A copy-edited version of this paper will appear in the journal Artificial Life, published by MIT Press.

Effect of Environmental Change Distribution on Artificial Life Simulations

John A. Bullinaria

School of Computer Science
University of Birmingham
Birmingham, B15 2TT, UK

j.a.bullinaria@cs.bham.ac.uk

Abstract: It is already well known that environmental variation has a big effect on real evolution, and similar effects have been found in evolutionary artificial life simulations. In particular, a lot of research has been carried out on how the various evolutionary outcomes depend on the noise distributions representing the environmental changes, and how important it is for models to use inverse power-law distributions with the right noise colour. However, there are two distinct factors of relevance: the average total magnitude of change per unit time and the distribution of individual change magnitudes, and misleading results may emerge if those factors are not properly separated. This paper makes use of an existing agent-based artificial life modeling framework to explore this issue using models previously tried and tested for other purposes. It begins by demonstrating how the total magnitude and distribution effects can easily be confused, and goes on to show how it is possible to untangle the influence of these interacting factors by using correlation-based normalization. It then presents a series of simulation results demonstrating that interesting dependencies on the noise distribution remain after separating those factors, but many effects involving the noise colour of inverse power-law distributions disappear, and very similar results arise across restricted-range white-noise distributions. The average total magnitude of change per unit time is found to have a substantial effect on the simulation outcomes, but the distribution of individual changes has very little effect. A robust counter-example is thereby provided to the idea that it is always important to use accurate environmental change distributions in artificial life models.

Keywords: Agent-based models, Evolution, Variable environments, Learning strategies, Population collapse.

1. Introduction

Over recent years, simulations of artificial populations have been widely used to elucidate the evolution of real populations and their behaviours, employing both mathematical models and agent-based approaches. They have demonstrated how different physical structures and associated behaviours can evolve depending on the given tasks and environments, and how diverse life histories emerge naturally across distinct species (e.g., Heino, Ripa & Kaitala, 2000; Whitehead, 2007; Kamenev, Meerson & Shklovskii, 2008; Thornton & Raihani, 2008; Bullinaria, 2009, 2017, 2018; Borg & Channon, 2012; Smaldino, Newson, Schank & Richerson, 2013; Mathot & Frankenhuis, 2018; Bartoli et al., 2020; Grove, Borg & Polack, 2020; Haaland, Wright & Ratikainen, 2020; Morgan, Suchow & Griffiths, 2020). Such models normally require a considerable degree of abstraction and simplification, both to minimize potential confounding factors that may leave the conclusions unreliable, and to render the simulations computationally tractable. Keeping the models as simple as possible is a sensible standard strategy, but finding a good trade-off between simplification and realism of the models is one of the key challenges of artificial life modeling. The particular challenge addressed in this paper is how realistically do the environments need to be represented in the models, which is a topic that has attracted considerable scrutiny already (Morales, 1999; Heino, Ripa & Kaitala, 2000; Halley, 1996; Halley & Inchausti, 2004; Vasseur & Yodzis, 2004; Thornton & Raihani, 2008; Ruokolainen, Linden, Kaitala & Fowler, 2009; Grove, Borg & Polack, 2020; Grove, Timbrell, Jolley, Polack & Borg, 2022). The starting point for many models has been to assume simple fixed environments, but that can be problematic because natural environments normally vary over time, and what evolves in changing environments may be very different to what evolves in static environments (Feldman, Aoki & Kumm, 1996; Whitehead, 2007; Botero, Weissing, Wright & Rubenstein, 2015; Bullinaria, 2018; Grove, Borg & Polack, 2020; Haaland, Wright & Ratikainen, 2020).

The key issue here is that populations required to survive in a changing environment will generally need to have a greater degree of adaptability (Potts, 1996, 1998, 2013; Grove, 2011; Haaland, Wright & Ratikainen, 2020). If the changes are sufficiently slow, it may be sufficient for physical adaptations to proceed by evolution, particularly when high levels of population diversity are maintained, but many changes are faster and larger and best coped with by behavioural adaptations over shorter timescales (Borg & Channon, 2012). For many species, that adaptability will naturally require the learning of modified behaviours, and that learning could be individual, social or some combination (Feldman, Aoki & Kumm, 1996; Kameda & Nakanishi, 2002; Whitehead, 2007; Borenstein, Feldman & Aoki, 2008; Boyd, Richerson & Henrich, 2011). An important open research question to be addressed in this paper is how do the particular learning strategies that evolve, and the population dynamics that emerge, depend on the details of the environmental variability. For the purposes of most realistic scenarios, the principal details of relevance are the average total magnitude of change per unit time and the distribution of individual change magnitudes. One fruitful existing approach for representing these and modeling the consequences of environmental change has been to perform agent-based simulations of evolving populations of abstract learning individuals. This paper extends previous work in this area (Bullinaria, 2018), looking specifically at the effect of the environmental change distributions, and how crucial it is to control for any interacting factors. It can also be seen as following on from the recent studies in this area by Grove, Borg & Polack (2020) and Grove et al. (2022). In particular, it explores the prevailing view that getting the distribution of individual changes right is crucial for reliable models, and asks whether previous simulation results in this area may be misleading because the effect of the distribution has not been separated from the effect of the average total magnitude of change.

Obviously, it will not be possible to explore this matter on all possible artificial life models, nor even for all possible parameter settings in one particular model. Instead the aim here will be to use an existing tried and tested model to demonstrate the general issues involved and to provide a simple robust counter-example to the idea that the chosen distribution of environmental changes will generally have an important effect on the simulation outcomes. This initial study can then form the starting point for future investigations into what aspects of the model will need to be changed to cause significant effects to appear. The next three sections describe the abstract agent-based learning model to be used and how the environmental changes can be efficiently represented within it. Then, a series of simulation results are presented that explore the key issues, in particular, how the crucial interacting factors relating to the change distributions can be separated, and what effects remain when that has been done. Throughout, care will be taken to point out which modeling simplifications have little effect on the outcomes, and which should be avoided because they do result in significant differences. The paper ends with some conclusions and discussion.

2. Agent-Based Learning Model

A fairly standard agent-based evolutionary modeling approach will be used for this study that involves maintaining a population of abstract simulated individuals, each of which is specified by a simplified set of evolvable parameters that represent their key innate properties (Eiben & Smith, 2016). Those properties then map, indirectly, on to an environment-dependent measure of performance which is used to determine selection for survival and procreation. The populations mimic real species in that they are comprised of competing individuals of all ages, with the oldest and worst-performing dying each simulated year, to be replaced by children of the best performing. Simulating such populations over many generations using previously optimized crossover and mutation operators for the procreation leads to the evolution by natural selection of properties that enable better performance. If the environment changes over time, that may drive the evolution of more appropriate properties for survival, or it may result in reduced performance levels that in extreme cases could be severe enough to result in population collapse. This section will describe the general learning agents on which the models to be studied are based, leaving the particular evolutionary simulation and analysis details used for this study to be specified in Section 4 after the issues surrounding environmental variability have been discussed in Section 3.

2.1 Representing the agents

In practice, all the models need to do at each simulation time-step is update the record (i.e., list of numbers) specifying each individual agent in the population. For this study, the record for each agent simply consists of their innate properties (that are only updated when a dead agent is replaced by a new-born agent), their age, and a list of what (if anything) they have learned so far during their lifetime (which may change at each time-step). It is the representation used for those records and the specification of the various update processes that differentiate one model from another, and they will typically require details of all the learning, survival and procreation processes, which in turn will typically depend on how the agents interact with each other and their environment.

For some species, the relevant innate properties may simply be their physical characteristics, and that alone is what determines their performance throughout their lives, but for many the ability to learn during their lifetimes to improve or adapt their behaviour is also important. Agent-based models have previously been developed to explore how lifetime learning can be an important factor in life history evolution (Bullinaria, 2009), and how different learning approaches will evolve to improve the performance of the particular species being modeled (Bullinaria, 2017). Moreover, extensions of those models have shown how different learning strategies will evolve depending on whether and how the environment changes over time (Bullinaria, 2018). This paper will further extend those models to explore more carefully how the evolutionary outcomes depend on the noise magnitudes and distributions underlying the environmental variations, and test the widely held belief that using the correct noise distribution is crucial for achieving realistic results.

The general modeling framework to be used here has previously been described, tested and optimized (Bullinaria, 2017, 2018), so only the details necessary for understanding and replicating the new results will be presented here. For most real species, it is not feasible to simulate the co-evolution of all their features, so the standard simplification is to separate out a small subset of factors of interest for simulated evolution, with the rest either fixed at their observed evolutionary endpoints or abstracted out completely. So here the physical characteristics of the modeled species will be taken to have already evolved (and will continue to evolve) to be broadly suitable for survival in their (possibly changing) environment, so they need not be explicitly included in the models. It is also assumed that the whole population is sufficiently similar physically that the individuals' behaviours are the main factor in determining their performance. Where the individuals can differ is in their learning strategies, and it is what they have each learned at each stage throughout their lifetimes that leads to the different behaviours and performance levels that drive the evolutionary selection processes. So, to summarize: the only innate properties that need to appear in each agent's record are parameter values specifying their learning processes.

2.2 Representing the learned behaviours

There are two crucial types of learning involved: direct individual learning from the environment, and imitative social learning from other individuals. The direct learning has the advantage that it can automatically track any environmental changes, but has the disadvantage that it can be slow, and attempting to speed it up can lead to errors. Social learning has the advantage of being quicker, and if the individuals learned from are selected according to their performance levels, the learned behaviours could be better than those learned directly. Indeed, it has been argued that social learning is essential for human adaptation (Boyd, Richerson & Henrich, 2011). However, if the environment changes, some socially learned behaviours may actually be out of date and detrimental rather than helpful. Moreover, if the population becomes too reliant on socially learned behaviours, environmental changes can lead to those behaviours becoming insufficient to sustain survival and result in population collapse (Whitehead & Richerson, 2009; Diamond, 2011; Botero, et al., 2015).

Learning in real species generally occurs in complex neural network systems, and although building evolutionary neural network based models is possible, the simulation run times for even the simplest of tasks are enormous (Bullinaria, 2009). A more computationally tractable approach, that also provides a more convenient way to accommodate the rather different social and direct learning approaches in the same simulation framework, has previously been found (Bullinaria, 2017). This involves maintaining a fixed pool of discrete abstract units of acquired information or behaviour that (for want of a better word) are called 'memes', and allowing individuals to acquire them via evolved learning strategies. This was inspired by an earlier meme-based approach that had been developed by Higgs (2000) to simulate the evolution of imitative social learning and the mimetic transition, though further testing and the need to incorporate individual learning and changing environments led to many changes from those original models (Bullinaria, 2017, 2018). Comparisons with models based on more realistic neural network learning (Bullinaria, 2009) confirmed that the new meme-based simulations of individual learning were able to capture the same life history evolution effects (Bullinaria, 2017).

The key idea with this approach is that some memes are 'good' and contribute positively to an individual's performance, while others are 'bad' and contribute negatively, and there will be a range of performance contributions that may depend on the environment and hence change with it. The aim of each new-born individual is to acquire the best subset of the available memes they can to perform well in their given environment by whichever learning strategy evolution has given them. The key evolutionary outcomes of the simulations are the associated evolvable parameters, namely the separate social and individual learning rates, and the individual performance levels that are simply the sum of the performance contributions of the memes they have acquired. The social learning rate specifies how quickly memes can be acquired from other individuals, and the individual learning rate specifies how quickly memes can be acquired by direct interaction with the environment. These two parameters summarize the effect of an enormous number of details, such as how the various types of learning take place in practice and how they depend on the species' associated physical characteristics, but that abstraction proves to be a useful simplification for the current study (Bullinaria, 2017, 2018). So, to summarize: each individual agent's record simply consists of their two innate leaning rates, their age, and a list of learned memes.

2.3 Implementing the learning processes

The simplest workable implementation of this approach (Bullinaria, 2017) has a fixed set of M learnable memes that each only need to be specified by their unique identification number and current performance contribution. Then each individual i can learn up to $\alpha_i \phi B_i$ memes each time-step by some form of social interaction with other individuals, where α_i is the evolved innate social learning rate in the range [0, 1] representing the degree to which social learning is adopted, ϕ specifies the maximum acquisition rate for the given species, and the effective brain capacity B_i is included to establish a scale independence for the learning rate. Similarly, each individual i can learn $\delta_i \psi B_i$ random good memes each time-step directly from their environment, where δ_i is the evolved individual learning rate, and ψ is a measure of ease of direct learning for the given species. The social learning may be limited by the number of different memes that are carried by the population at the given time, but the only limit on what individuals can learn directly from their environment is their brain capacity B_i. Both forms of learning also involve a further parameter that acts to limit the associated evolved learning rate. For the social learning there is a species dependent transmission fidelity f such that a fraction $1-f$ of copied good memes are copied incorrectly. For the direct individual learning, the limiting factor is that learning too quickly introduces errors, for which the simplest meme-based implementation has a probability $\rho\delta_i$ of learning a bad rather than good meme, where ρ is a species-dependent measure of direct learning difficulty. In more detailed evolutionary simulations of the emergence of different species, the parameters B, ϕ, ψ, f and ρ would summarize the effect of a number of associated physical characteristics and be co-evolved along with the learning rates (Bullinaria, 2017), but for the purposes of the current study it will be sufficient to fix them at specified values throughout the simulations to produce outcomes for one particular existing species that has the potential for both individual and social learning.

The details of this approach have already been tested, explored and optimized (Bullinaria, 2017), including an investigation of the dependence on factors such as the choice of who is socially learned from, the transmission fidelity f of the social learning, the effective brain capacity B in relation to the total number of learnable memes M, and which bad memes are acquired when learning errors occur, and life history characteristics such as periods of parental protection and restrictions on when reproduction can take place. An important advantage built into this abstract approach is that useful simulations can be carried out without the need to specify any details about what exactly the memes represent, nor how exactly they are acquired by the two forms of learning (Bullinaria, 2017). Moreover, it is easy to include numerous additional evolvable parameters to specify further life history traits, including more details of the learning processes, such as which other individual are learned from (e.g., random, best performing, oldest, nearest, parents, etc.), or if, when and how the innate learning rates should be adapted during a lifetime, but such complications will be avoided for the purposes of this paper by fixing them at typically observed values throughout (Bullinaria, 2017). So, to summarize: each agent i has a record consisting of their social leaning rate α_i, their individual learning rate δ_i, their age, and a list of identification numbers specifying which memes they have learned so far, and that record is updated each time-step the agent survives by incrementing their age and adding up to $\alpha_i\phi B_i + \delta_i\psi B_i$ more learned memes.

2.4 Accommodating environmental variability

The general approach above was extended (Bullinaria, 2018) to explore how the simulation outcomes were affected by having random environmental variations that cause the performance contributions of the available memes to change over time. These changes could be explicitly temporal (e.g., due to world-wide climate change) or indirectly temporal (e.g., due to moving to a new geographical region). Another key advantage built into the abstract meme-based approach is that such environmental changes can be implemented very efficiently if the memes and their performance contributions are represented in a particular way, and this approach also has the advantage of allowing a clearer analysis of the outcomes (Bullinaria, 2018). If there is a simplified set of M acquirable memes that have their performance contributions spread uniformly through the range [–1, +1], the memes can be conveniently ranked and numbered so that meme number i has contribution $1–2i/M$. Then to implement an environmental change just means re-ranking the M memes and renumbering the set of memes held by each individual in line with that. The basic unit of change is a simple swap of a pair of memes i and $j<i$, corresponding to having a meme i that becomes relatively less useful after an environmental change swapping its ranking with a meme j that becomes relatively more useful. The "magnitude" of the change is then simply $i–j$ and individuals will have their performance change by 0, $\pm 2(i–j)/M$ or 0 depending on whether they hold neither, one or both of those memes.

Crucially, at no point is there any need to keep track of what each meme is, nor how it depends on the environment – all that is required is a simple meme number swap to represent each random environmental change. Clearly, a sequence of such swaps is a gross simplification of what happens to real populations in real changing environments, but the earlier study (Bullinaria, 2018) has explored and optimized the details of this approach and presented simulations which demonstrate that it is sufficient to give useful results with minimal effort, which is what renders it feasible to test a broad range of environmental change distributions in this paper. Moreover, there is no reason to suspect that applying these particular simplifications will interfere with the current aim, which is not to further refine the simulation approach, but to simply identify a counter-example related to the prevailing view of how the simulation outcomes are affected by the choice of environmental change distribution. So, to summarize: by aligning the meme identification numbers with their changing uniformly-distributed meme performance contributions, each individual random environmental change can be efficiently modeled by simply swapping a random pair of meme identification numbers wherever they appear in the agents' records.

3. Environmental Change Distributions

Real environmental variability over time is made up of a sequence of many individual random changes drawn from some distribution of magnitudes. Elucidating the effect of different environmental change distributions has received considerable attention over recent years. There is significant evidence that natural change distributions are well approximated by inverse power laws (Halley, 1996; Halley & Inchausti, 2004; Newman, 2005), with their exponent (or noise colour) depending on whether marine or terrestrial environments are involved (Steele, 1985; Vasseur & Yodzis, 2004). The relation of environmental noise colour to population dynamics, population collapse and extinction risk have been widely studied (e.g., Ripa & Lundberg, 1996; Petchey, Gonzalez & Wilson, 1997; Halley & Kunin, 1999; Morales, 1999; Heino, Ripa & Kaitala, 2000; Wichmann et al., 2005; Kamenev, Meerson & Shklovskii, 2008; Ruokolainen et al., 2009; Gilljam et al., 2019), and the investigation of coloured noise in artificial life models is ongoing (Grove, Borg & Polack, 2020; Grove et al., 2022).

Inverse power law distributions exist whenever the number of occurrences (or probability) of an object or event of a particular size varies inversely as a power of that size (Newman, 2005). The examples discussed in detail by Newman (2005) include several that relate (either directly or indirectly) to natural environmental changes, including the magnitudes of earthquakes, the diameters of comet and asteroid craters, and the intensities of solar flares. In the simulation framework outlined in the previous section, the random environmental changes (which may correspond to any number of related or unrelated, natural or agent-driven, events) are ultimately represented as swaps of the corresponding meme performance contributions, and their distribution is therefore defined in terms of how the number of times n a particular magnitude of change m occurs is related to that magnitude (i.e., the meme performance contribution swap distance). So, to model the inverse power-law distributions that have been found to correspond to natural environments, that requires having $n \sim 1/m^{\beta}$ for a particular noise colour β.

This is rather different to the earlier models (Bullinaria, 2018) which instead took a simpler approach based on white noise ($\beta = 0$) distributions with the individual swap magnitudes taken uniformly from particular restricted ranges. To check whether that simplification has affected the simulation outcomes in any significant way, and to investigate the effect of noise colour on the more realistic distributions, requires the earlier simple swap range limits to be removed and replaced by inverse power-law swap distributions of particular noise colours β. For a discrete set of $M = 500$ memes with uniformly distributed performance contributions, that gives the familiar linear relation in the log number versus log magnitude plots seen on the left of Figure 1. For each distribution, the minimum individual change corresponds to a swap between a pair of memes with adjacent performance contributions, and the maximum individual change corresponds to the swap between the pair of memes with the best and worst performance contributions. So the distributions all have well-defined finite means that can be used for any normalization purposes that may be required. Newman (2005) presents similar-looking graphs for a range of real-world examples, and also provides a good explanation of how inverse power laws can be identified in practice and the mathematics underlying them. The right graph in Figure 1 shows how the distributions in the left graph compare with the earlier restricted range white-noise distributions of equivalent average magnitudes, in which the beta=0 and the maximum range=500 cases are identical.

A complication with the earlier simulations (Bullinaria, 2018) was how to best avoid the end effects inherent in having the relatively small number M of available memes (only 500) necessary for the simulations to run reasonably quickly. Given the meme-centered approach used in that study, it made sense there to start by picking a uniformly-random meme to swap, compute the restricted set of memes it could swap with to give an acceptable magnitude of change, and then uniformly-randomly pick one of them. However, because of the limited number of memes, the actual ranges of swaps available were sometimes less than the chosen "range size", and so the final distributions were not totally uniform and the average swap sizes were less than half the range size. That simplification was not important for the purposes of that investigation because it was only the variation of the outcomes over an abstract average magnitude scale that was of interest. Here, however, it is important to keep the distributions in the chosen form and have well-defined average magnitudes, so now the starting point needs to be choosing the actual swap distance from the required change-magnitude distribution, and then the key swapped meme is chosen uniformly-randomly from the possibilities allowed by the swap distance and meme set size. The first set of simulation results presented below will check whether this change alone leads to significantly different results.

Then the obvious next step is to go ahead and replace the restricted-range distributions in those simulations to determine whether inverse power-law distributions give different results, and to explore the effect of using different noise colours β, but there is a complication with that too. By definition, for lower β there will be a higher proportion of large magnitude swaps, so the average magnitude of change per swap will be larger, and the correlation of the original list of meme performance contributions with the resulting list will fall off more quickly with the number of individual swaps, as shown on the left of Figure 2. This graph follows exactly the same pattern as the corresponding plots showing how larger white-noise restricted ranges (that also have larger average magnitude of change) lead to faster fall-offs in the correlations, as shown on the right of Figure 2. Thus, while it is straightforward to repeat the earlier simulations (Bullinaria, 2018) using coloured noise distributions, the effects of colour and average magnitude may be confused. The second set of simulation results presented below will show exactly what happens in this case. Then, a third set of simulation results will be presented that uses a swap-frequency-based total magnitude of change normalization to disentangle the total change magnitude and noise colour effects in order to explore how the outcomes depend on each of them separately. Finally, a fourth set of simulations reveal how the same total magnitude normalization approach affects the restricted-range distribution results.

4. Simulation and Analysis Details

Apart from the environmental change distributions, all the simulations presented in this paper will follow exactly the same set-up and parameter values as optimized and specified in the earlier studies (Bullinaria, 2017, 2018). These were chosen there to keep the simulation run-times feasible and to roughly match the age distributions and learning timescales of real human populations with high fidelity social learning. The modeling framework described above was designed to be much more general than that, so, in principle, this study could have been based on a totally different species. However, there is an obvious advantage to using a model that has already been through a considerable degree of testing and optimization, and it is important that it is based on a species that is known to have a non-trivial dependence on environmental variability. For example, modeling a simpler species that has little or no social learning, or social learning of low fidelity, is unlikely to provide evolutionary outcomes of sufficient complexity to exhibit clear non-trivial effects of environmental noise (Bullinaria, 2017, 2018). Beyond that, there is no reason to suppose that the particular choice of settings will have a crucial effect on whether the current study finds an effect of noise distribution in the simulation outcomes.

Each simulation maintains a fixed population size of 200 individuals with one update cycle per simulated year, and every simulated year 10% of individuals die due to unsuccessful competition, and a random 20% of those aged over 60 years die of old age. Real populations are usually much larger, but the important issue here is that this provides sufficient diversity that bigger (and hence computationally more costly) populations are found to result in similar patterns of evolved parameter averages, though they do tend to be more robust against population collapse (Bullinaria, 2018). Individuals dying are replaced by children of pairs of the best performing survivors, with competition for survival and procreation both controlled by standard pair-wise tournament selection (Eiben & Smith, 2016). It is the need to compete with older individuals in this way that drives the evolution of faster learning, so simpler models with once-per-generation updates should be avoided (Bullinaria, 2009). The number of available memes M was set to be 500 and each individual is born with none and has fixed effective brain capacity B to learn up to 100 of them. The other learning parameters described above were fixed at $\phi = 0.1$, $\psi = 0.1$, $f = 1$ and $\rho = 0.04$ which leads to the evolved individual learning rates δ typically falling in the range [0, 2.5].

All the simulations are run for up to 10 million simulated years and carried out 20 times (with different random number seeds) to provide reliable statistics which are presented as means and standard deviations in all the graphs. The populations actually normally stabilize over much shorter timescales, but having very long runs maximizes the chances of identifying any slowly evolving features, and also provides better statistics for the population collapse times and evolved parameters (Bullinaria, 2018). Factors such as the learning rates and performance levels which vary during individual runs typically go through three phases: first their initial evolution from zero or very low random values towards near-optimal values, then a "stable" period during which they only vary in line with randomness in the learning and procreation processes and environment, and finally (but not always) a phase where the population is on the verge of collapse and large atypical variations may occur. These phases have different lengths that can vary quite considerably across runs, so to provide a reliable representation of the evolved stable state in each case, these values are averaged over 30 points spread evenly over the middle 20% of each run, which has previously been found (Bullinaria, 2018) to be sufficient to avoid the start and end effects.

The initial populations have zero social learning rates to avoid any bias towards social learning, and small random individual learning rates drawn from the uniform distribution [0, 0.1] to provide an initial set of memes that any social learning can work with. These then evolve throughout each run via crossover that has each child's innate parameter values chosen randomly from the ranges spanned by their two parents, with mutations adding random constants so the inherited values can fall outside the parental ranges. It has been found previously (Bullinaria, 2018) that simulations with mutations that are too small can experience population collapse before the key evolving parameters reach their near-optimal stable states, and mutations that are too large can also slow the evolution, so some care is needed in choosing the mutation rates. A related complication is that mutations for parameters that lie within a uniform fixed range (e.g., [0, 1] for the social learning rates) are best treated differently from those which naturally range over many orders of magnitude and are only restricted to be positive (e.g., the individual learning rates). Fortunately, good mutation details for a wide class of models, which includes all those of this paper, have already been established for the simulations of Bullinaria (2009, 2017, 2018): the social learning rate mutations are direct uniformly random additions from the range [–0.02, 0.02] with the final rates capped to lie within their fixed range [0, 1], while the individual learning rate mutations take place on a natural logarithm scale with randomly chosen additions from a Gaussian distribution of zero mean and variance 0.02. One of the reasons for running each simulation for such a large number of simulated years, and only using the middle portion of each run to compute averages, is to make sure that reliable results are obtained across all the simulations even if the mutation details are slightly sub-optimal in some cases.

Population collapse corresponds to a sudden drop in average population performance to a level that is not survivable in the given environment (e.g., due to inability to acquire sufficient food, or failure to cope with changed climate or competing species). The precise performance cutoff level has previously been found to have little effect on the results (Bullinaria, 2018**)** and is taken to be 20 on the abstract scale used here.

5. Models With Restricted-Range White-Noise Distributions

To begin, the effect of using the new distribution-centered approach for the environmental changes needs to be checked, and this also provides a base-line against which to compare the other models. Figure 3 presents the new population-average results as a function of average change magnitude for the case of one restricted-range change per simulated year. As anticipated, all the evolutionary outcomes are in broad agreement with the corresponding results of the earlier (Bullinaria, 2018) study (e.g., the Figure 3 plots here are in line with the Random and Best plots in Figure 10 there), though the average change magnitudes here are not simply related to the change range magnitudes used there.

For the standard case of social learning from the best performing other individuals (Best), the average individual performance levels not surprisingly fall as the environmental change magnitudes increase. At low and intermediate levels of change, high levels of social learning alone evolve. However, when there are higher levels of change, evolution still leads to high levels of social learning, but it is then accompanied by significant levels of direct individual learning which allows individuals to cope better with the changes. Only when the levels of change and individual learning are highest, do the social learning rates begin to drop. Interestingly, for increasing intermediate levels of environmental change, evolution leads to social learning dominating to the extent that adverse changes more frequently result in a sudden drop in average performance leading to shorter times to population collapse, which beyond a certain point are increasingly mitigated by the adoption of individual learning.

When the social learning is from random other individuals (Random), that learning is less effective, so significant levels of individual learning evolve even for low change magnitudes, and the performance advantage over that range is lost. However, this learning strategy has the advantage that it almost totally avoids the problem of population collapse.

Considering all these cases together confirms that the chosen simulation set-up is providing a sufficiently complex pattern of outcomes that can reasonably be expected to highlight any effects of the environmental noise distribution. Further learning strategies and their evolution, together with details of the evolutionary time-courses, dependences on transmission fidelities, and associated distributions of learned memes were investigated in the earlier study (Bullinaria, 2018), and are certainly interesting, but for the purposes of the current study the standard cases described above should suffice.

6. Models With Coloured-Noise Distributions

As discussed above, it is widely believed that the variations in real environments are best described by inverse power-law distributions rather than by the kind of simpler restricted ranges of white noise simulated in the last section and the earlier study (Bullinaria, 2018). To see how modeling those more realistic environmental changes affect the evolutionary outcomes requires re-running all the simulations using inverse power-law change distributions for a range of different noise colours β, but with everything else the same. Again with one environmental change per simulated year for the cases of high fidelity social leaning from random and best performing other individuals, that leads to the β dependent outcomes shown in Figure 4. The results for the $\beta = 0$ cases here are the same as for the maximum magnitude cases in the restricted-range models of Figure 3, because those simulations are identical. What is maybe less expected is how closely the pattern of results for $\beta > 0$ follow the equivalent plots of Figure 3 but with the *x*-axis reversed.

By definition, the whiter the noise is (i.e., the lower β is), the more high magnitude changes will be present, and that will obviously result in a higher average magnitude of change. That difference is certainly evident in the meme-performance correlation plots on the left of Figure 2. It is feasible that this fact alone could be the reason why increasing β in Figure 4 should match the effect of decreasing the average magnitude in Figure 3. One easy way to explore whether Figure 4 may simply be showing an average-magnitude effect, rather than a direct noise-colour effect, would be to adjust the average total magnitude of change per simulated year for each colour by increasing or decreasing the number of individual changes that occur each update cycle. Re-running all the simulations with ten changes per simulated year, instead of one, increases the average total magnitude of change each year by a factor of ten. That leads to very similar plots to Figure 4, but with the *x*-axis shifted, so the new plots over the range $\beta = 0$ to $\beta = 3$ look like the Figure 4 plots over the range $\beta = 1$ to $\beta = 4$. And re-running all the simulations with one hundred changes per simulated year leads to the plots shifting even further to the right. Having the effects of colour and magnitude interact in this way is clearly problematic, and, to get a clearer understanding of the effect of noise colour alone, a more careful investigation is required.

7. Models With Normalized Coloured-Noise Distributions

Varying the number of meme-swaps per simulated year is a simple factor that affects the average total magnitude of change, and, as just noted, increasing or decreasing that swap frequency results in the pattern of outcomes in Figure 4 shifting to higher or lower β values. Conveniently, there is no reason why the environmental change frequency cannot be set differently for each change colour. That means the average total change magnitude per simulated year can be set to any value for any change colour β, simply by setting the number of meme-swaps per simulated year appropriately, and that may be expected to be sufficient to provide an effective way to separate the effects of total magnitude and colour.

The first complication is that, although simply multiplying the average magnitude of a single swap by the required number of swaps per simulated year can bring the average total magnitudes into line across colours, that does not quite bring the meme performance contribution correlations into line too. For example, the left graph of Figure 5 shows the correlations resulting from applying the set of swap-frequency adjustments required to produce one particular average yearly total magnitude (namely, 250) for all colours β, with the same pattern being easily moveable to higher or lower numbers of years simply by applying the same proportional increase or decrease in the number of swaps per simulated year across all colours. The colour-based differences are now certainly a lot less than those seen in Figure 2 for uniform swap frequencies, but there remain some effects of colour because many small swaps for high β will not simply add up to give a bigger swap, but will instead perform a random walk with less overall effect. The consequence of re-running all the earlier simulations with the swap frequencies set in this way is that the effect of colour still ends up looking like a magnitude effect.

A potentially better way to proceed would be to adjust the number of swaps per simulated year for each colour by the necessary amount to bring the correlations themselves into line across colours, with the total environmental change magnitudes then corresponding to how quickly the correlations fall. That is straightforward to do in the models used above simply by reading off from plots like those on the left of Figure 2 the proportion of extra swaps required to bring the $\beta = 0$ correlation curve into line with each of the higher-β curves of interest, and increasing the swap frequency for each of those higher-β cases by the corresponding amount. Then the effect of varying the average total change magnitude per simulated year can easily be explored by imposing a further uniform scaling of the swap frequencies across all colours. The only remaining complication is that the required numbers of swaps per simulation update cycle are rarely integer, so a stochastic process is required to give the required numbers on average. The earlier simulations could thus be repeated with β–dependent swap frequencies set to bring all the correlations into line with the $\beta = 0$ case, and correlation-based average total change magnitudes defined to be the actual average total change magnitudes associated with the $\beta = 0$ case. There are then no significant differences in the correlation plots across colours, as shown in the right graph of Figure 5, and the inverse-power-law plots of Figure 1 are simply shifted vertically to those shown in Figure 6.

The new key simulation results for the standard case involving high fidelity social learning from the best performing other individuals are shown in Figure 7 for four representative environmental change colours β, for comparison with Figures 3 and 4. It is now evident that, after removing the effect of correlation-based total change magnitude, very little dependence on the change colour remains. As in the earlier study (Bullinaria, 2018) and the simulations presented above, the evolutionary outcomes fall into three broad categories with smooth transitions between them:

- For low average total change magnitudes (< ~0.02), the evolved social learning rates are near maximal (~1) and the individual learning rates are very low (~0) for all change colours β, and the environmental changes are insufficient to cause a high likelihood of population collapse for any β. This is consistent with what was found previously, though, with the new correlation-based definition of total change magnitude used here, what now constitutes 'low' (i.e., magnitude ~0.01) is a lot less than that (magnitude ~1) in the earlier study (Bullinaria, 2018).
- For intermediate average total change magnitudes (~0.1 to ~10), the evolved social learning rates are still near maximal (~1) and the individual learning rate are still very low (~0) for all change colours β, but the environmental changes are now sufficient for the reliance on social learning to result in rapid population collapses for all β. This situation is consistent with that found for intermediate magnitudes (of ~15 to ~60) in the earlier study (Bullinaria, 2018).
- For high average total change magnitudes (> ~100), the evolved individual learning rates are high (~2) for all change colours β, which provides sufficient adaptability to avoid population collapses, again as found in the earlier study (Bullinaria, 2018). The social learning rates remain significant, but not near maximal, for all β, increasing from ~0.4 for $\beta = 0$ to ~0.8 for $\beta = 3$, which is consistent with the earlier results (Bullinaria, 2018) and the idea that for high magnitudes of change the social learning becomes less reliable the whiter the noise is.

For all total change magnitudes, the average performance levels have little dependence on colour, and get worse with increasing magnitude as expected.

8. Models With Normalized Restricted Range White-Noise Distributions

While the simulation outcomes shown in Figure 7 for the high average total change magnitude cases with correlation-based normalized coloured noise distributions are broadly in line with the corresponding restricted-range distribution results shown in Figure 3, the 'intermediate' and 'low' average total change magnitude regimes appear shifted to much lower average magnitudes. The details of the normalized inverse power law individual-change distributions shown in Figure 6 suggest why this may be happening. To get the correlation-based average total magnitudes to line up across colours requires quite considerable adjustments to the number of changes per simulated year for the high β cases. This means that even for high β there are still as many fairly high magnitude swaps (~250) as for the maximum magnitude restricted-range case, and these were not present at all for the lower magnitude restricted-range cases.

To fully understand this issue, yet another set of simulations are required to explore how the presence or absence of large magnitude individual changes affect the outcomes. This is straightforward to do in the current framework because the various restricted-range distributions can also equally easily be normalized by adjusting the number of environmental change swaps per simulated year to give matching correlation-based average total magnitudes. Rerunning all the simulations doing that for various restricted ranges for the individual white-noise change magnitudes leads to very similar outcomes across range sizes, shown for the maximum, minimum and intermediate ranges in Figure 8. These outcomes are also remarkably similar to those in Figure 7, so it seems that the apparently shifted low and intermediate change regimes are simply due to the differences in the total change magnitude scale definition rather than anything to do with the presence of large magnitude individual changes.

9. Conclusions and Discussion

The overall conclusion from the simulations presented above is that there is at least one class of artificial life models (namely the evolutionary agent-based learning models upon which they are based) for which choosing realistic degrees of environmental change is important for achieving results that provide a reliable account of the natural system in question, but it is the average total magnitude of change per simulated year that is most important, rather than the distribution of individual changes. Simulations using more realistic inverse power-law distributions have been found to give evolutionary outcomes that are only slightly different quantitatively from those using simpler restricted-range white-noise distributions, and the qualitative patterns of results are remarkably similar. Moreover, when the average total change magnitude is controlled for, the particular noise colour of the inverse power-law distributions, and the particular range size for the restricted-range distributions, have very little effect on what evolves.

Perhaps the most important finding is how easily the effect of noise colour for the inverse power-law distributions can be totally swamped by, and confused with, the effect of total noise magnitude. In particular, the simulations here have shown how the obvious environmental change frequency-based normalization approach based on simple magnitude averaging fails to fully separate the colour and total change magnitude effects, but a similar frequency-based normalization approach based on temporal correlations can. Only once the environmental change colour and total change magnitude effects have been separated in this way could the simulation results presented show clearly how the total magnitude is by far the most important factor determining the evolutionary outcomes, but the colour does still have a small effect on some aspects, particularly in the transitions between outcome categories and for very high magnitude changes. The same pattern of evolutionary outcomes could then be shown to exist across the range sizes if the inverse power-law distributions are replaced by simpler restricted-range white-noise distributions.

Another advantage of separating the total environmental change magnitude per unit time from the noise distribution colour or restricted range is that first approximations of various timescale or pace-of-life effects (Mathot & Frankenhuis, 2018) can be obtained easily without needing to change or re-run any of the simulations. For example, a species with lifespans that are shorter than the original species by a factor of two can be modeled by the same simulations simply by redefining the simulated time per simulation update cycle to be half as long. Then, without changing the number of environmental changes per update cycle, the average total magnitude of change per simulated year will be doubled, and the evolutionary outcomes can simply be read off from the new points of the same graphs, such as those in Figures 7 and 8. Of course, accurately switching between species will also require adjusting various other parameters too, such as those underlying the learning, reproductive and death processes, but the broad pattern of how the environmental changes interact with life-spans will be evident.

The same models can also be applied to another aspect of environmental variability that this paper has not yet considered, namely the issue of environmental change persistence or memory. So far, all the changes have persisted and accumulated over time resulting in the environment performing a kind of random walk, with the noise building up and reducing the performance contribution correlations as seen in Figures 2 and 5. This is a reasonable approximation for many real environments, but there also exist natural systems that automatically revert back to earlier states after being subject to a change event (for example, the atmosphere after a meteorite impact). In the extreme case of this, each random change is rapidly replaced by a new random change. This aspect of variability was the focus of Halley (1996) and Grove et al. (2022), and inverse power laws appear here in a different way to that discussed above. By treating the environmental variability as a randomly changing noise signal laid upon a stable (or predictable) underlying state, the noise can be represented as a time series with a particular degree of autocorrelation and analyzed using spectral density decomposition. It is found that such power spectra are well approximated by inverse power laws, with the extreme persistent-change random-walk case corresponding to red/Brown noise (exponent 2) and the opposite extreme case of rapid-decay repeated-new-change corresponding to white noise (exponent 0). Vasseur and Yodzis (2004) have surveyed the different spectral exponents found across a wide range of natural environmental variables. In the framework of the models discussed in this paper, changes lacking persistence can easily be implemented by adding change reversals at appropriate points in the sequence of random changes. The models then run in exactly the same way as above, but the performance contribution correlations will clearly fall more slowly if earlier changes are eventually reversed. The difficulty in this case is that setting the average total change magnitude per unit time is not so straightforward. Given the speed of lifetime learning, many change reversals will come after the change has already been accommodated by learning, and will hence be as disruptive as a totally random new change. So the evolutionary outcomes will be the same as in the models above, but for a higher-than-actual effective correlation-based average magnitude. If the environmental changes do come and go so fast that learning cannot deal with them, then the learning models of this paper are not relevant. In that case, one of the non-lifetime-learning strategies will be needed to deal with them, such as maintaining a diverse population so that there are always some individuals that are able to cope whatever environmental change may occur, or having all individuals develop a tolerance to the changing environmental conditions, but these approaches will rarely reach the performance levels achievable when lifetime learning is able to track the changes. So here, and also more generally, whether it is appropriate to build models based on the average magnitude of the whole noise time series, or something closer to that of the individual changes, will depend on the timescales of the noise events and the processes by which the changes are being accommodated by the modeled species. A final complication is that particularly intelligent species may be able to learn to recognize which changes will eventually reverse and on what timescales, and hence learn how best to adjust their behaviours to deal with them, but simulating patterns of learned behaviour of such complexity will clearly require more sophisticated models than those discussed in this paper.

The practical consequence of all this is that, for these artificial life simulations with changing environments to be reliable, it will be important to have a good idea of what the average total environmental change magnitude should be in relation to the natural time-scales of the modeled species, as that is the main determiner of the outcomes. However, this study has found that the other aspects of the change distribution have surprisingly little effect, meaning that choosing the right colour for an inverse power-law distribution, or the right restricted range for a white-noise distribution, may not be as important as previously thought. Knowing that there is just one key parameter affecting the broad pattern of results, rather than many, will certainly make it easier to carry out such simulations.

Finally, it is important to remember that all the results presented in this paper have been based on only one particular agent-based learning model, for only a small portion of its underlying parameter space, and for only the principal evolutionary outcomes, and it is possible (maybe even likely) that different conclusions will arise from alternative models. As highlighted in the introduction, it was never going to be feasible to test the effect of environmental noise distributions on all possible artificial life models. The original aim of this study was to explore this matter using one existing tried and tested model with parameter settings that had been optimized in the past for unrelated purposes. However, instead of finding interesting dependences on the colour of the inverse power-law noise distributions, it was found that the expected dependences virtually disappear when the average total environmental change magnitudes were properly controlled for. This study therefore provides a robust counter-example to the idea that using the right environmental noise distribution is always important for artificial life models. It is hoped that others will now take on board the issues outlined in this paper and apply the key ideas to their own artificial life models, such as separating out the total change magnitude effect using correlation-based normalization of coloured inverse power-law or white restricted-range noise models of environmental change, and hence explore how general the results are and what kinds of models *do* exhibit interesting environmental change distribution effects independent of the average total magnitude.

Acknowledgement

I am grateful to James Borg for suggesting that the effect of environmental noise distributions could be fruitfully investigated using my models.

References

Bartoli, A., Catto, M., De Lorenzo, A., Medvet, E. & Talamini, J. (2020). Mechanisms of social learning in evolved artificial life. *Proceedings of the 2020 Conference on Artificial Life*, 190-198.

Borenstein, E., Feldman, M. W. & Aoki, K. (2008). Evolution of learning in fluctuating environments: When selection favors both social and exploratory individual learning. *Evolution*, *62*, 586-602.

Borg, J. & Channon, A. D. (2012). Testing the variability selection hypothesis: The adoption of social learning in increasingly variable environments. *Proceedings of the 2012 Conference on Artificial Life*, 317-324.

Botero, C. A., Weissing, F. J., Wright, J. & Rubenstein, D. R. (2015). Evolutionary tipping points in the capacity to adapt to environmental change. *Proceedings of the National Academy of Sciences USA*, *112*, 184-189.

Boyd, R., Richerson, P. J. & Henrich, J. (2011) The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences USA*, *108*, 10918-10925.

Bullinaria, J. A. (2009). Lifetime learning as a factor in life history evolution. *Artificial Life*, *15*, 389-409.

Bullinaria, J. A. (2017). Imitative and direct learning as interacting factors in life history evolution. *Artificial Life*, *23*, 375-405.

Bullinaria, J. A. (2018). Evolution of learning strategies in changing environments. *Cognitive Systems Research*, *52*, 429-449.

Diamond, J. (2011). *Collapse: How societies choose to fail or succeed*. New York, NY: Penguin.

Eiben, A. E. & Smith, J. E. (2016). *Introduction to Evolutionary Computing*. Berlin, Germany: Springer.

Feldman, M. W., Aoki, K. & Kumm, J. (1996). Individual versus social learning: Evolutionary analysis in a fluctuating environment. *Anthropological Science*, *104*, 209-232.

Gilljam, D., Knape, J., Lindén, A., Mugabo, M., Sait, S. M. & Fowler, M. S. (2019). The colour of environmental fluctuations associated with terrestrial animal population dynamics. *Global Ecology and Biogeography*, *28*, 118-130.

Grove, M. (2011). Speciation, diversity, and Mode 1 technologies: The impact of variability selection. *Journal of Human Evolution*, *61*, 306-319.

Grove, M., Borg, J. M. & Polack, F. (2020). Coloured noise time series as appropriate models for environmental variation in artificial evolutionary systems. *Proceedings of the 2020 Conference on Artificial Life,* 292-299.

Grove, M., Timbrell, L., Jolley, B., Polack, F. & Borg, J. M. (2022). The importance of noise colour in simulations of evolutionary systems. *Artificial Life,* to appear.

Haaland, T. R., Wright, J. & Ratikainen, I. I. (2020). Generalists versus specialists in fluctuating environments: A bet-hedging perspective. *Oikos*, *129*, 879-890.

Halley, J. M. (1996). Ecology, evolution and 1/f noise. *Trends in Ecology and Evolution*, *11*, 33-37.

Halley, J. M. & Inchausti, P. (2004). The increasing importance of 1/f-noises as models of ecological variability. *Fluctuation and Noise Letters*, *4*, R1-R26.

Halley, J. M. & Kunin, W. E. (1999). Extinction risk and the 1/f family of noise models. *Theoretical Population Biology*, *56*, 215-230.

Heino, M., Ripa, J. & Kaitala, V. (2000). Extinction risk under coloured environmental noise. *Ecography*, *23*, 177-184.

Higgs, P. G. (2000). The mimetic transition: A simulation study of the evolution of learning by imitation. *Proceedings of the Royal Society B: Biological Sciences, 267*, 1355-1361.

Kameda, T. & Nakanishi, D. (2002). Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior*, *23*, 373-393.

Kamenev, A., Meerson, B. & Shklovskii, B. (2008). How colored environmental noise affects population extinction. *Physical Review Letters*, *101*, 268103.

Mathot, K. J. & Frankenhuis, W. E. (2018). Models of pace-of-life syndromes (POLS): A systematic review. *Behavioral Ecology and Sociobiology, 72*, 41.

Morales, J. M. (1999). Viability in a pink environment: why “white noise” models can be dangerous. *Ecology Letters*, *2*, 228-232.

Morgan, T. J., Suchow, J. W. & Griffiths, T. L. (2020). Experimental evolutionary simulations of learning, memory and life history. *Philosophical Transactions of the Royal Society, B375*, 20190504.

Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, *46*, 323-351.

Petchey, O. L., Gonzalez, A. & Wilson, H. B. (1997). Effects on population persistence: The interaction between environmental noise colour, intraspecific competition and space. *Proceedings of the Royal Society of London B*, *264*, 1841-1847.

Potts, R. (1996). Evolution and climate variability. *Science*, *273*, 922-923.

Potts, R. (1998). Variability selection in hominid evolution. *Evolutionary Anthropology, 7*, 81-96.

Potts, R. (2013). Hominin evolution in settings of strong environmental variability. *Quaternary Science Reviews*, *73*, 1-13.

Ripa, J. & Lundberg, P. (1996). Noise colour and the risk of population extinctions. *Proceedings of the Royal Society of London B*, *263*, 1751-1753.

Ruokolainen, L., Linden, A., Kaitala, V. & Fowler, M. S. (2009). Ecological and evolutionary dynamics under coloured environmental variation. *Trends in Ecology and Evolution*, *24*, 555-563.

Smaldino, P. E., Newson, L., Schank, J. C. & Richerson, P. J. (2013). Simulating the evolution of the human family: Cooperative breeding increases in harsh environments. *PLOS ONE, 8(11),* e80753.

Steele, J. H. (1985). A comparison of terrestrial and marine ecological systems. *Nature*, *313*, 355-358.

Thornton, A. & Raihani, N. J. (2008). The evolution of teaching. *Animal Behaviour, 75*, 1823-1836.

Vasseur, D. A. & Yodzis, P. (2004). The color of environmental noise. *Ecology*, *85*, 1146-1152.

Whitehead, H. (2007). Learning, climate and the evolution of cultural capacity. *Journal of Theoretical Biology*, *245*, 341-350.

Whitehead, H. & Richerson, P. J. (2009). The evolution of conformist social learning can cause population collapse in realistically variable environments. *Evolution and Human Behavior*, *30*, 261-273.

Wichmann, M. C., Johst, K., Schwager, M., Blasius, B. & Jeltsch, F. (2005). Extinction risk, coloured noise and the scaling of variance. *Theoretical Population Biology*, *68*, 29-40.

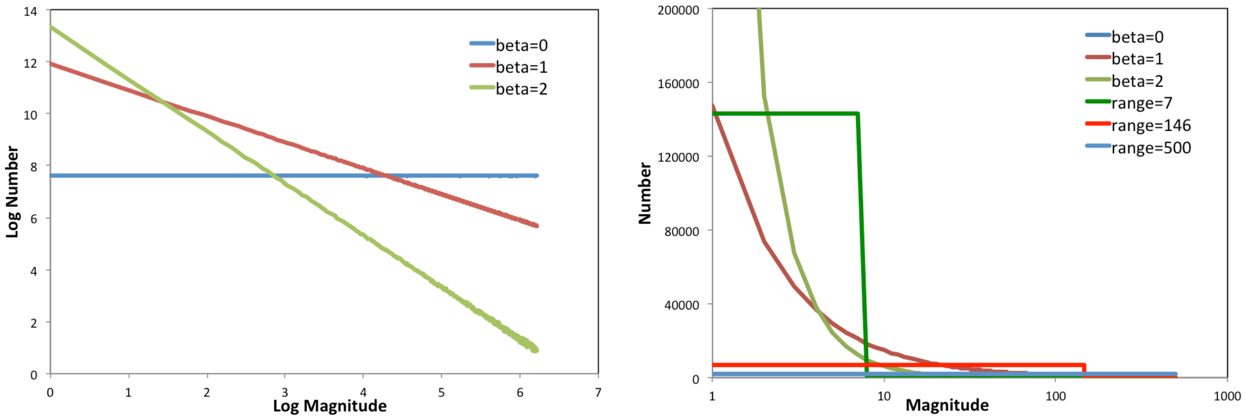

Figure 1. Number of individual environmental changes as a function of their magnitude for three representative inverse power-law distribution colours β (left), together with restricted-range white-noise distributions of the same average magnitudes (right), showing means from 100 runs of 1000000 swaps.

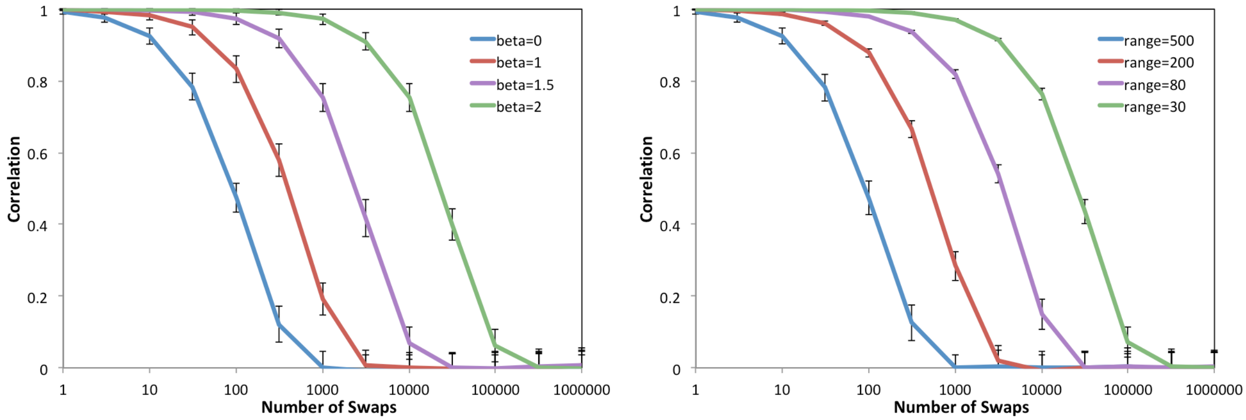

Figure 2. Pearson correlations of the initial and resulting meme performance contributions as a function of the number of individual environment changes, for each of four representative inverse power-law change colours β (left), and four representative restricted white-noise ranges (right), showing means and standard deviations from 100 runs.

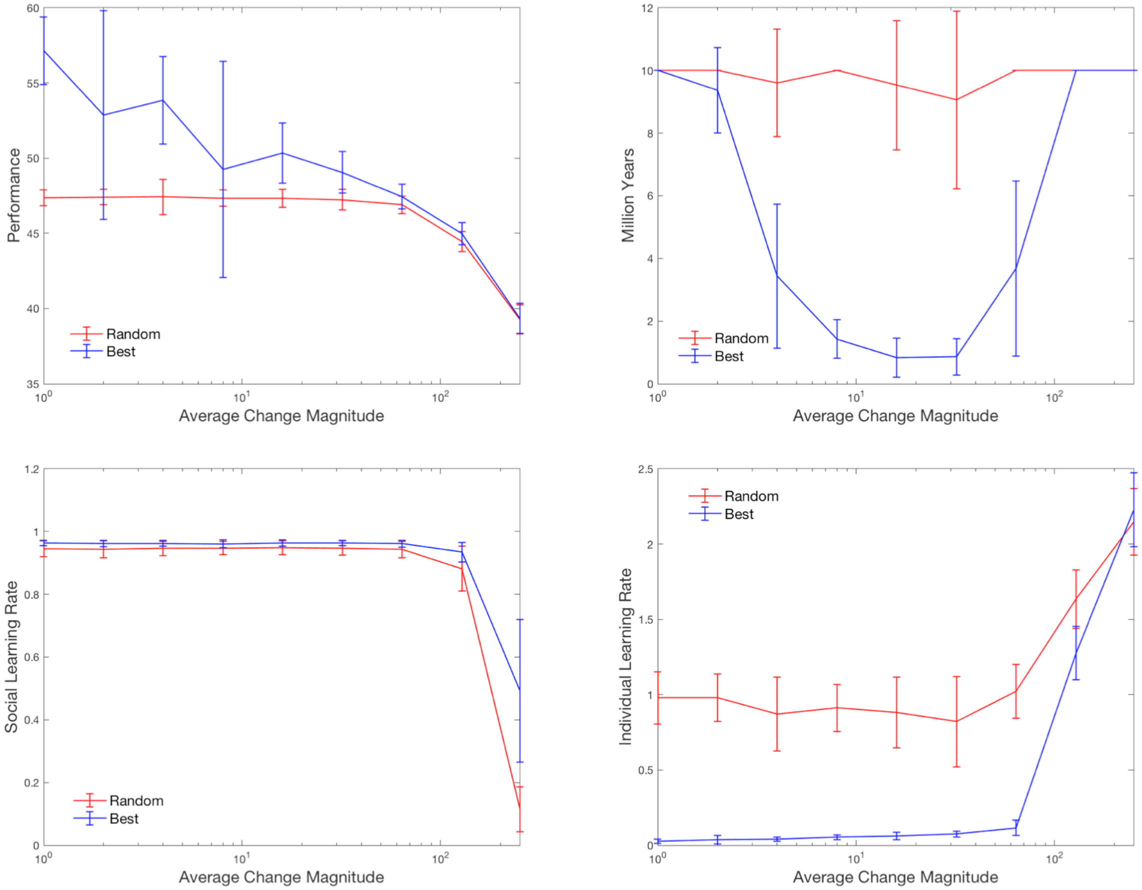

Figure 3. Effect of the restricted-range white-noise environmental change magnitude on the average performance (top left), times to population collapse (top right), social learning rates (bottom left), and individual learning rates (bottom right), for one environmental change per simulated year and high fidelity social leaning from random and best performing other individuals.

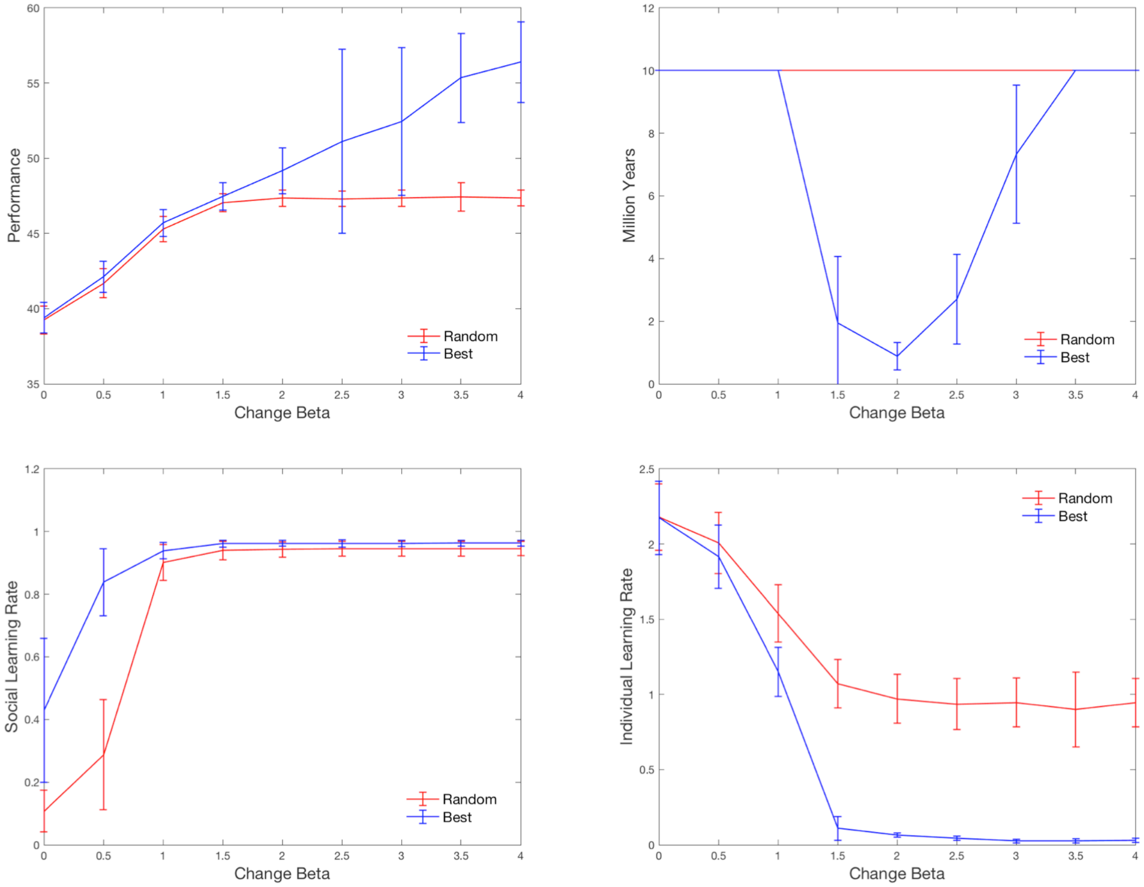

Figure 4. Effect of the inverse power-law environmental change colour β on the average performance (top left), times to population collapse (top right), social learning rates (bottom left), and individual learning rates (bottom right), for one environmental change per simulated year and high fidelity social leaning from random and best performing other individuals.

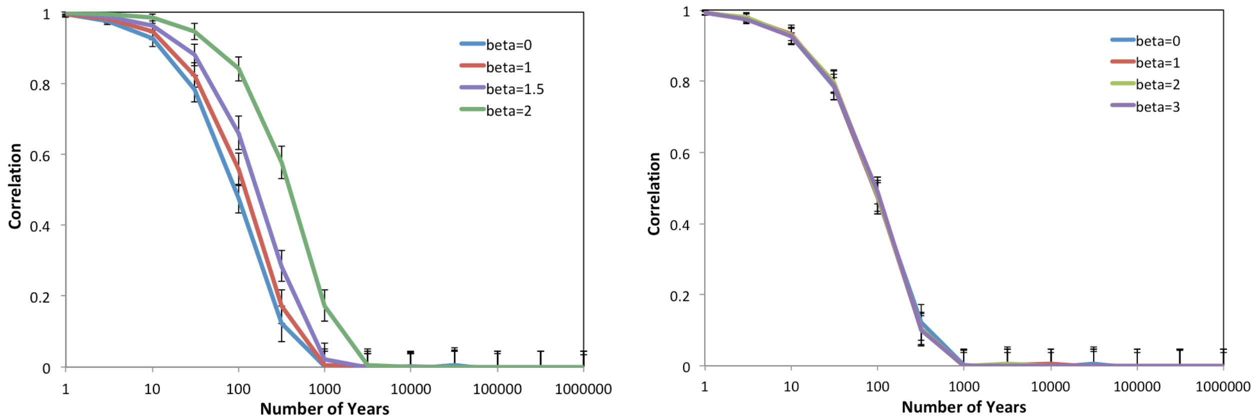

Figure 5. Pearson correlation of the initial and resulting meme performance contributions as a function of the number of years of environmental change, with fixed average total change magnitude per simulated year (left) and β dependent numbers of changes per simulated year chosen to normalize the correlations (right), each for four representative inverse power-law change colours β.

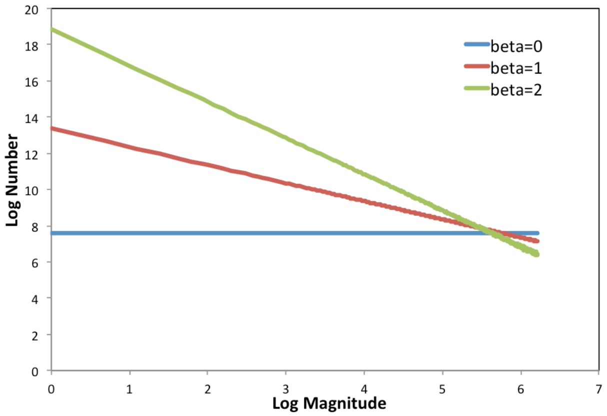

Figure 6. Inverse power-law distributions of Figure 1 after the number of environmental changes per simulated year have been adjusted to bring the meme performance contribution correlation curves of Figures 2 and 5 totally into line, thus separating the effect of change colour β from the correlation-based average total change magnitude.

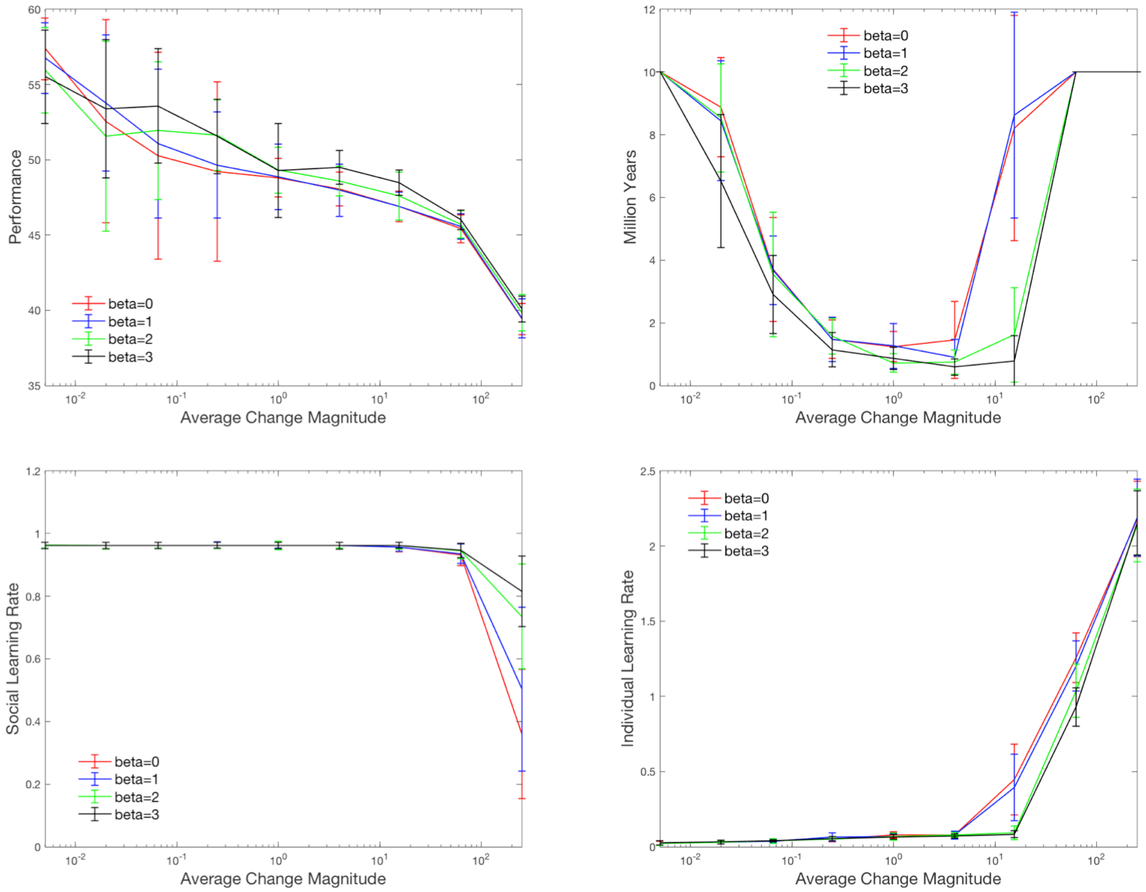

Figure 7. Effect of the correlation-based average total environmental change magnitude on the average performance (top left), times to population collapse (top right), social learning rates (bottom left), and individual learning rates (bottom right), for social leaning from the best performing other individuals, and the number of environmental changes per simulated year set to normalize the meme performance contribution correlations across four representative inverse power-law change colours β.

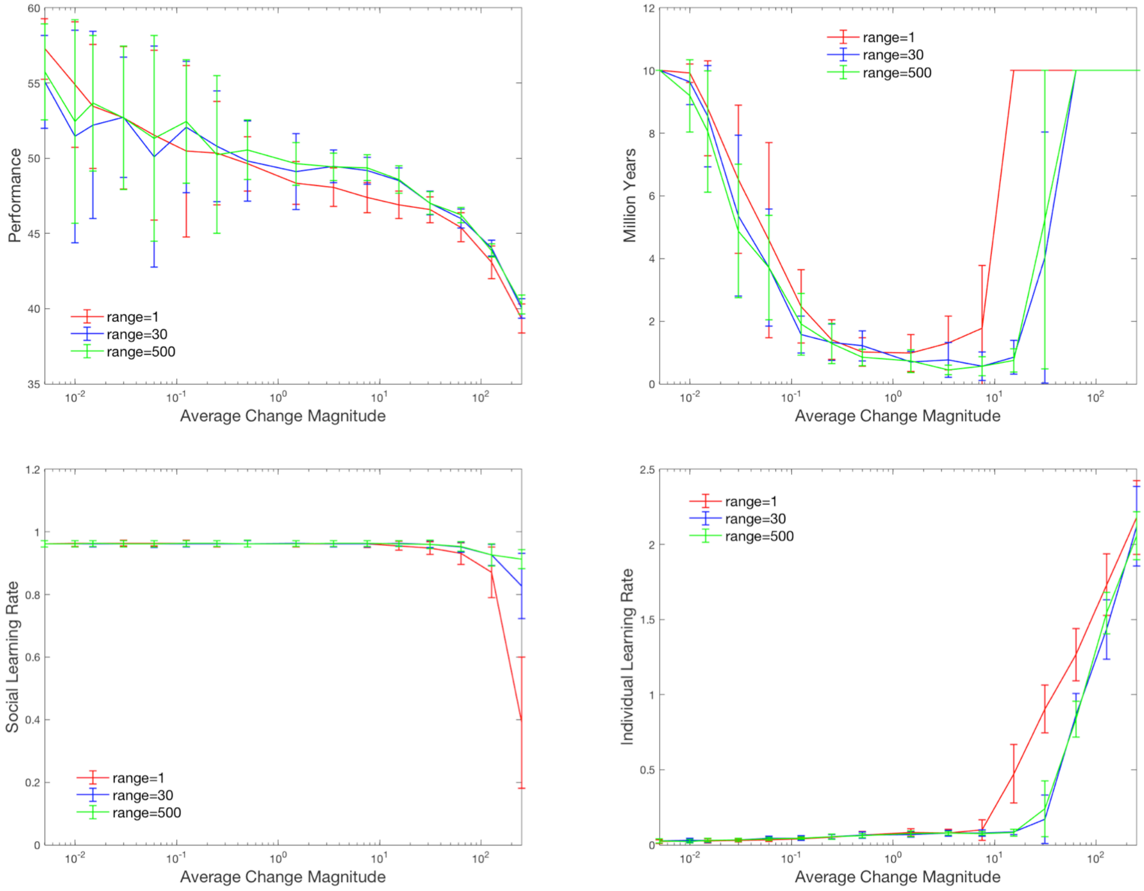

Figure 8. Effect of the correlation-based average total environmental change magnitude on the average performance (top left), times to population collapse (top right), social learning rates (bottom left), and individual learning rates (bottom right), for social leaning from the best performing other individuals, and the number of environmental changes per simulated year set to normalize the meme performance contribution correlations across three representative restricted ranges for the individual white-noise change magnitudes.