

Fitness Transmission: A Genealogic Signature of Adaptive Evolution

Thomas Miconi
School of Computer Science
University of Birmingham
Birmingham B15 2TT, UK
txm@cs.bham.ac.uk

Abstract

We introduce *fitness transmission* as a simple statistical signature of adaptive evolution within a system. Fitness transmission is the correlation between the fitness of parents and children, where fitness is evaluated after the number of grandchildren, suitably normalised. This measure is a direct calculation based on a genealogical record, rather than on genetic or phenotypic observation. We point out that the Bedau-Packard statistics of evolutionary activity cannot be used as a reliable system-wide signature of adaptive evolution, because they can produce positive signals when applied to certain “random”, non-evolutionary systems. We apply fitness transmission to simple evolutionary algorithms (as well as neutral equivalents) and demonstrate its capacity to accurately detect the presence or absence of Darwinian evolution.

Introduction: Are we evolving yet?

Consider the following problem: imagine that you are observing a simulation, in which a population of agents move, interact and reproduce. The simulation is complex, or its output is obscure (or both), and it is not easy to grasp what, if anything, is going on. Knowing that these agents reproduce, we may ask ourselves the question: are they also *evolving*? Are they undergoing genuine natural selection and adaptive evolution? Or are they just perpetuating random genetic traits, following a chaotic trajectory through genotype space without ever undergoing any meaningful evolution?

This question arises from the fact that when a population of reproducing agents is observed, it is not always immediately clear whether the dynamics of the population result from Darwinian evolution, or merely from random variations and stochastic effects such as genetic drift. The particular system at hand may also introduce its own effects, which may bias or alter the dynamics of the population in unpredictable ways. When this system is sufficiently complex, determining whether a population is evolving in a Darwinian sense may not be a trivial task.

Besides its conceptual implications, the question is of practical interest. It is often desirable to determine whether natural selection and evolutionary adaptation are occurring within a given system, especially in the fields of evolution-

ary computation and artificial life. Indeed in some situations, the onset of significant adaptive evolutionary activity is by itself a major objective of the system: for example, artificial environments such as Echo (Hraber et al., 1997) and Geb (Channon, 2006) were explicitly designed with the aim of exhibiting meaningful evolutionary activity. Being able to detect the presence of genuinely adaptive evolution is a fundamental pre-requisite for the validation of such systems.

Related Work

Traditional methods for detecting natural selection

The problem of detecting natural selection has a long history in biology. Endler’s authoritative treatment (Endler, 1986) describes the traditional (that is, non-molecular) methods for detecting natural selection. However, all these methods are based on *phenotypic observation* of chosen traits: they require collecting statistics on the frequencies of certain, pre-defined traits, and then performing some calculations to determine whether or not natural selection has acted on these traits. This is precisely what we seek to avoid here: we do not ask whether natural selection has acted on this or that trait, but simply whether it is active in the population. Also we want to dispense with detailed phenotypic observation.

The molecular revolution in biology has made it possible to collect vast amounts of genetic data. This creates new possibilities for the detection of natural selection, based on direct assessment of nucleotide variation (Sabeti et al. (2006) provide a recent review). But these approaches require access to a full genetic record. Furthermore, biological genomes are simple sequences of symbols from a four-letters alphabet; but artificial life models need not be so simple in their structure, and this may affect the applicability of these methods.

The Bedau-Packard measure of evolutionary activity

Bedau and Packard (Bedau and Packard, 1992; Bedau et al., 1998) have developed a groundbreaking set of concepts and methods to “discern whether or not evolution is taking place in an observed system.” Bedau and Packard are specifically

interested in the *innovations* produced by evolution, and in the capacity of various systems to keep on producing adaptive innovations over time - or not. This requires a method to determine whether an apparent innovation is indeed adaptive or merely the result of random fluctuations, which clearly relates to our own concerns. To this end, Bedau and Packard introduce a set of methods to compute the “*evolutionary activity*” of components and, by extension, of systems.

The Bedau-Packard measures of evolutionary activity are based on *persistence* of adaptive innovations: they identify components that persist over time at a level that exceeds what would be expected under purely random conditions. In the words of Bullock and Bedau (Bullock and Bedau, 2006), “if a particular element persists in the system for a long time, this is likely to be because it is being maintained by selection.”

If we are to use persistence “for a long time” as a criterion for detecting evolution, we need a method to determine what “a long time” is. When do we decide that a given element has persisted long enough to be regarded as ‘adapted’? To tackle this problem, Bedau and Packard introduced the idea of using a neutral “shadow” of the system under study: a replication of the original system, in which birth, reproduction and death of individuals occur in synchronisation with the real system, but are applied to randomly chosen individuals. More precisely, every time a new individual is being created in the real system under study, a new individual is also created in the shadow; but with the difference that, in the shadow, the parents of the new individual are chosen randomly. Thus the neutral shadow is expected to show the behaviour that would be seen in the system, in the absence of any selective pressure. By comparing the persistence data obtained in this “shadow” to that obtained in the real system, Bedau and Packard argue, it should be possible to detect whether selection and adaptive evolution are present.

Building upon the concept of enduring persistence as a measure of evolutionary activity, Packard and Bedau have developed a series of evolutionary statistics based on persistence information. These statistics include diversity D (the number of different components present at a given time in the population), activity $a_i(t)$ (the age of component i at time t , indicating how long it has persisted so far), cumulative activity $A_{cum}(t)$ (the sum of the ages of all components present at time t), and new activity $A_{new}(t)$ (the sum of the ages of all components present in the system at time t that are new, but sufficiently aged to indicate adaptive value, divided by diversity at time t).

Bedau-Packard statistics and non-evolutionary systems

Bedau and Packard’s measures are arguably the most widely known of their kind. They have been applied to several systems, including artificial ecologies such as Echo, and natural components such as the genera within the fossil record

(Bedau et al., 1998). Other researchers have successfully applied them to various systems (Standish, 2002; Channon, 2006; Taylor, 1999). However, it is *not* suitable as a test to detect the presence of adaptive evolution within a system. The basic reason why the Bedau-Packard statistics cannot be used as a detector of evolution by natural selection is that they may attribute a positive score to “random” processes, which are clearly not evolutionary. Importantly, this is the case even if a shadow is used to normalise activity scores. The crux of the matter is that these statistics essentially track “excess” variance in the persistence of components, which is used as a proxy for selection and therefore (it is argued) for adaptive value. The shadow is used to define the level of persistence which can be termed “excess.” But excess variance in persistence may be caused by other factors than natural selection. What if, due to some quirk in the rules of the system, some high variance in persistence occurs that is not related to heritable characteristics? If we apply the Bedau-Packard statistics to such a system, we may find that the Bedau-Packard measure classifies such a system as adaptive, even though it is not - even if we use a shadow.

It is easy to devise examples of systems which illustrates this distinction. For instance, consider a population in which reproduction, selection and evolution occurs normally, except for the fact that fitness is randomly attributed to each individual at birth, independently of its genome. That is, while genetic material is transmitted as expected from parent to offspring, this genetic material has no influence over fitness, which is chosen randomly for each new individual. Note that no heritable variance in fitness occurs, nor does any adaptation take place. However, those individuals that happen to be highly “fit” (out of sheer luck) will tend to persist for a long time, and may flood the population with their (short-lived, but nevertheless genetically similar) offspring. No such thing will be observed in the shadow, where reproduction and survival will be random, leading to random diffusion of the genetic material throughout genotype space. Therefore, a difference will occur between the activity counts (and diversity counts) of the shadow and of the real system, creating a positive signal on the Bedau-Packard measure and associated tests.

In figure 1 we describe the results of Bedau-Packard statistics applied precisely to such a system.¹ The system is a simple steady-state genetic algorithm in which, at every “generation”, 10 out of the 100 individuals are eliminated and replaced by new individuals, created by copying and mutating a surviving parent. Survivors are selected by fit-

¹In these experiments we have applied the Bedau-Packard statistics to entire genotypes, in order to follow the authors’ method (Bedau et al., 1998). However we are not at all certain that whole genotype persistence is a reliable indicator of evolution. We note that in nature, as soon as recombination and mutation are involved, it is very unlikely that any genotype ever persists for more than one generation.




Figure 1: Graphs showing the results of Bedau statistics for a non-Darwinian system, as well as for a corresponding shadow system. The leftmost graph indicates the cumulative frequency counts for each genome over time (that is, the running sum of the frequency of each genome within the population at each generation.) The middle graph indicates the cumulative distribution of persistence counts for all genomes over the history of a run (that is, for each value, the number of individuals that survived longer than this value.) The rightmost graph shows the average cumulative activity $\bar{A}_{cum}(t) = \frac{A_{cum}(t)}{D(t)}$ - the sum of all persistence counts of genomes present at a time, divided by the number of genomes present at that time. These graphs are consistent with what is expected from real evolutionary systems. (Bedau and Packard, 1992; Bedau et al., 1998)

ness ranking, and selection of parents occur through tournament selection, very much as in a normal genetic algorithm. However, the fitness of individuals is randomly chosen at birth, independently of their genome. The actual method to “calculate” fitness is to increase a certain counter repeatedly until a random number picked between 0 and 9 is equal to 0 (thus the distribution of fitnesses is exponential.) In the shadow systems for both experiments, selection of survivors and parents are random (thus the shadow systems for both experiments are essentially identical, which predictably results in similar graphs).

Note, in particular, the onset of high average activity, the flattening of the cumulative distribution of persistence counts (with an order of magnitude difference between the longest-living genotypes of real and shadow systems), and perhaps most significantly the appearance of large “telic waves” (Bedau and Packard, 1992) (tall, lengthy lines) in genome frequency plots, despite the decidedly non-teleological nature of these environments. All these are regarded as positive signals of evolutionary activity in the Bedau-Packard framework. (Bedau and Packard, 1992; Bedau et al., 1998)

Surely many other examples could be found. More generally, these simple systems illustrate the fact that high variance in persistence can be caused by many other processes than natural selection. “Random” systems, in which no meaningful evolution or adaptation occurs, can still obtain high marks on the Bedau-Packard measure if they produce high variance in genetic persistence.

Of course, in our toy system, it is easy to see (just by looking at the rules) that variance in persistence is due to random fluctuations, and that no true natural selection exists. But this is precisely the heart of the matter. First, when we study a real system, we may not have access to its internal rules, so clearly in this case we cannot use the Bedau-Packard statis-

tics as a test of Darwinian evolution. But even if we *do* have full access to the rules of the system, the complexity of even mildly elaborate systems may prevent us from asserting with absolute certainty whether or not a “random force” generates strong variance in persistence. For example, considering a system similar to Echo (Hraber et al., 1997), can we really exclude, a priori, that such a factor could come into play? Can we offer absolute guarantee, simply by looking at the rules of the system, that no weird effect will arbitrarily and significantly favour certain individuals rather than others (without being based on these individuals’ heritable features)? The answer, of course, is that we cannot. It follows that, if we apply the Bedau-Packard statistics on such a system and obtain a positive result, we cannot (in the absence of further information) use this fact alone to conclude that adaptive evolution is active in the system.

It is important to be clear about the meaning of this result: this should not be interpreted as a minimisation of the importance of Bedau-Packard statistics. Rather, this is a reminder that these statistics should not be used to detect adaptive, Darwinian evolution within a system, even by normalising against a shadow. *If* we know, a priori and through other means, that the system is indeed affected by genuine adaptive evolution, and if we can rest assured that “weird” effects will be nil or negligible, *then* we can fruitfully apply the Bedau-Packard measure to assess the dynamics of long-term evolutionary innovation within this system. The valuable contribution of these statistics in this regard has often been pointed out. However we cannot use these statistics to determine the presence of evolution by natural selection within a system, as opposed to any system-induced dynamics which create high variance in persistence: the Bedau-Packard statistics are not designed to distinguish the former from the latter, even by using a shadow system.

Fitness Transmission: A test statistic for natural selection

Darwinian evolution: randomness, selection and heredity

In general, evolution is simply defined as a change in the frequencies of heritable innate characteristics within a reproducing population, from one generation to the next. Natural selection, one of the mechanisms that guide evolution, is broadly defined as variance in reproductive success caused by heritable innate characteristics. Darwin realised that adaptive evolution automatically results from the existence of fitness-impacting, heritable variations. Variations that improve fitness will be propagated quickly, initiating thriving lineages; while those that reduce fitness will hinder their own propagation, creating feeble (or even quickly extinct) lineages. Thus lineages constantly branch out into variants, and the uneven distribution of these branches, being dramatically skewed towards those which result from fitness-enhancing variations, will result in the overall effect that the newer descendants of the original lineage will tend to be those better adapted to their current, local environment: heritable fitness-affecting variation will have “steered” the original lineage towards adaptive directions among all those encountered by mutational variations.²

Note that although this process will usually result in a modification of the species over sufficiently long periods of time, it will also often result in temporary stasis. If a species happens to be located at a convenient local optimum in the fitness landscape, then variations which depart from the optimum will mostly reduce the fitness of the individual. In this case the differential transmission of characteristics enforced by natural selection will actively maintain the population around the optimum: the population will be constantly steered back towards its current position. This phenomenon, known as ‘stabilising selection’, is actually regarded as more common than directional selection (see (Ridley, 1993), Chap. 4.4).

Fitness Transmission: A genealogic signature of Darwinian evolution

From this discussion we can deduce a method to detect the active presence of natural selection. If fitness-impacting, heritable traits are actually being transmitted and propagated, then this should have an impact on the genealogical record: individuals sharing a common lineage, being more likely to inherit common fitness-impacting characteristics, should therefore tend to exhibit slightly similar fitnesses in comparison to the rest of the population. In other words, if some fitness-affecting traits are being transmitted, then there should be some degree of *correlation* between the fitnesses

²Or, in short: as creatures replicate, genes mutate, adaptations proliferate, and species originate.

(that is, the reproductive success) of individuals from a common lineage: the transmission of heritable, fitness-affecting traits should result in some degree of *differential transmission of fitness*.

Fitness transmission is our proposed signature for natural selection. It is, quite simply, the statistical correlation between the fitness of children and parents. The basic idea of fitness transmission is that, when natural selection is active in a population, parents and children should exhibit a *tenuous*, but *persistent* correlation in fitness.

Calculation of Fitness Transmission

Number of grandchildren as a measure of fitness

The term “fitness” is notoriously ambiguous and can be a significant source of confusion (Dawkins, 1982, Chap. 10). A common practical measure of an individual’s fitness is its number of grandchildren, rather than number of children. To have many grandchildren, an individual must not only have many children, but these children themselves must also be successful in reproducing; this corresponds to the intuitive notion of fitness as ability to pass on one’s genes. We will use the number of grandchildren as a measure of individual fitness. Therefore, to measure fitness transmission, we measure the statistical correlation between the number of grandchildren (NOGC) of an individual, and that of its children.

Fitness correlation is a local measure in time. That is, we divide the record in time periods, or “slices,” and calculate fitness transmission independently for each period. This is done by only considering individuals born within this time period for the “child” data set of each period (the parents of these individuals are then collected in the “parents” data set, independently of their time of birth). However, the reproductive success for a given individual may be collected over its entire history, even if it goes beyond the time-slice being considered.

Comparing what is comparable

As usual when calculating statistical correlations, care should be taken in only comparing what is comparable: conflating data from widely different distributions may result in artificial, spurious correlations. In some artificial systems, selective conditions may change widely over the course of an evolutionary run, even with a fixed fitness function. This may wreak havoc on undiscerning evaluations of statistical correlation. For example, in a simple genetic algorithm, if strict ranking is used, surviving and reproducing entails dislodging a previous survivor; but as evolution proceeds towards an optimum, and new champions are increasingly well-adapted, it becomes increasingly difficult (and thus rare) for new individuals to dislodge previous champions. This means that the children’s fitness will tend to go down (because more of them disappear without a descent) and the parent’s fitness will tend to go up (because they remain in the population longer) over time. This alone is sufficient to

create a strong, *negative* correlation between the fitnesses of parents and children over the whole process: earlier parents would have a moderate number of grandchildren, each with a good chance to reproduce; while later parents would accumulate enormous numbers of grandchildren, which would have comparatively low reproductive success.

To avoid this, we must ensure that we only consider quantities (that is, fitnesses) obtained under similar conditions. To this end, the periods over which reproductive successes are measured should start at the same point in time, so that we can ensure that they are obtained over equivalent conditions. In practice, this means that when we compare the NOGC of an individual X and its parent, we should only consider the grandchildren of the parent that were born at the same time as X or later. This ensures a “fair game” between the parent and the child: both scores will be obtained under similar circumstances, and results obtained by the parents in earlier (possibly harsher or milder) circumstances will not spoil the data.

Necessary normalisations

Unfortunately, the choice of using NOGC as a measure of fitness introduces an obvious problem: the NOGC of an individual and that of its children are clearly not independent quantities. Saying that A has many grandchildren is saying that A’s children have many children, and therefore, *out of this fact alone*, are likely to have many grandchildren themselves, even with random reproduction. This problem can be easily addressed by normalisation to make the considered values independent. To do this, we do not use the raw NOGC for the parents; rather, for every parent-child pair, we consider the parent’s NOGC *minus the number of children of this particular child*. This modified NOGC is an estimation of the parent’s fitness that is not biased by this particular child’s own success, and thus any correlation represents a true correlation in fitness.

Another, less significant problem is that, in general, the population of interest will be finite. The consequence is that the reproductive successes of individuals living during the same period of time are not independent: any child for a given individual is one less opportunity for another individual to have a child. Even with random mating and reproduction, if one individual happens to have more children than average, then any other randomly picked individual is mechanically more likely to have fewer children than average. In other words, limited population introduces a slight negative correlation between the modified NOGC of parents and children. This effect is much less important than the previous one, but may be noticeable, especially with small populations. A simple solution to this problem is to normalise the modified NOGC of the parent: for every parent-child pair (P_i, C_i) from the slice, we divide the modified NOGC of P_i by the total sum of all grandchildren of all other parents within the slice - minus C_i ’s children. The resulting

proportion is independent of this child’s own success.

Those normalisations are made necessary by the fact that the quantities under scrutiny are not independent. They would become unnecessary if, instead of evaluating fitness transmission from parents to children, we attempted to calculate it between grandparents and grandchildren. The problem, of course, is that any signal would be much weaker due to the increased indirection - often to the point of being drowned in noise.

Calculation method for fitness transmission

Where does this leave us? From all these considerations, we can deduce the following calculation method for fitness transmission:

- Divide the entire genealogic record into discrete periods of time. If the system is generational, generations may be used as time periods.
- For every time period within the genealogic record, perform the following operations:
 1. For every individual C_i born during this time period, find its parent P_i (which may be born at any time before C_i , not necessarily during this time period) and store the resulting parent-child pair (P_i, C_i) . Note that any given individual may occur in several pairs.
 2. For every stored parent-child pair (P_i, C_i) , retrieve their respective total number of grandchildren (NOGC) $N(P_i)$ and $N(C_i)$, born during or after (not before) this time period.
 3. Elimination of dependency: for every pair (P_i, C_i) , subtract the number of children of C_i from the $N(P_i)$, resulting in the new value $N'(P_i)$.
 4. Normalisation: for every parent P_i in the set of parent-children pairs for this time period, divide $N'(P_i)$ by the sum of all grandchildren of all other parents $P_{j \neq i}$ - carefully excluding C_i and its descendants from the count. This results in a final value $N''(P_i)$.
 5. Calculate the statistical correlation between the $N''(P_i)$ and the $N(C_i)$ variables over all parent-child pairs for this time period, using the standard Pearson formula:

$$\text{Corr}(X, Y) = \frac{\sum_{i=1}^N (x_i - \bar{X})(y_i - \bar{Y})}{(N - 1)\sigma_X \sigma_Y}$$

The resulting value $\text{Corr}(N''(P_i), N(C_i))$, for every time period, is our estimator for the intensity of fitness transmission during that time period.

Experiments

Experimental settings

Our purpose in this section is to set up a couple of experiments in order to determine whether fitness transmission

is indeed a reliable indicator of Darwinian evolution. To do this, we will use simple evolutionary systems with predictable dynamics, in which the presence or absence of evolution can be easily controlled. We will apply our calculation method to these systems and determine whether the presence or absence of Darwinian evolution was successfully detected.

To perform our experiments, we used genetic algorithms involving a population of 1000 individuals, over 100 generations. We considered two optimisation problems: the Rosenbrock function $100(x^2 - y^2)^2 + (1 - x)^2$ (using genomes of 2×12 bits) and a very simple OneMax problem over 20 bits. The Rosenbrock function is a commonly used test function in the field of optimisation. The purpose of the simple OneMax problem is to examine the behaviour of different algorithms on very easy problems, when the the global optimum is discovered quickly. In our algorithms, at each generation, a new population is created either by applying bitwise mutation to a parent selected from the previous generation, or (with 66 % probability) by applying one-point crossover between two parents, and then applying bitwise mutation to the resulting offspring. The probability of mutating (flipping) each bit is the inverse of the total number of bits in the genome, rounded to the closest higher percent; thus, on average, each genome should undergo about one mutation. As explained below, we tested different methods of selection and replacement.

As a point of comparison, we need a “neutral” version of the genetic algorithm, which preserves as many features of the algorithm as possible, while effectively removing Darwinian evolution. We chose to use a system in which every new individual was attributed a random genotype (and therefore a random fitness) at birth, regardless of the genetic make-up of its parents. This is different from purely random selection in that selection still occurs, and is still based on fitness; however the randomness of the reproductive process prevents any meaningful evolution: fitness-affecting traits are still present, but not heritable. A satisfactory measure of evolutionary activity should be able to detect the absence of real evolution and return a zero value for this situation.

A simple genetic algorithm

We first describe the calculation of fitness transmission in a standard simple genetic algorithm, using tournament selection. In this algorithm, each new individual is created by selecting parents from the previous generation (using tournament selection), and generating offspring as previously described. The process is iterated until the new population is filled.

Figure 2 shows the results of these calculations, applied to the “fossil record” generated by our simple genetic algorithm. This figures shows the results for the Rosenbrock function optimisation problem with 20 bits, both with normal reproduction and with reproduction based on random




Figure 2: Rosenbrock function, non-overlapping generations, 5 different runs (top) and average of 50 different runs (bottom).




Figure 3: OneMax function, non-overlapping generations, 5 runs (top) and average of 50 runs (bottom).

phenotypes. The top graph shows the results of 5 different run for each of these reproduction methods, while the bottom graph shows average curves over 50 runs. Figure 3 shows the same data for the OneMax problem. In the normal selection case, the correlation between the number of children of parents and children is distinctly positive (especially at the very beginning at the run) and stabilises to a positive plateau. The enduring positive value indicates that the population is constantly and actively maintained in the

vicinity of the global optimum (which is reached quite early in the OneMax problem) through active evolutionary forces. Even though the optimum has been reached, mutation constantly disperses the population, and Darwinian evolution constantly drives it back. Stabilising selection results in a positive value for differential fitness transmission. In the case of random genotypes, as expected, no meaningful fitness transmission occurs.

That the enduring presence of fitness transmission in this case is caused by mutation can be seen quite readily. If we set the mutation rate to zero, then the population converges totally: all individuals end up sharing the exact same genome, and diversity disappears. From this point on, all individuals having exactly the same genotype, evolution simply stops. The result is that evolutionary activity, as indicated by fitness transmission, quickly goes to zero (with noise oscillations) after an initial phase of high activity (see Figure 4). This illustrates the capacity of fitness transmission to distinguish between active stabilising selection on the one hand, and passive stillness caused by absence of genetic variation on the other (though this ability breaks down in extreme situations, as discussed in section .)




Figure 4: OneMax function, non-overlapping generations, without mutation, 5 runs (top) and averages of 50 runs (bottom).

Removing selective gradient among parents

Here we try to make the problem more challenging problem by reducing the scope of selection. To do this, we modify our algorithm as follows: at every generation, a small set of survivors is selected from the population through strict ranking selection, and the parents for the next generation are then *randomly* selected from among this set of survivors. Offspring are created as previously mentioned (66% crossover,

mutation, etc.) The effect of this modification is to effectively remove any selective gradient among parents. This is because the only effect of selection in this system is to decide which individuals become parents in the first place. Once individuals have been selected as parents, their number of children is random, and as a result is not affected by natural selection. In particular, note that if we had tried to evaluate fitness by the number of children alone, then no fitness transmission could be detected: no correlation can exist between the number of children of parents and children, simply because all parents have a random number of children. However, as shown in figures 6 and 5, our measure for fitness transmission is able to detect the signal created by this more indirect form of natural selection.




Figure 5: Rosenbrock function, non-overlapping generations with ranking-based survival and random parent selection, 5 runs (top) and averages over 50 runs (bottom). The initially high signal goes to a very low, but still noticeably non-zero value.

Limitations of fitness transmission

Although fitness transmission is valuable as a signature of adaptive evolution, several limitations must be mentioned.

Extreme stabilising selection: While fitness transmission is able to detect moderate stabilising selection, it breaks down in the extreme situation of *absolute* stabilising selection - that is, when only one genotype is viable, and any individual that differs from the optimum systematically fails to reproduce. In this case, no heritable variation in fitness exists. In this situation, stabilising selection has the effect of effectively freezing the reproducing population, and therefore becomes invisible to fitness transmission.

Only one extant lineage: More generally, there are pathological situations in which genealogic methods can not be

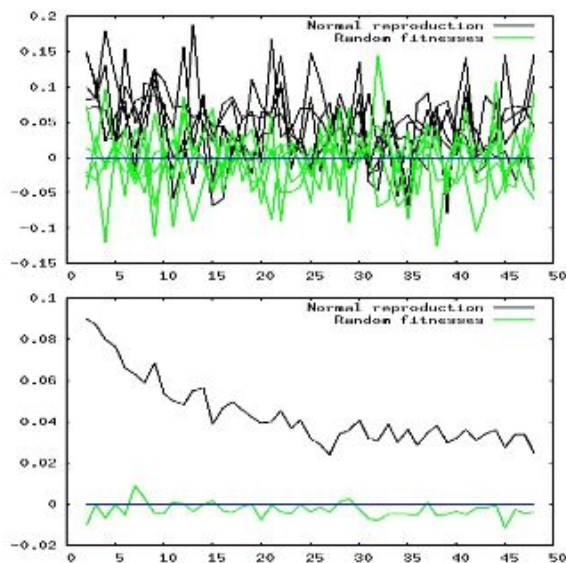


Figure 6: OneMax function, non-overlapping generations with ranking-based survival and random parent selection, 5 runs (top) and averages over 50 runs (bottom).

used at all. One such situation occurs when all individuals present at any given time share the exact same genealogic tree - in other words, when there is never more than one lineage in the population. In this case, while Darwinian evolution can certainly occur, the presence of only one lineage at any time within the population prevents the possibility of inter-lineage comparison, upon which genealogic analysis relies. For example, consider a non-overlapping generational system, such that at every generation, two individuals are selected to serve as parents for the next generation, and *all* the individuals from the new generation are children of *both* of those selected parents. Since all individuals will always share the exact same set of parents, grandparents, and so on, fitness transmission cannot be applied. We believe that this situation is sufficiently exotic to preserve the usefulness of genealogical analysis. In addition, such situations can be easily detected in any system for which a genealogical record exists.

Non-biological selection: A more subtle aspect of fitness transmission is that it detects natural selection in the most general sense, applying to *any* heritable character, including those that we might not think of as “biological”. Any kind of heritable trait that affect reproductive success (genetic, epigenetic, cultural, etc.) will be detected by fitness transmission. If the objective is to detect biological natural selection *alone*, then fitness transmission should not be used on its own.

Conclusion

We have shown that differential fitness transmission is a useful signature of Darwinian evolution, which can be detected in genealogical record by using simple statistics. We believe that this signature may be more suitable for this purpose than previously suggested methods for detecting evolution. We have applied this statistic to the genealogical records generated by real evolutionary algorithms, demonstrating its capacity to detect the presence or absence of adaptive evolution.

Acknowledgements

This research was conducted at (and funded by) the School of Computer Science at the University of Birmingham, under the supervision of Alastair Channon. This paper benefited greatly from the comments of anonymous reviewers for the *Artificial Life* journal.

References

- Bedau, M. A. and Packard, N. H. (1992). Measurement of evolutionary activity, teleology, and life. In Langton, C., Taylor, C., Farmer, D., and Rasmussen, S., editors, *Artificial Life II: Proceedings of the Workshop on Artificial Life*, pages 431–461, Redwood City, CA. Addison-Wesley.
- Bedau, M. A., Snyder, E., and Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In Adami, C., Belew, R. K., Kitano, H., and Taylor, C. E., editors, *Artificial Life VI: Proceedings of the Sixth International Workshop on the Synthesis and Simulation of Living Systems*, pages 228–237, Cambridge, MA. MIT Press.
- Bullock, S. and Bedau, M. (2006). Exploring the dynamics of adaptation with evolutionary activity plots. *Artificial Life*, 12:1–5.
- Channon, A. D. (2006). Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines*, 7(3):253–281.
- Dawkins, R. (1982). *The extended phenotype: The gene as the unit of selection*. Freeman.
- Endler, J. (1986). *Natural Selection in the Wild*. Princeton University Press.
- Hraber, P. T., Jones, T., and Forrest, S. (1997). The ecology of Echo. *Artificial Life*, 3(3):165–190.
- Ridley, M. (1993). *Evolution*. Blackwell Scientific.
- Sabeti, P., Schaffner, S., Fry, B., Lohmueller, J., Vailly, P., Shamovsky, O., Palma, A., Mikkelsen, T., Altshuler, D., and Lander, E. (2006). Positive Natural Selection in the Human Lineage. *Science*, 312(5780):1614–1620.
- Standish, R. K. (2002). Diversity evolution. In Standish, R. K., Bedau, M., and Abbas, H. A., editors, *Artificial Life VIII: Proceedings of the Eighth International Workshop on the Synthesis and Simulation of Living Systems*, pages 131–137. MIT Press.
- Taylor, T. (1999). *From artificial evolution to artificial life*. PhD thesis, University of Edinburgh.