the fixation of an individual’s traits is determined not just by its immediate number of offspring but by its total number of descendants after a long period of time. The relationship between the two is not necessarily straightforward and depends on higher-order effects such as evolvability. In our system, a low magnitude of u is negatively correlated with short-term fitness but positively correlated with long-term number of offspring, which leads to its maintenance in the population.

Here we have presented results only for the case of a one-dimensional spatial structure. For the well-mixed case we have only preliminary results. These show that the performance is much worse than in the one dimensional case over the range shown in Figure 2 (that is, $N \leq 5000$), and the scaling with population size over this range is much weaker, as we would expect from our argument based on absorption time. However, much longer runs will be needed, with much larger population sizes, in order to understand its form. Full results on the well mixed case will be presented elsewhere.

Discussion

We have argued that lineage selection should lead to selection for evolvability when the absorption time is comparable

to (or greater than) the rate at which natural selection affects fitness within lineages. The resulting selection pressure for evolvability should thus be stronger when absorption time is high, which in particular should be the case for very large populations and when spatial diffusion is a limiting factor. We constructed a simple model of a population evolving on a two dimensional fitness landscape in order to demonstrate this, allowing the effects of population size to be studied through computational experiments. These results are important because they give us some clues about when lineage selection is likely to select for evolvability in biology: they suggest that it will become much more important at larger spatial and temporal scales.

In general, the evolution of evolvability is the result of an interplay between two things: on the one hand the topology of the fitness landscape (or, more broadly, all of the biological details of how changes in genotype and phenotype occur and how they relate to ecological function), and on the other the dynamics of evolution at the level of the population. One approach is to focus on the genotype-phenotype level, asking what it is about biological landscapes that allows evolvability to exist and to be controllable genetically (Wagner, 2008); see also Stadler et al. (2001) for a formal treatment of the topology of biological landscapes.

In contrast, our approach has been to construct a simple “toy” landscape with a particular set of properties, in order to focus on the population side of the equation. A key challenge in future work is to focus on the interplay between the two: what does lineage selection look like on more realistic high-dimensional landscapes, and how does the more complex evolutionary dynamics on such landscapes interact with absorption time at the population level?

Complex high-dimensional landscapes are an inherent property of ALife models such as Avida and other genetic algorithms, and we hope the community will take up the challenge of investigating the effect of population size on selection for evolvability in such systems. We suspect that lineage selection plays a strong role in their evolutionary dynamics already, and the challenge is to measure and clearly understand its effects.

Another future challenge is in understanding how this is affected when other time scales come into play. A common theme in previous studies is a periodically changing environment; in this case the time scale of environmental fluctuations is presumably also an important parameter, whose relationship to the other two time scales must be important for the evolutionary dynamics. More broadly, one can model more dynamic situations, in which fitness arises through ecological interactions rather than being imposed directly. The evolution of evolvability can also be important in these situations (Bedau and Packard, 2003; Watson et al., 2016), and there are presumably new phenomena that can occur when ecological time scales also come into play.

We emphasise that lineage selection is not the only mech-

anism through which evolvability can come under selective pressure. In one recent development, Kounios et al. (2016) demonstrate that in a high-dimensional system evolution can improve over time (both converging faster and reaching fitter local peaks), even within a single line of descent. This occurs not via lineage selection but through complex dynamics at the level of the genotype-phenotype map. The mechanism is similar to the way in which deep neural networks are able to generalise, with local fitness peaks playing the role of data points and selection playing the role of gradient descent. More broadly, Watson et al. (2016); Watson and Szathmáry (2016) have argued that such learning-related phenomena are important for evolution at the ecosystem as well as at the individual level.

Although this represents a completely different mechanism from lineage selection, the two effects together might magnify one another. Presumably the effectiveness of generalisation in evolution is a complex function of many parameters, as it is in machine learning. If multiple such lines of descent coexist we might expect those that generalise better to attain higher fitnesses than those that don't, and hence to be fixed in the population. In this way lineage selection could naturally play the role of hyperparameter evolution, tuning the genetic representation to one that facilitates the kind of generalisation that Kounios et al. observe.

Finally it is worth returning to our initial question about the role of lineage selection in the origin of life. The modern genetic system is very complex and was presumably preceded by something else. Regardless of whether this was RNA templating, some other pre-RNA polymer (Hud et al., 2013), attractor-based heredity in autocatalytic networks of peptides (Vasas et al., 2012) or one of the many other mechanisms that have been proposed, there must be some means through which it could give rise to the “central dogma” architecture of modern cells.

Whatever the original mechanisms of heredity were, they must have been much less evolvable than those in existence today. Nevertheless they must have been open-ended enough that there were mutational paths leading to much more sophisticated mechanisms, closer to those of the modern cell. One possibility is that the advantages conferred by evolvability were in themselves enough to drive the selection of better mechanisms of heredity, i.e. that genes and the ribosome — and their precursors — were selected precisely because of the evolutionary benefits they conferred.

Our work adds plausibility to this hypothesis by showing that lineage selection imposes a greater selection pressure for evolvability at larger spatiotemporal scales. On the scale of the evolution of all life on Earth — the largest scale of all — it could have been a strong force indeed.

It has been argued on the basis of phylogenetic data that LUCA (the Last Universal Common Ancestor of extant life on Earth) was not isolated but existed within a global ecosystem comprised of other lineages deriving from still earlier

common ancestors (Zhaxybayeva and Gogarten, 2004). In our model each population is descended not from typical members of old populations but from among the most evolvable. We conclude by hypothesising that the same might be true of life on Earth: perhaps life tried not one but many different cellular architectures, and today's life is descended from those that were the most evolvable. In this way the major features of the modern cell can perhaps be explained in terms of the effect they have on its evolvability.

Acknowledgements

This work was supported the ELSI Origins Network (EON), which is supported by a grant from the John Templeton Foundation. The work has benefited greatly from discussions with Nicholas Guttenberg and others at ELSI, and from comments on the manuscript by Lee Altenberg and Richard Watson.

References

- Aboitiz, F. (1991). Lineage selection and the capacity to evolve. *Medical Hypotheses*, 36(2):155–156.
- Akçay, E. and Van Cleve, J. (2016). There is no fitness but fitness, and the lineage is its bearer. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1687):20150085.
- Altenberg, L. (1995). Genome Growth and the Evolution of the Genotype-Phenotype Map. *Evolution and Biocomputation*, 899(Chapter 11):205–259.
- Bedau, M. A. and Packard, N. H. (2003). Evolution of evolvability via adaptation of mutation rates. *Biosystems*, 69(2-3):143–162.
- Braught, G. W. (2005). Learning and Lineage Selection in Genetic Algorithms. In *IEEE SoutheastCon, 2005.*, pages 483–488. Dickinson College, Carlisle, United States, IEEE.
- Canino-Koning, R., Wisner, M. J., and Ofria, C. (2016). The Evolution of Evolvability: Changing Environments Promote Rapid Adaptation in Digital Organisms. In *International Conference on the Synthesis and Simulation of Living Systems*, pages 268–275, Cambridge, MA. MIT Press.
- Cooper, W. S. and Kaplan, R. H. (1982). Adaptive “coin-flipping”: a decision-theoretic examination of natural selection for random individual variation. *Journal of Theoretical Biology*, 94(1):135–151.
- Crombach, A. and Hogeweg, P. (2008). Evolution of Evolvability in Gene Regulatory Networks. *PLOS Computational Biology*, 4(7):e1000112.
- Dawkins, R. (1989). The Evolution of Evolvability. In Langton, C., editor, *Artificial Life: Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, pages 201–221, Santa Fe, NM.
- Demetrius, L. (1987). An extremal principle of macromolecular evolution. *Physica Scripta*, 36(4):693–701.
- Earl, D. J. and Deem, M. W. (2004). Evolvability is a selectable trait. *Proceedings of the National Academy of Sciences*, 101(32):11531–11536.
- Eshel, I. (1973). Clone selection and the evolution of modifying features. *Theoretical Population Biology*, 4(2):196–208.
- Ewens, W. J. (2004). *Mathematical Population Genetics I. Theoretical Introduction*. Springer, New York, 2 edition.
- Graves, C. J. and Weinreich, D. M. (2017). Variability in fitness effects and the limitations of fitness optimization. *bioRxiv.org* <https://doi.org/10.1101/107847>.
- Harvey, I. (2009). The Microbial Genetic Algorithm. In *Proceedings of the 10th European Conference on Artificial Life, part II*, pages 126–133. Springer, Berlin, Heidelberg.
- Hud, N. V., Cafferty, B. J., Krishnamurthy, R., and Williams, L. D. (2013). The Origin of RNA and “My Grandfather’s Axe”. *Chemistry & biology*, 20(4):466–474.
- Kingman, J. F. C. (1982). On the genealogy of large populations. *Journal of Applied Probability*, 19(A):27–43.
- Kounios, L., Clune, J., Kouvaris, K., Wagner, G. P., Pavlicev, M., Weinreich, D. M., and Watson, R. A. (2016). Resolving the paradox of evolvability with learning theory: How evolution learns to improve evolvability on rugged fitness landscapes. ArXiv preprint <https://arxiv.org/abs/1612.05955>.
- Leibler, S. and Kussell, E. (2010). Individual histories and selection in heterogeneous populations. *Proceedings of the National Academy of Sciences of the United States of America*, 107(29):13183–13188.
- Pigliucci, M. (2008). Is evolvability evolvable? *Nature Reviews Genetics*, 9(1):75–82.
- Stadler, B. M. R., Stadler, P. F., Wagner, G. P., and Fontana, W. (2001). The Topology of the Possible: Formal Spaces Underlying Patterns of Evolutionary Change. *Journal of Theoretical Biology*, 213(2):241–274.
- Vasas, V., Fernando, C., Santos, M., Kauffman, S., and Szathmáry, E. (2012). Evolution before genes. *Biology Direct*, 7(1):1.
- Wagner, A. (2008). Robustness and evolvability: a paradox resolved. *Proceedings of the Royal Society of London B: Biological Sciences*, 275(1630):91–100.
- Wagner, G. P. and Altenberg, L. (1996). Perspective: Complex Adaptations and the Evolution of Evolvability. *Evolution*, 50(3):967–976.
- Watson, R. A., Mills, R., Buckley, C. L., Kouvaris, K., Jackson, A., Powers, S. T., Cox, C., Tudge, S., Davies, A., Kounios, L., and Power, D. (2016). Evolutionary Connectionism: Algorithmic Principles Underlying the Evolution of Biological Organisation in Evo-Devo, Evo-Eco and Evolutionary Transitions. *Evolutionary Biology*, 43(4):553–581.
- Watson, R. A. and Szathmáry, E. (2016). How Can Evolution Learn? *Trends in Ecology and Evolution*, 31(2):147–157.
- Wilke, C. O., Wang, J. L., Ofria, C., Lenski, R. E., and Adami, C. (2001). Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature*, 412(6844):331–333.
- Zhaxybayeva, O. and Gogarten, J. P. (2004). Cladogenesis, coalescence and the evolution of the three domains of life. *TRENDS in Genetics*, 20(4):182–187.