REvoSim: Organism-level simulation of macro- and microevolution

DOI:

[10.1111/pala.12420](https://doi.org/10.1111/pala.12420)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Garwood, R., Spencer, A. R. T., & Sutton, M. D. (2019). REvoSim: Organism-level simulation of macro- and microevolution. *Palaeontology*, 62(3), 339-355. <https://doi.org/10.1111/pala.12420>

Published in:

Palaeontology

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

REvoSim: Organism-level simulation of macro- and microevolution

Russell J. Garwood^{1,2*}

Alan R. T. Spencer^{2,3}

Mark D. Sutton^{3*}

¹School of Earth and Environmental Sciences, University of Manchester, Manchester M13 9PL, UK

²Earth Sciences Department, Natural History Museum, London SW7 5BD, UK

³Department of Earth Science and Engineering, South Kensington Campus, Imperial College London, SW7 2AZ, UK

*These authors have contributed equally to this work.

Abstract	2
Introduction	2
Simulation Model	4
Grid-cell iteration phase	4
Reproduction phase	5
Dispersal phase	5
Calculation of fitness	6
Species tracking	8
Model implementation	9
Software packages	9
Random Numbers	10
Visualization and Output	11
Population visualization modes	11
Population count	11
Mean fitness	11
Coding genome as colour	11
Non-coding genome as colour	11
Gene frequencies	11
Settles	11
Breed/Settle Fails (R=Breed; G=Settle)	11
Species	12
Output options	12
Results: Analysis, Validation, and Limitations	12
Limitations	14
An example simulation	15
Discussion	16
Acknowledgements	17
Data Archiving Statement	17
References	17
Figure captions and table	21

Abstract

Macroevolutionary processes dictate the generation and loss of biodiversity. Understanding them is a key challenge when interrogating the earth-life system in deep time. Model-based approaches can reveal important macroevolutionary patterns, and generate hypotheses on the underlying processes. Here we present and document a novel model called REvoSim (Rapid Evolutionary Simulator) coupled with a software implementation of this model. The latter is available here as both source code (C++/Qt, GNU General Public License), and as distributables for a variety of operating systems.

REvoSim is an individual-based model with a strong focus on computational efficiency. It can simulate populations of 10^5 – 10^7 digital organisms over geological timescales on a typical desktop computer, and incorporates spatial and temporal environmental variation, recombinant reproduction, mutation and dispersal. Whilst microevolutionary processes drive the model, macroevolutionary phenomena such as speciation and extinction emerge. We present results and analysis of the model focussing on validation, and note a number potential applications. REvoSim can serve as a multipurpose platform for studying both macro- and microevolution, and bridges this divide. It will be continually developed by the authors to expand its capabilities and hence its utility.

Introduction

The use of simplified numerical models is a well-established approach to studying complex natural systems. It provides: a means of investigating which aspects of complexity are important for the production of observed patterns; an experimental system in which the connections between system parameters and behaviour can be studied; and a means of predicting system behaviour in the past or future. Examples of the utility of this approach are widespread, and include climate modelling (see overview by McGuffie and Henderson-Sellers 2014), stellar evolution (e.g. Hansen et al. 2004), modelling of the Earth's internal processes (e.g. Turcotte and Schubert 2014), and far more. The processes of biological evolution are inherently algorithmic. Indeed, we note that in some instances the distinction between evolution and a simulation of the process is challenging to draw (Pennock 2007). Hence the processes lend themselves well to computational modelling. This is not a novel observation, and many studies have used various forms of evolutionary modelling to attempt to reproduce and explain patterns observed in the biosphere. These include studies of speciation mechanisms (e.g. Yamaguchi and Iwasa 2013), taxonomic patterns (e.g. Sigwart et al. 2018), ecological complexity (e.g. Laird et al. 2008), evolution of morphological complexity (e.g. Adami et al. 2000), kin-selection (e.g. Clune et al. 2010), and many others.

For palaeontologists, a primary interest is macroevolution, which occurs over geological time spans that are measured in millions of years. Observed macroevolutionary patterns have been well-studied and much debated, although their interpretation is complicated by the nature of the fossil record (e.g. incompleteness, stratigraphic averaging, etc). Microevolution — population-scale evolution taking place over much shorter periods of time — is, by contrast, better understood because it is amenable to direct observation over historical times (e.g. Tataronov and Avise 2007), experimental investigation (e.g. Good et al. 2017), and mathematical modelling (for an overview see Templeton 2006). Scaling microevolutionary theory into macroevolutionary timescales has long been considered problematic, as an epistemological gap exists where study is inhibited by the timescales involved, which are too long for experimental or observational studies, but too short to be resolved in the fossil record. This gap is where, for example, the processes of speciation in complex eukaryotes take place, and is one reason why these — and other broad evolutionary patterns — remain contentious and imperfectly understood.

Evolutionary models and the software packages that implement them vary in many details, but can usefully be considered in two categories. Firstly some operate at an organismal level, simulating populations of individuals whose reproductive efficiency is dependent on complex interactions with either each other, a non-biotic environment, or both (individual-based modelling, e.g. *Tierra*, *Avida*, *aevo*, the Tangled Nature model (Adami 1998; Batut et al. 2013; LaBar and Adami 2016; Arthur et al. 2017)). These have typically been used to simulate and study intraspecific microevolutionary processes, either because their focus does not lie on addressing macroevolutionary questions, or because computational speed limitations have made the simulation of large populations (> 200,000 individuals) over long timescales (millions of years) challenging. Secondly, more abstracted models exist that operate at a species or population level, simulating models of speciation rather than natural selection (Genieys et al. 2009) or modelling ecological interactions (Stauffer et al. 2005). These simulations are intrinsically macroevolutionary in scale, but require assumptions about emergent behavior of complex evolutionary systems. Such systems have a wide range of important applications, but do not bridge the epistemological gap between micro- and macroevolution.

As such, there are a wide range of potentially fruitful and important applications for an organismal level model, built using simple, abstracted principles, which can operate on macroevolutionary scales. Here we describe such a model: an organism-level simulation of sufficient computational efficiency that population sizes of 10^5 – 10^7 individuals can be studied over geological timescales. It was created specifically to bridge the gap between micro- and macroevolution, and also incorporates concepts of spatial relationships, spatial and temporal environmental variation, and recombinant reproduction. Isolation of gene-pools (i.e. speciation; see discussion below) emerges naturally within simulations under appropriate conditions, as do other patterns observed in real world systems. Both the model and its software implementation (Figure 1) are termed REvoSim (Rapid Evolutionary Simulator). Other individual-level models for macroevolution do exist, and show a great deal of promise in elucidating patterns (Costa et al. In press; Gavrillets et al. 2000; Aguiar et al. 2009; Rosindell et al. 2015), but generally run on shorter timescales with fewer individuals than used here.

REvoSim has been in development since 2008. We release it here (along with a utility program which generates environments for the software, EnviroGen) with the intention that it can be used as a multipurpose platform for the study of many evolutionary phenomena. While it was designed with macroevolutionary studies in mind, it is also applicable to microevolutionary problems. It is complementary to the many other approaches of studying evolution on a range of different timescales. The aim of this contribution is to introduce and document the model and software. Because the software has been designed for versatility, a large number of user-defined variables are provided and need to be defined (as opposed to providing hard-coded defaults). This allows REvoSim to be applied to a broad range of questions, and as such the overview provided here is also intended to allow user to make informed choices for simulation settings. We present results and analysis to validate the model using microevolutionary studies, and briefly discuss potential applications. These are diverse, and might include, for macroevolutionary studies: investigation of the preponderance of speciation modes under different conditions; the influence of rates of environmental change on the frequency of stasis; recovery from mass extinctions; and controls on biodiversity. Microevolutionary studies could incorporate the interaction of mutation, space, and population density, or the emergence Dobzhansky-Muller incompatibilities, for example. To maximise the usefulness and accessibility of our software, we release it both as source code and as ready-to-run builds for major operating systems. We intend development of the package as an ongoing process, and welcome bug reports, comments, and suggestions for future directions and features.

Simulation Model

REvoSim uses a simplified first-principles evolutionary model to facilitate high computational efficiency, enabling the simulation of large populations incorporating space, over geological time, using modest computer hardware. Speeds attainable depend on the computer hardware in use, the size of the populations simulated, and details of the experimental setup (most especially on whether species tracking and fitness recalculation are activated). With a typical 2018 desktop computer, speeds of between 500,000 and 1,000,000 iterations per hour can be achieved for populations of around 250,000.

The REvoSim simulation is abstract, but could be envisaged as representing a population of sessile invertebrates inhabiting a marine basin. Under this conceptualisation, the length of time represented by one iteration is of the order of one year. With a static (unchanging) environment and space turned on the simulation could be considered as a parallel of a petri dish of microorganisms, and with a dynamic environment and space disabled, it would be closer to a chemostat (in both cases an iteration would represent less time; minutes - days). REvoSim uses a two-dimensional grid of configurable size (default 100x100). Evolutionary agents (digital organisms) do not move during their existence, so occupy a single grid-square. Each grid-square has a configurable number of 'slots' (by default 100); at any one time each slot can either hold one digital organism, or can be empty.

Grid-squares each possess an 'environment', which consists of three integer parameters in the range 0-255. These are visualized as colour, the three parameters representing red, green and blue values. The environment for the entire grid can thus conveniently be visualised as a raster (bitmapped) 24-bit colour image. Environments can be static (specified by a single raster image), or dynamic (specified by a sequence of raster images). For dynamic environments, the number of simulation-iterations between image updates is configurable but by default 50 (i.e. the environment image advances to the next in sequence every 50 iterations). REvoSim can optionally interpolate between environmental images to provide smooth transitions.

REvoSim digital organisms possess a genome of 64 binary digits (bits; 0 or 1). This genome is split into two halves of 32 bits each. The first 32 bits (referred to as the 'coding genome') are used to calculate the fitness of the digital organism in an environment (see below; note also that throughout this contribution 'fitness', when used in relation to REvoSim, represents the calculated fitness within the software, rather than reproductive success, although the two are linked). The second 32 bits (the 'non-coding genome') does not contribute to fitness, but all 64 bits are used to compute genetic difference for the purpose of breeding and species identification. No concept of morphology or ontogeny is simulated. The only properties digital organisms possess are their genome and three integers; a calculated fitness (see below), an 'age counter' (initially 15 by default), and an energy level (initially 0).

REvoSim runs are initiated by seeding with a single individual in the central pixel of the simulation; this can be a user-defined or random genome (but capable of surviving in the given environment at the start of the simulation). The simulation then consists of repeated iterations, each of which consists of a *grid-cell iteration phase*, a *reproduction phase*, and a *dispersal phase*, the details of which are described in the next section. Simulations can be run indefinitely (within the memory limits of a system), or for a set duration. A batch-facility exists to enable unsupervised repeat runs.

Grid-cell iteration phase

Each cell is considered independently (and hence parallel processing is used at this stage). For each digital organism living in each cell:

1. The age counter is decremented; if the age counter reaches 0 the digital organism dies and is removed from the simulation immediately.
2. Energy e_t is added to this specific digital organism according to the formula:

$$e_t = f_t (1 + (\frac{e}{\sum f}))$$

Where e is a global constant, set by default to 2000, $\sum f$ is the sum of the calculated fitness (see below) of all individuals living in the grid square, and f_t is the calculated fitness f for this particular digital organism. e_t uses integer arithmetic (i.e. a floor function) for speed, and the '1 +' term hence exists to ensure that some energy is always added even where integer approximation reduces the $e / \sum f$ term to zero.

3. If the energy possessed by a digital organism is greater than the sum of the 'Breed Threshold' and the 'Breed Cost' (by default both 500), the 'Breed Cost' is deducted from the digital organism's energy, and it is flagged for this iteration as 'attempting to breed'.

Reproduction phase

The default mode of reproduction in REvoSim is recombinant (sexual, although there are no genders); an asexual mode is also available. As for grid-cell iteration, each cell is considered separately and a parallelised algorithm is used. Each digital organism flagged as 'attempting to breed' in point three of the grid-cell iteration phase is paired at random with another that is attempting to breed in the same grid square. There is no bar on self-breeding (a digital organism may randomly select itself to breed with). At this point self-breeding is enforced if the simulation is set to asexual reproduction. Compatibility is then determined for each pair; this can be by genomic (hamming) difference, by tracked species identification (see species section below), or both. Genomic difference compatibility is computed by counting the number of bits differing between the two genomes. If this is greater than the 'Max Diff To Breed' setting (by default 3), the pairing is incompatible. Tracked species identification compatibility simply requires that both members of the pair belong to the same species (see below). If the pairing is compatible, a child genome is generated. This is achieved by taking a bit comb, i.e. a random number where - on average - every other bit is a 1, using an AND operation ($\wedge$) to copy through the genome of one parent for all bits in the comb that are 1. The bit comb is then inverted and the same operation used to copy the inverse bits from the second parent. This newly constructed genome - which inherits half its genetic material from each parent - then has a chance of mutation; a single random integer is generated in the range 0-255, and if this is less than the 'Chance of Mutation' setting (by default 10), then a randomly selected bit in the genome is flipped (from 0 to 1 or 1 to 0). Once generated, the new genome is entered into the 'settling list' for the dispersal phase (see below).

Where breeding is aborted no further attempt is made in the current iteration, but the digital organism that initiated the breed has the 'Breed Cost' refunded to its energy pool; this ensures that it has enough energy to attempt to breed again next iteration (as energy never decreases), should it not have died of old age. 'Breed Cost' refunds are made only to the initiator of breeding not to the partners, to avoid double-refunds. While it is possible for a digital organism to breed many times in an iteration, once as an initiator and many times as a randomly selected partner, a mean of two attempts will be made each iteration that it is flagged as 'attempting to breed'. The reproduction model used by REvoSim is intended to be computationally efficient, to simulate sexual reproduction in sufficient detail that gene-pools can exist, to enforce genetic isolation between gene-pools that lack intermediate forms, and to simulate difficulty but not impossibility of finding a mate for relatively rare forms.

Dispersal phase

REvoSim digital organisms do not move between grid squares: this choice was made for computational efficiency. Allowing movement would require a second sweep through every organism each iteration to implement their movement. Instead, mobility occurs only following breeding, when a progeny digital organisms 'settles' in a selected grid square that may be different from that in which its parents reside. We envisage this process as the settling of a larval propagule. A random angle α is generated in the range

0-360 degrees, and a random integer b in the range 0-255. δx , δy are calculated as follows.

$$\delta x = (\sqrt{\frac{65536}{b+1}} - 16) \cdot \frac{\sin(a)}{d}$$

$$\delta y = (\sqrt{\frac{65536}{b+1}} - 16) \cdot \frac{\cos(a)}{d}$$

Where d is the setting value 'Dispersal' (by default 15). Calculations are approximated using integer arithmetic and use pre-calculated look-up tables for speed, with the final result rounded down to the nearest integer.

δx and δy represent the offset position from the parent's grid-square. With the default $d=15$ this calculation provides rare large jumps (up to 16 squares with $d=15$), but the mean offset is ~ 0.6 squares, and nearly 75% of children will attempt to settle in the same square as their parents.

If this offset places the child off the edge of the grid, it is removed from the simulation immediately (i.e. it dies). To avoid resultant boundary effects, the experimental 'world' can be set to be toroidal (wrapped around — in both directions). If a child is placed off the top in a toroidal world, it appears on the bottom. A 'nonspatial' setting is also available; this replaces the calculation described above with random placement on the grid.

Once the grid-square for settling has been selected, the child is placed into the first available free slot in that square — if there are no free slots, it fails to settle and dies immediately. Otherwise its initial energy level is set to 0, and its fitness to the environment of the square is calculated (see below). If the fitness is 0, the child dies immediately.

By default, fitness is calculated only on settling — it is not recalculated if or when the environment changes. Where environmental change is slow with respect to the lifespan of individual digital organisms this simplification is not expected to cause artefacts. A recalculate fitness option is provided in REvoSim settings, which causes all fitnesses to be recalculated every iteration prior to step 1 of the grid-cell iteration phase — all digital organisms who are found to have a fitness of 0 after this calculation are immediately removed from the simulation. This option increases computational load by approximately 50%, and should only be used where rapid environmental change is simulated.

Calculation of fitness

The calculation of the fitness of a digital organism within an environment is at the core of the REvoSim model, and the approach is described here in depth. At the start of the simulation tables of 'masks' are generated to enable the calculation of fitness for each level of red, green and blue. For each of these three primary colours, 256 32-bit combinations are computed at the start of the simulation, to correspond to each of the possible levels of red, green or blue (0–255).

These are generated by choosing a random 32-bit number (i.e. a random pattern of 32 bits) as a combination to correspond to a red level of 0, randomly flipping one bit of this number (i.e. turning a single 0 to a 1 or vice versa) to generate a combination to correspond to a red level of 1, and continuing to randomly flip one bit to generate each subsequent combination for all red levels 2–255. 'mask' combinations are generated in the same way for the green and blue elements of the environment, starting again with a random number for 0 level in each case — see Figure 2.

To calculate the fitness of a genome in a particular environment, a boolean exclusive OR ($\oplus$) operation is carried out between the 32-bit coding genome (see above) and each of the three 'masks' specified by the red, green and blue levels of the environment (see above). The $\oplus$ operation results in a 0 where both bits of the operands are 0 or both are 1, but a 1 where the bits of the operands are different. The total number of '1's in these three results is summed to give a value x in the range 0-96. Fitness f is an integer

initially calculated as:

$$f = s - |x - t| \quad (1)$$

where t and s are constants configurable in settings. t is termed the 'fitness target', and represents the optimum value of x to maximise f . It is set by default at a value of 66; while it has a theoretical range of 0-92, values very different to the default are not recommended (see below for rationale). s is termed the 'settle tolerance', as it controls the distance from the optimal fitness target at which the fitness will be 0; digital organisms which attempt to settle in an environment within which they have a fitness of 0 die instantly, and are removed from the simulation. s also provides the maximum attainable value of f . By default it is set to 15.

Negative values of f are not permitted and are computed as 0, so more completely

$$f = \max(0, s - |x - t|) \quad (2)$$

See also figure 3.

The REvoSim algorithm for f is computationally inexpensive. $\oplus$ operations require only a single processor instruction. Counting of '1's in a 32-bit number is performed efficiently by pre-computing counts for all possible 16-bit numbers (of which there are only 65,536), and using these precomputed tables to separately count the '1's in each half of the 32-bit number. The remainder of the calculation uses only integer subtraction, modulus computation, and integer comparison, all of which are very fast operations. The algorithm also has three important properties.

1. **Small environment changes result in small fitness changes.** By virtue of the means in which the 'masks' are constructed, fitness decreases gradually as environment changes gradually; for instance a digital organism with a maximal f value (s) for the environment (128,128,128) will also be well adapted to (128,128,127), with an f of $s-1$.
2. **Small genomic changes result in small fitness changes.** A single point-mutation in the genome (flipping a 1 or 0) will result in a change in f of 1, 2 or 3 points.
3. **There are many ways to be maximally fit for a given environment, but more ways to be less fit.** There are 2^{32} possible 32-bit genomes. The probability of a randomly selected genome being maximally fit for this environment is extremely small, but nonetheless there are a large number of different maximally fit genomes possible (see Table 1). The probability of a randomly selected genome having being able to survive in Table 1 (i.e. $f > 0$) is around 0.24; this can be reduced by lowering s .

Property 3 ensures that 'trivial' convergence, where the simulation repeatedly converges to the same coding genome for any given environment regardless of the starting point, is very unlikely to occur. It also ensures that substantial evolutionary change will be required to converge on an optimal solution when digital organisms enter new environments that they are not initially adapted to; they are very unlikely to find themselves well adapted, although there is a fair chance that they will be able to survive. The decreasing counts with increasing fitness further ensure that adaptive peaks are just that — the vast majority of mutations in a relatively fit digital organism will reduce rather than increase its fitness. The values in Table 1 are sensitive to t — higher values of t result in relatively few ideal solutions, and lower values result in a relatively high probability of a randomly selected genome being very fit. For example, $t=71$ gives only 24 genomes where $f=15$; $t=61$ gives over 9 million, around 0.2%. This sensitivity necessitates a narrow practical range for t , and analyses of data of this type forms the rationale behind the selection of the default $t=66$.

Property 1 allows species to have environmental ranges. Properties 1 and 2, taken together, are necessary for stepwise evolution to take place; small changes in environment can drive changes in

genomic structure without rendering digital organisms non-viable, and single genetic changes will typically have small enough effects to allow fitness ‘hill climbing’.

Species tracking

Species, in the sense of isolated gene-pools that are typically adapted to different environments, emerge naturally from REvoSim simulations. The software includes a mechanism to identify and track these species, assigning each species a unique ID number which is visualized with a randomly assigned colour during the simulation. Internally it also records the parent species of each new species found, and can log data such as origination times, extinction times, and modal genome of the species. REvoSim includes a facility to reconstruct and output phylogenetic trees from these logs (see below).

Species identification is relatively computationally expensive; under certain circumstances it can represent the slowest element of the entire simulation. For this reason it is not enabled by default, and is only performed if species visualization is on, or if the user selects the ‘track species’ option. Figures 4 and 5 present an exemplar visualization of the species identification problem using 8 x 8-bit genomes for simplicity. Figure 4 shows these genomes, and their genetic dissimilarities.

Figure 5 shows all ‘connections’ with a genetic dissimilarity of only 1 (left) — if the simulation was set up to allow breeding only between genomes with a difference of 1 or less, two reproductively isolated species would be present in this sample (genomes 1,2,3 and 4 forming one, and 5,6,7 and 8 the other). Note that although 1,2,3 and 4 are part of the same gene-pool, many pairings (e.g. 2 and 3) are too dissimilar to breed directly with each other directly; this does not preclude them from belonging to the same species. If the simulation allows breeding between partners with genetic dissimilarity of 2 or less (the default), then the connection shown on the right of Figure 5 apply. Here all eight genomes form one rather variable species; despite distances between pairings being as high as 6 (between 7 and 4), these are potentially reproductively connected via intermediates. A partial isolation between groupings (1,2,3,4) and (5,6,7,8) is still apparent in this situation, but a treatment of these as a single species is nonetheless appropriate as they are not reproductively isolated — each grouping does contain members who can breed with the other.

A naïve but thorough species-identification algorithm (for a single iteration) would hence examine all genomes present, determine the genetic distance between each pair, and separate out groupings of genomes that lack connectivity to other groupings. This algorithm would, however, be prohibitive computationally for large numbers of genomes (complexity at least of $O(n^2)$). Further complications are the identification of species connectivity *between* iterations, and, where new species arise, the identification of their parent species.

REvoSim solves these problems by tracking species continuously through genealogical descent, and investigating them to determine if any have split (i.e. speciated) at regular intervals — by default every 50 iterations. Each species is assigned a unique ID number; all individuals in the simulation are tagged with a species number, and offspring inherit the species ID of their parent. If the two parents have different IDs, which is possible but not likely, the offspring inherits one at random; alternatively reproduction can be set to enforce conspecific parentage (see above), circumventing this issue. Each investigation iteration (every 50 iterations by default), a full clustering analysis is carried out for each living species, as defined by the species ID tags of individuals. Within each species, a set of genomes is constructed, and an exhaustive grouping algorithm determines how many isolated clusters of genomes exist (using the concepts of Figures 4,5). If a single cluster exists, the species persists (with the same ID number). If multiple clusters exist, the most diverse one (i.e. the cluster with the largest number of different genomes) preserves the ID number from the last analysis, and other clusters are assigned new numbers (and their parentage noted). This algorithm is fast where the total population is split into many species, adding little more than 10–20% of computational overhead to the simulation, but in scenarios with a small number of variable species that contain a large number of genomes, this figure can reach

500% or higher. Potential exists for speedups from heuristic approximations and parallelization; these are not currently implemented.

Numerous species concepts with which to classify living organisms exist, and are a subject of ongoing discussion. The root cause of numerous current issues relate to the need to create clusters within one time slice to understand a lineage-based process occurring through time. These issues exist also in REvoSim. The species concept described above is designed with modelling macroevolution in mind, and has been developed as a parallel for eukaryote species concepts. It is closely aligned with the biological species concept by identifying interbreeding populations that are reproductively isolated. Other species concepts exist, and are debated, both in eukaryotes (Wheeler and Meier 2000), and in prokaryotes, where concepts in recent years have relied on 16S rRNA sequence analysis and DNA hybridization (Rosselló-Mora and Amann 2001; Brenner et al. 2015). In prokaryotes active debates are ongoing regarding the utility of competing means of delineating species (Gevers et al. 2005; Konstantinidis and Tiedje 2005; Koeppe et al. 2008; e.g. Kim et al. 2014). Under a number of bacterial classification schemes the species within REvoSim might better be thought of as ecotypes: a population within a single ecological niche, but which share some of the properties associated with eukaryote species (Cohan 2001).

We have chosen the REvoSim species concept outlined herein on the basis several considerations: it is relatively computationally efficient and, in contrast to many concepts and methods used to identify species in organic systems, is tractable using organisms comprising a 64-bit genome. Furthermore, the REvoSim species concept is comparable to the most widely used and well established eukaryote species definition. As well as being difficult to implement, other definitions within both eukaryotes and prokaryotes continue to be developed (e.g. Bobay and Ochman 2017). They are changing quickly and none have reached widespread acceptance. An additional reason behind a eukaryote-like species concept is that many experimental evolution studies using microorganisms are species agnostic. In studies including, for example, long term evolutionary dynamics (Lenski et al. 2015; Lenski and Burnham 2018), mutation rates (Krašovec et al. 2014, 2017), and parallelism (Vogwill et al. 2016) the patterns observed and reported are not impacted by species definitions used. Furthermore, in some studies looking at, for example, radiations (Kassen et al. 2004; Meyer et al. 2011), disturbances (Buckling et al. 2000), and diversity (Kassen et al. 2000) morphs/genotypic lineages in asexual organisms are considered valuable analogues to species. In contrast, for studies looking at long term trends in macroevolution (Benson et al. 2010, 2014; Dunne et al. 2018) species are the fundamental unit, and required to probe these questions.

Model implementation

Software packages

We provide here a software implementation of this evolutionary model (REvoSim, summarised in Figure 6) alongside a package for generating environments for the software (EnviroGen). The software is intended to be easy to use, and in order to make it available to the widest range of researchers possible, it is fully documented, and is published under the released under the GNU General Public License. It is parallelised, and coded in C++. The source code for each package is available in a GitHub repository, to which new releases will be pushed as the software continues to be developed (<https://github.com/palaeoware>). Documentation hosted by Read the Docs is linked from the repository and programs. The software has a graphical user interface designed using the Qt cross-platform software development kit, and can thus be compiled for Windows, Mac OS, and Linux/Unix. For ease we provide executables for each of operating system in the Palaeoware GitHub repository. REvoSim is the more demanding of the two software packages, but will run on any modern desktop/laptop: we recommend a minimum of 1GB RAM and a 1.8 GHz or faster, ideally multicore processor. We also recommend a minimum screen resolution of 1280x720 if using the software without the genome comparison docker (and 1920x1080 if this is enabled).

Random Numbers

Both the above packages use a pre-generated table of 65,536 random numbers 0–255 for randomization during the simulation, both for speed and to avoid potential biases from pseudo-random number generators. 10Mb of quantum-generated random numbers from <http://www.randomnumbers.info/> are packaged into the executable; these can be replaced with any other random number file preferred by the user.

Visualization and Output

Population visualization modes

REvoSim simulations can be tracked in many ways. Several population visualizations are provided; these are calculated at regular intervals (by default 50 iterations), and displayed onscreen as a raster image alongside an image representing the current environment (Figure 1, visualisation mode is coding genome as colour).

1. Population count

Each grid-square is visualized with a grey-level representing the current number of digital organisms alive in the square.

2. Mean fitness

Each grid-square is visualized with a grey-level representing the mean fitness of all digital organisms alive in the square, scaled so that white is maximum fitness (=‘Settle Tolerance’).

3. Coding genome as colour

For each grid-square, the most commonly occurring (modal) genome is computed. The 32-bit coding part of this genome is then converted to a colour, where the level of red is the (scaled) count of 1’s in the least significant 11 bits, the level of green is the (scaled) count of ‘1’s in the next least significant 11 bits, and the level of blue is the (scaled) count of ‘1’s in the most significant 10 bits. This visualization approach provides a quick means of distinguishing genomes, ensuring that small genomic changes result in small changes in colour. It does not, however, guarantee that the same colour will in all cases represent the same genome (as many very different genomes may possess the same bit-counts in each of the three 11/10 bit segments).

4. Non-coding genome as colour

As for 3 above, but the non-coding half of the genome is visualized.

5. Gene frequencies

For each grid-square, the mean value of the first three coding bits is computed, and a composite colour created using the mean of the least significant bit as red, the mean of the next least significant bit as green, and the mean of the third least significant bit as blue.

6. Settles

Each grid-square is visualized with a grey-level representing the number of successful ‘settling’ events in that square since the last visualization.

7. Breed/Settle Fails (R=Breed; G=Settle)

Each grid-square is assigned a colour representing the number of ‘fails’ since the last visualization. The number of ‘breed fails’ (attempts to breed that were aborted due to compatibility) provides the red level, and the number of ‘settle fails’ (attempts to settle that resulted in a fitness of 0 and hence the death of the settling digital organism) provides the green level. Fail visualizations are scaled non-linearly (using $v = 100f^{0.8}$, where f = mean fails per iteration, and v is visualized intensity on a scale 0-255). Fail visualization maps the edge of species or subspecies ranges, as it highlights cells where gene-flow is restricted by any mechanism.

8. Species

REvoSim simulations often produce discrete species, i.e. reproductively isolated gene-pools, and for many macroevolutionary studies the identification and tracking of the fate of species is a requirement. This visualisation option assigns a unique colour to each species, and colours each grid-square with the species in the lowest occupied slot of that square.

Output options

REvoSim can output files to record runs and facilitate subsequent analysis, which are placed within the folder specified in the output tab, in a subfolder titled *REvoSim_output*. Any of the above visualisations can be saved as an .png image stack as required, providing a visual record of a run (see, e.g. SI movie; Garwood et al. 2019). In addition to this, REvoSim has two forms of text logging, written at different points during a run. Both these files have a standardised, CSV-based format to allow parsing using, for example, the R programming language, or any other analytical framework desired.

The file *REvoSim_log.txt* is written at every polling interval (i.e. at the same time as the visualisation is updated and species are calculated). This starts with a timestamp and the settings of the run, and then provides an ongoing record of the run, listing for each poll: iteration number; population grid data for the entire simulation — the number of living digital organisms, their fitness, the number of entries on the breed list, a count of failed breed attempts, and the number of species; and then finally species data in the form of the species ID for every living species above a user-defined number of individuals — the origin of that species, its parent, the number of individuals alive in the simulation, and and the genome of a randomly sampled individual from this species. Because it is written as a simulation progresses, this log is not able to record information such as the duration of species' existence, or their peak population size/geographic range, for example. To overcome this, there is a second log, saved as *REvoSim_end_run_log.txt* written at the end of a run or when the user requests this using an option in the software. This features the full tree from the run in Newick format with branch lengths, and then detailed species data including dynamics and geographic range data for all species above a certain size.

Results: Analysis, Validation, and Limitations

REvoSim is a highly abstracted model that incorporates numerous elements of evolution from organic systems, for example: reproduction, mutation, and dispersion. The abstracted nature of REvoSim makes quantitative validation of the model challenging: direct observation of evolution on these temporal scales and population numbers would obviate the need for a model. Here we document and analyse a simple REvoSim run, provide qualitative comparisons of REvoSim results with those from experimental microevolutionary studies, and discuss emergent macroevolutionary phenomena that find parallels in real evolution, providing confidence in the validity and applicability of our model and implementation.

Figure 7 shows the relationship for the first 2500 iterations of a REvoSim run between the mean fitness of the organisms across the whole environment, the number of breeding individuals, and then the probability that a living organism is attempting to breed. The run was seeded by 100 copies of a genome capable of surviving in the central pixel (the organisms with these genomes differ in their start age: on initialisation this is a random integer up to the start age defined in the settings, and which applies to the rest of the run for all progeny). This initial genome possessed a fitness of 2. As the run commenced, no breeding occurred in the first few iterations as organisms in the central square built up sufficient resource to allow them to breed. When this threshold was reached, some organisms began to breed as others died (of old age) and the population decreased, causing a rapid peak in breeding probability. At this point the mean fitness dropped slightly, reflecting the occurrence of mutation in offspring. This is typically deleterious, but in this initial expansion the organisms settling in empty or near-empty cells

experienced little competition. The fine perturbations in the early phase of this run seen near the origin of the graph relate to the organisms breaking into the left and right stripes of the default environment. As the simulation reached carrying capacity (around 1000 iterations) the number of breeding individuals plateaued, and breed probability stabilised at around 0.1. This was slightly lower while the organisms were still expanding their range. This may have resulted from lower population numbers on the edge of the range reducing the likelihood of finding a mating partner. Mean fitness was around 11 at the point where carrying capacity was reached, and subsequently gradually increased towards the maximum possible value (15). The existence of some deleteriously mutated offspring of optimally adapted parents prevented the mean from reaching this optimal value.

REvoSim runs can be compared with studies of experimental evolution conducted using model organisms. Precise equivalence is not a realistic goal; not only is determining 'correct' parameter settings highly problematic, but real systems are inherently more complex than those of REvoSim. Even in well studied organisms such as *Escherichia coli*, new discoveries surrounding evolutionary mechanisms are occurring regularly (for example density-associated mutation-rate plasticity (Krašovec et al. 2014, 2017)). These are both developing our understanding of microevolutionary processes, and increasing the known complexity of even these model systems. Furthermore, important relationships such as that between mutation rate, peak fitness, and fitness landscape are poorly constrained and difficult to assess using living systems (Clune et al. 2008). Despite these limitations, clear parallels exist between REvoSim results and those from 'wet-lab' experimental work. One example is the bacterial (MEGA)-plate experiment of Baym et al. (Baym et al. 2016). This comprised a large acrylic plate holding agar with exponentially higher concentrations of antibiotic in strips from the edges towards the centre. The SI movie shows a run that uses REvoSim's dual reseeding functionality, within an environment of stripes grading from blue to pink. The spatiotemporal patterns of the two systems are very similar, with mutation-driven adaptation allowing new populations to establish themselves in different stripes, and thus multiple coexisting lineages to evolve (SI movie). These compete with each other within the stripe as the simulation/experiment continues.

Another useful comparison is with one of the largest scale microbial experiments, the long-term evolution experiment (LTEE) (Good et al. 2017). This has been running with 12 populations of *E. coli* for 25 years, or >66,000 generations. It has been used to study and quantify adaptation through natural selection in these populations, and the periodic sampling then freezing of each population allows a range of analyses. The experiment has revealed stably coexisting and also competing ecotypes, as also seen in REvoSim runs (SI movie), and facilitates the extent of adaptation to be measured across time (by competing frozen ancestral strains with later ones: relative fitness is the ratio of realized growth rates between the strains)(Lenski 2017). This latter process has allowed relative fitness curves to be derived for each strain over the duration of the study, and a mean fitness for all populations to be produced. The same can be achieved by calculating relative fitness every 60 iterations over several hundred REvoSim runs (both these curves are shown in Figure 8; but see discussion in the section Limitations). In both instances the relative fitness curve starts with an initial value of 1. In the LTEE this increases immediately and rapidly, whilst in REvoSim runs there is a lag before breeding starts whilst organisms build up energy to allow breeding. However, in both cases the pattern shown is an initial, rapid increase in relative fitness, which then levels off as the experiment/simulation continues. There are significant differences between these systems (e.g. REvoSim has a higher chance of mutation, far simpler organisms, and fundamentally different fitness landscape). Nonetheless, the similarity of the relative fitness patterns through time suggests that REvoSim is simulating a comparable process to the wet-lab experiment.

REvoSim simulations exhibits emergent properties, at a level above microevolution, that are known to occur in the real world. Speciation — the formation of reproductively isolated clusters of individuals — occurs in REvoSim under certain conditions. This occurs both across environmental boundaries (i.e. allopatric speciation), and within a single environment (sympatric); both are visible at various points in

the SI movie (see also below), which has a species panel where reproductively isolated groups are given contrasting colours. Adaptive radiations occur across sharp environmental boundaries. These are shown in the SI movie, where individuals capable of living in adjacent environmental stripes evolve, and then radiate to fill the otherwise empty environmental region. Extinctions also occur within REvoSim: these are clear in the SI movie where a new region of grey environment is introduced across a dark red circle. An initial colonisation from the adjacent red region results in two new species which then compete with the advancing colonisation from the grey regions. REvoSim runs also have a carrying capacity (SI movie) which is dictated by settings such as the amount of resource available to each grid square and the number of grid squares. In natural environments the existence of a carrying capacity for population size has achieved broad consensus (McGlade 1999); its emergence in REvoSim is hence compatible with real-world systems. Carrying capacity at a species level may also occur in both the real world (del Monte-Luna et al. 2004) and in REvoSim runs, but as real-world patterns are imperfectly understood (see Close et al. 2017 and references therein), we consider this an area in which simulation may inform understanding of nature rather than as an opportunity for validation.

Limitations

By necessity models simplify the system they are used to investigate, and limitations result. The focus in REvoSim on computational efficiency results in a level of abstraction which dictates that some of these are potentially significant. A number of simplifications are apparent, and whilst they may be addressed with future software development, they could meanwhile impact on hypotheses developed using REvoSim in its current form. We do not provide an exhaustive list (evolving organic systems are orders of magnitude more complex than a 64-bit digital organism based computer simulation), but we highlight a number of these herein and some potential implications. The lack of direct inter-organism interactions rules out parasitism, predation, symbiosis and many other important ecological considerations, and thus currently this system is better suited to studying environmentally-driven factors than it is biotically-driven (i.e. red queen) patterns. The constant genome size and point mutations in breeding, preclude a number of important modes by which molecular evolution proceeds in organic systems, including insertions, deletions, duplications, frameshift mutations, translocations, inversions, horizontal gene transfer and repeat expansions. This will make major genomic shifts in REvoSim less likely than the real world. We note that some alternative digital evolution packages, e.g. Avida, do allow genome size and factors such as this to be studied (e.g. Gupta et al. 2016), and thus may be preferable for such topics.

Another important consideration, especially when comparing REvoSim to experimental evolution studies (and a caveat to our discussion above), is that the LTEE suggests fitness trajectories in microorganisms can be fitted with a power law (Lenski et al. 2015; Lenski and Burnham 2018). An implication of this is that fitness will permanently increase for a given population and environment, albeit at decreasing rates (we know of no studies to tackle whether this pattern is likely to occur in complex eukaryotes). In contrast, for a static environment within REvoSim the fitness has a hard-coded maximum, which the results above demonstrate is reached within 3000 iterations. Achieving a computationally efficient mechanism of calculating fitness with no cap is challenging, however, the impact of this issue can be ameliorated by running REvoSim experiments with a dynamic environment. This allows simulations to be run for millions of iterations without the population reaching maximal fitness: rather evolution allows the digital organisms to follow fitness peaks. Whilst this does introduce an additional variable, it also better reflects the situations in which organisms evolve in the real world where the environment is rarely static.

Another implication of the fitness algorithm in REvoSim is that it relies on integer operations for speed. This enforces a discrete number for fitness constrained to be an integer for individual digital organisms. It is thus possible to have a number of organisms within a single environmental colour or a given pixel with different genomes but precisely the same fitness. Multiple different optimal species ($f=15$ under

default settings) might easily and co-exist. Differential migration pressure effects may slowly make rarer species lose ground, but they are not capable of out-competing each other on fitness grounds alone as their fitnesses are precisely equal. In these cases their evolution will be driven by random elements (selection for breeding and then the impact of mutation). The impact of this can be minimised by selecting environments without solid blocks of colour, and by using dynamic environments; here the actual fitness of a species to its environment is the mean f from all grid squares occupied, which will not be integral.

In real-world organisms, genes can interact and epistasis occurs, altering the fitness landscape (e.g. Bershtein et al. 2006). REvoSim lacks mechanisms by which this occurs directly, rather mutations impact only on the fitness algorithm, and ability to compete with other digital organisms. This reduces the realism of the fitness landscape employed herein, but the exact nature of epistasis in both micro- and macro-organisms (and thus its impact on the fitness landscape) is a topic of current active research (Sanjuan and Elena 2006). Given this uncertainty, the difficulty of achieving epistasis using a computationally efficient algorithm, and the fact that evolutionary simulation software already in existence is well-suited to studying this topic (Lenski et al. 1999; Wilke et al. 2001), it is unlikely to be a focus for development in the near future. Rather, for studies in which epistasis is likely to have a significant impact, *Avida* is better suited.

In a related fashion, the lack of organismal development and associated pathways in REvoSim will limit the dynamics of evolution within the system: in many eukaryotes, complex life cycles increase the niches available to organisms and the nature of competition between these (Wilbur 1980). This could impact on studies of diversity given that developmental mechanisms facilitate niche differentiation (Moran 1994) and indeed speciation (e.g. Kamiya 1992). Similar impacts may result from a lack of phenotypic plasticity (Pfennig et al. 2010).

In conclusion, there are numerous simplification and abstractions within REvoSim for computational efficiency. Whilst this necessitates a critical analysis of results collected using the model, we do not believe that it reduces the potential of the software. One of the benefits of a simplified model is that it provides an opportunity to probe whether these missing complexities are required for real-world patterns to appear. If real-world patterns appear, they can be investigated and quantified, and if they do not this too can be informative.

An example simulation

In addition to the above results, we present herein an example run of REvoSim to demonstrate its potential for studying evolutionary dynamics with large populations. This is presented in the SI movie, but we also provide stills as SI figures. This run proceeds over 20,500 iterations, and has a dynamic environment in which the simulation is seeded within a red circle surrounded by grey environment. During the run, a grey rectangle is used to split this red circle. The simulation uses default settings, other than: dispersal, which is set to 70 to accommodate the dynamic environment; energy input, which is set to 750 to restrict the population size to ~100,000 individuals; mutation which is 1 to prevent significant genomic changes between polling iterations. When the run is initiated, the initial individual from seeding begins to breed, and disperses to fill the red circle. Almost immediately, however, there is a sympatric speciation event within this circle (SI movie, 36 seconds; SI figure 2). The two species coexist, and together fill the central zone of red environment; they are initially unable to settle in the grey region outside this zone, however. A number of small species appear, but generally do not survive for a significant duration. At member of the dark green species manages settles in the grey zone (SI movie, 46 seconds; SI figure 3), and maintain a breeding population across this environmental divide, whilst expanding into the vacant ecospace. In contrast a second breakout event occurs shortly after with individuals from the lighter blue species establishing themselves in grey (still vacant) space, and then

immediately undergoing allopatric speciation (SI movie, 48 seconds; SI figure 4). Another such event occurs immediately above this, and the three resulting species coexist, ultimately filling the grey region with no organisms (SI movie, 58 seconds; SI figure 5). At this point, a grey rectangle is placed across the red circle, driving many of the organisms living here to extinction (SI movie, 59 seconds; SI figure 6), and making several smaller species extinct, as reflected by the species diversity curve of this run (figure 9; environmental perturbation occurs at 10,000 iterations).

Several species break into this newly vacated space: one via allopatric speciation, and a population of the first species from the simulation which moves in from the red region adjacent via dispersal. There is also subsequent recolonisation by two species from the surrounding grey region, coupled with a number of speciation events on the advancing wave of colonisation of these two species. One such event occurs in a species that survives from the beginning simulation. This includes individuals from populations on the advancing front of the species colonising from the grey environment but also includes those individuals left in remaining red regions (SI movie, 60 seconds; SI figure 7). The smaller population is considered the new species here, as described in the REvoSim species documentation (see above), even though the individuals within the red environment represent the first established population. These species then compete at the centre of the vacated space in the simulation with one of the new species ultimately going extinct. In the fragmented red areas of the original environment, high levels of speciation begin to occur as the run continues (SI movie, 75 seconds onwards; SI figure 8). This example simulation has been chosen for several reasons: it demonstrates several important aspects of macroevolution, as outlined previously (speciation, extinction, adaptive radiations), but also the versatility and scope of investigations in REvoSim. By including both space and time, even relatively simple runs such as this show a number of complex and interesting patterns which warrant investigation and quantification.

Discussion

REvoSim runs reproduce phenomena observed at both macro- and microevolutionary scales. All simulation models are simplified/idealised models of nature. In view of the reproduction of these phenomena we consider REvoSim to be similar enough to 'true' evolutionary processes to possess predictive value, i.e. for its results in scenarios beyond those discussed above to be (potentially) applicable to the real world. We recommend cautious interpretation of REvoSim results, however: studies should be cognisant of the model's abstracted nature and the impact this may have on the conclusions drawn from its use. REvoSim lacks, for example, direct ecological interactions, ontogeny, hierarchical organisation, evolvability, epigenetics, and a niche construction. Continued development of the software package will allow many of these to be achieved within REvoSim, albeit at a computational cost. These developments are planned to include larger genomes, ecological interactions, pathogens, evolvable mutation rates, and variable breeding (i.e. the genome dictates whether recombination occurs).

The un-augmented version of REvoSim presented herein nonetheless provides scope for many investigations. For example, the distinction between macro- and microevolution results in part from the significant difference in temporal scales between palaeontological studies and those focussing on ecological patterns or employing wet lab experimentation (Gontier 2015 provides discussion of what is a multifaceted distinction). There is a paucity of systems and methods that bridge this divide: lab-based systems are currently subject to severe limitations (Bell 2016). Debate thus continues as to whether macroevolution is fundamentally different to microevolution, or whether it represents the action of microevolutionary processes writ large (Erwin 2001; Simons 2002; Jablonski 2017; this is also touched upon by the discussion of Hannisdal and Liow 2018 within this special issue). REvoSim may shed light on this by bridging this temporal divide: adaptive radiations, extinctions, and speciation can be studied in this context, and the impact of space on these investigated. In comparison to microbial experimental evolution, this may be considered closer to the evolution of ecotypes, but see discussion in the species

section. REvoSim could be used to assess the impact of biotic as opposed to environmental forcing, and thus the impact of the Red Queen within the system (Strotz et al. 2018). We do, however, note the limitations above regarding the lack of organismal interaction and introducing this will be a focus of future development. How environmental, biological, and temporal factors impact on speciation and/or stasis is easily quantified within the system, as are the patterns of recovery post mass extinctions, which can be induced through sudden environmental perturbations. Subsampling from log files would also allow the creation of a fossil record, and any desired artificial biases to this could be introduced. This could be used in studying the impact of an imperfect fossil record on analysis of radiations, extinctions, and speciation; the reconstruction of phylogenies and the myriad roles these can play (see Slater and Harmon 2011 and references therein); taxonomic diversity studies (Close et al. 2017, 2018); and patterns associated with palaeobiogeography (Mannion et al. 2014).

Acknowledgements

We are grateful to Christopher Knight, Roger Benson, Guillaume Gomez, Joe Keating, and Robert Sansom for discussion, and to Andrew Smith for inviting this contribution. We would like to acknowledge the assistance of the Software Sustainability Institute. We are grateful for the comments of two anonymous reviewers that greatly improved this work. The work carried out by the Software Sustainability Institute is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) through grant EP/H043160/1 and EPSRC, BBSRC and ESRC Grant EP/N006410/1.

Data Archiving Statement

Data for this study are available in the Dryad Digital Repository:
<https://datadryad.org/review?doi=doi:10.5061/dryad.7tg18mb>

References

- ADAMI, C. 1998. *Introduction to Artificial Life*. Springer-Verlag, New York.
- , OFRIA, C. and COLLIER, T. C. 2000. Evolution of biological complexity. *Proceedings of the National Academy of Sciences*, **97**, 4463–4468.
- AGUIAR, M. A. M. de, BARANGER, M., BAPTESTINI, E. M., KAUFMAN, L. and BAR-YAM, Y. 2009. Global patterns of speciation and diversity. *Nature*, **460**, 384–387.
- ARTHUR, R., NICHOLSON, A., SIBANI, P. and CHRISTENSEN, M. 2017. The Tangled Nature Model for organizational ecology. *Computational and Mathematical Organization Theory*, **23**, 1–31.
- BATUT, B., PARSONS, D. P., FISCHER, S., BESLON, G. and KNIBBE, C. 2013. In silico experimental evolution: a tool to test evolutionary scenarios. *BMC Bioinformatics*, **14**, S11.
- BAYM, M., LIEBERMAN, T. D., KELSIC, E. D., CHAIT, R., GROSS, R., YELIN, I. and KISHONY, R. 2016. Spatiotemporal microbial evolution on antibiotic landscapes. *Science*, **353**, 1147–1151.
- BELL, G. 2016. Experimental macroevolution. *Proceedings of the Royal Society B: Biological Sciences*, **283**, 20152547.

- BENSON, R. B. J., BUTLER, R. J., LINDGREN, J. and SMITH, A. S. 2010. Mesozoic marine tetrapod diversity: mass extinctions and temporal heterogeneity in geological megabiases affecting vertebrates. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, **277**, 829–34.
- BENSON, R. B. J., FRIGOT, R. A., GOSWAMI, A., ANDRES, B. and BUTLER, R. J. 2014. Competition and constraint drove Cope's rule in the evolution of giant flying reptiles. *Nature Communications*, **5**.
- BERSHTEIN, S., SEGAL, M., BEKERMAN, R., TOKURIKI, N. and TAWFIK, D. S. 2006. Robustness–epistasis link shapes the fitness landscape of a randomly drifting protein. *Nature*, **444**, 929–932.
- BOBAY, L.-M. and OCHMAN, H. 2017. Biological Species Are Universal across Life's Domains. *Genome Biology and Evolution*, **9**, 491–501.
- BRENNER, D. J., STALEY, J. T. and KRIEG, N. R. 2015. Classification of Procaryotic Organisms and the Concept of Bacterial Speciation. In *Bergey's Manual of Systematics of Archaea and Bacteria*, American Cancer Society, 1–9 pp.
- BUCKLING, A., KASSEN, R., BELL, G. and RAINEY, P. B. 2000. Disturbance and diversity in experimental microcosms. *Nature*, **408**, 961–964.
- CLOSE, R. A., BENSON, R. B. J., UPCHURCH, P. and BUTLER, R. J. 2017. Controlling for the species-area effect supports constrained long-term Mesozoic terrestrial vertebrate diversification. *Nature Communications*, **8**, 15381.
- , EVERS, S. W., ALROY, J. and BUTLER, R. J. 2018. How should we estimate diversity in the fossil record? Testing richness estimators using sampling-standardised discovery curves. *Methods in Ecology and Evolution*, **9**, 1386–1400.
- CLUNE, J., GOLDSBY, H. J., OFRIA, C. and PENNOCK, R. T. 2010. Selective pressures for accurate altruism targeting: evidence from digital evolution for difficult-to-test aspects of inclusive fitness theory. *Proceedings of the Royal Society of London B: Biological Sciences*, **278**, 666–674.
- , MISEVIC, D., OFRIA, C., LENSKI, R. E., ELENA, S. F. and SANJUÁN, R. 2008. Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. *PLoS computational biology*, **4**, e1000187.
- COHAN, F. M. 2001. Bacterial Species and Speciation. *Systematic Biology*, **50**, 513–524.
- COSTA, C. L. N., LEMOS-COSTA, P., MARQUITTI, F. M. D., FERNANDES, L. D., RAMOS, M. F., SCHNEIDER, D. M., MARTINS, A. B. and AGUIAR, M. A. M. In press. Signatures of microevolutionary processes in phylogenetic patterns. *Systematic Biology*.
- DUNNE, E. M., CLOSE, R. A., BUTTON, D. J., BROCKLEHURST, N., CASHMORE, D. D., LLOYD, G. T. and BUTLER, R. J. 2018. Diversity change during the rise of tetrapods and the impact of the 'Carboniferous rainforest collapse'. *Proc. R. Soc. B*, **285**, 20172730.
- ERWIN, D. H. 2001. Macroevolution is more than repeated rounds of microevolution. *Evolution & Development*, **2**, 78–84.
- GARWOOD, R. J., SPENCER, A. R. T. and SUTTON, M. D. 2019. Data from REvoSim: Organism-level simulation of macro- and microevolution. Dryad Digital Repository. <https://datadryad.org/review?doi=doi:10.5061>
- GAVRILETS, S., LI, H. and VOSE, M. D. 2000. Patterns of parapatric speciation. *Evolution*, **54**, 1126–1134.

- GENIEYS, S., BESSONOV, N. and VOLPERT, V. 2009. Mathematical model of evolutionary branching. *Mathematical and Computer Modelling*, **49**, 2109–2115.
- GEVERS, D., COHAN, F. M., LAWRENCE, J. G., SPRATT, B. G., COENYE, T., FEIL, E. J., STACKEBRANDT, E., DE PEER, Y. V., VANDAMME, P., THOMPSON, F. L. and SWINGS, J. 2005. Re-evaluating prokaryotic species. *Nature Reviews Microbiology*, **3**, 733–739.
- GONTIER, N. 2015. Uniting Micro- with Macroevolution into an Extended Synthesis: Reintegrating Life's Natural History into Evolution Studies. In *Macroevolution*, Springer, Cham, 227–275 pp.
- GOOD, B. H., MCDONALD, M. J., BARRICK, J. E., LENSKI, R. E. and DESAI, M. M. 2017. The dynamics of molecular evolution over 60,000 generations. *Nature*, **551**, 45–50.
- GUPTA, A., LABAR, T., MIYAGI, M. and ADAMI, C. 2016. Evolution of Genome Size in Asexual Digital Organisms. *Scientific Reports*, **6**, 25786.
- HANNISDAL, B. and LIOW, L. H. 2018. Causality from palaeontological time series. *Palaeontology*, **61**, 495–509.
- HANSEN, C. J., KAWALER, S. D. and TRIMBLE, V. 2004. *Stellar Interiors: Physical Principles, Structure, and Evolution*. *Astronomy and Astrophysics Library*. Springer-Verlag, New York.
- JABLONSKI, D. 2017. Approaches to Macroevolution: 1. General Concepts and Origin of Variation. *Evolutionary Biology*, **44**, 427–450.
- KAMIYA, T. 1992. Heterochronic Dimorphism of *Loxoconcha Uranouchiensis* (Ostracoda) and its Implication for Speciation. *Paleobiology*, **18**, 221–236.
- KASSEN, R., LLEWELLYN, M. and RAINEY, P. B. 2004. Ecological constraints on diversification in a model adaptive radiation. *Nature*, **431**, 984–988.
- , BUCKLING, A., BELL, G. and RAINEY, P. B. 2000. Diversity peaks at intermediate productivity in a laboratory microcosm. *Nature*, **406**, 508–512.
- KIM, M., OH, H.-S., PARK, S.-C. and CHUN, J. 2014. Towards a taxonomic coherence between average nucleotide identity and 16S rRNA gene sequence similarity for species demarcation of prokaryotes. *International Journal of Systematic and Evolutionary Microbiology*, **64**, 346–351.
- KOEPEL, A., PERRY, E. B., SIKORSKI, J., KRIZANC, D., WARNER, A., WARD, D. M., ROONEY, A. P., BRAMBILLA, E., CONNOR, N., RATCLIFF, R. M., NEVO, E. and COHAN, F. M. 2008. Identifying the fundamental units of bacterial diversity: A paradigm shift to incorporate ecology into bacterial systematics. *Proceedings of the National Academy of Sciences*, **105**, 2504–2509.
- KONSTANTINIDIS, K. T. and TIEDJE, J. M. 2005. Genomic insights that advance the species definition for prokaryotes. *Proceedings of the National Academy of Sciences of the United States of America*, **102**, 2567–2572.
- KRAŠOVEC, R., BELAVKIN, R. V., ASTON, J. A. D., CHANNON, A., ASTON, E., RASH, B. M., KADIRVEL, M., FORBES, S. and KNIGHT, C. G. 2014. Mutation rate plasticity in rifampicin resistance depends on *Escherichia coli* cell–cell interactions. *Nature Communications*, **5**, ncomms4742.
- , RICHARDS, H., GIFFORD, D. R., HATCHER, C., FAULKNER, K. J., BELAVKIN, R. V., CHANNON, A., ASTON, E., MCBAIN, A. J. and KNIGHT, C. G. 2017. Spontaneous mutation rate is a plastic trait associated with population density across domains of life. *PLOS Biology*, **15**, e2002731.

- LABAR, T. and ADAMI, C. 2016. Different Evolutionary Paths to Complexity for Small and Large Populations of Digital Organisms. *PLOS Computational Biology*, **12**, e1005066.
- LAIRD, S., LAWSON, D. and JENSEN, H. J. 2008. The Tangled Nature Model of Evolutionary Ecology: An Overview. In *Mathematical Modeling of Biological Systems, Volume II*, Birkhäuser Boston, 49–62 pp.
- LENSKI, R. E. 2017. Experimental evolution and the dynamics of adaptation and genome evolution in microbial populations. *The ISME Journal*, **11**, 2181–2194.
- and BURNHAM, T. C. 2018. Experimental evolution of bacteria across 60,000 generations, and what it might mean for economics and human decision-making. *Journal of Bioeconomics*, **20**, 107–124.
- , OFRIA, C., COLLIER, T. C. and ADAMI, C. 1999. Genome complexity, robustness and genetic interactions in digital organisms. *Nature*, **400**, 661–664.
- , WISER, M. J., RIBECK, N., BLOUNT, Z. D., NAHUM, J. R., MORRIS, J. J., ZAMAN, L., TURNER, C. B., WADE, B. D., MADDAMSETTI, R., BURMEISTER, A. R., BAIRD, E. J., BUNDY, J., GRANT, N. A., CARD, K. J., ROWLES, M., WEATHERSPOON, K., PAPOULIS, S. E., SULLIVAN, R., CLARK, C., MULKA, J. S. and HAJELA, N. 2015. Sustained fitness gains and variability in fitness trajectories in the long-term evolution experiment with *Escherichia coli*. *Proc. R. Soc. B*, **282**, 20152292.
- MANNION, P. D., UPCHURCH, P., BENSON, R. B. J. and GOSWAMI, A. 2014. The latitudinal biodiversity gradient through deep time. *Trends in Ecology & Evolution*, **29**, 42–50.
- MCGLADE, J. M. 1999. *Advanced ecological theory: principles and applications*. Blackwell Science, Malden, Mass.
- MCGUFFIE, K. and HENDERSON-SELLERS, A. 2014. *The Climate Modelling Primer*. Wiley-Blackwell, Chichester, West Sussex, UK ; Hoboken, NJ, UK.
- MEYER, J. R., SCHOUSTRA, S. E., LACHAPELLE, J. and KASSEN, R. 2011. Overshooting dynamics in a model adaptive radiation. *Proceedings of the Royal Society of London B: Biological Sciences*, **278**, 392–398.
- DEL MONTE-LUNA, P., BROOK, B. W., ZETINA-REJON, M. J. and CRUZ-ESCALONA, V. H. 2004. The carrying capacity of ecosystems. *Global Ecology and Biogeography*, **13**, 485–495.
- MORAN, N. A. 1994. Adaptation and Constraint in the Complex Life Cycles of Animals. *Annual Review of Ecology and Systematics*, **25**, 573–600.
- PENNOCK, R. T. 2007. Models, simulations, instantiations, and evidence: the case of digital evolution. *Journal of Experimental & Theoretical Artificial Intelligence*, **19**, 29–42.
- PFENNIG, D. W., WUND, M. A., SNELL-ROOD, E. C., CRUICKSHANK, T., SCHLICHTING, C. D. and MOCZEK, A. P. 2010. Phenotypic plasticity's impacts on diversification and speciation. *Trends in Ecology & Evolution*, **25**, 459–467.
- ROSINDELL, J., HARMON, L. J. and ETIENNE, R. S. 2015. Unifying ecology and macroevolution with individual-based theory. *Ecology Letters*, **18**, 472–482.
- ROSSELLÓ-MORA, R. and AMANN, R. 2001. The species concept for prokaryotes. *FEMS Microbiology Reviews*, **25**, 39–67.
- SANJUAN, R. and ELENA, S. F. 2006. Epistasis correlates to genomic complexity. *Proceedings of the National Academy of Sciences*, **103**, 14402–14405.

- SIGWART, J. D., SUTTON, M. D. and BENNETT, K. D. 2018. How big is a genus? Towards a nomothetic systematics. *Zoological Journal of the Linnean Society*, **183**, 237–252.
- SIMONS, A. M. 2002. The continuity of microevolution and macroevolution. *Journal of Evolutionary Biology*, **15**, 688–701.
- SLATER, G. J. and HARMON, L. J. 2011. Unifying fossils and phylogenies for comparative analyses of diversification and trait evolution. *Methods in Ecology and Evolution*, **4**, 699–702.
- STAUFFER, D., KUNWAR, A. and CHOWDHURY, D. 2005. Evolutionary ecology in silico: evolving food webs, migrating population and speciation. *Physica A: Statistical Mechanics and its Applications*, **352**, 202–215.
- STROTZ, L. C., SIMÕES, M., GIRARD, M. G., BREITKREUZ, L., KIMMIG, J. and LIEBERMAN, B. S. 2018. Getting somewhere with the Red Queen: chasing a biologically modern definition of the hypothesis. *Biology Letters*, **14**, 20170734.
- TATARENKOV, A. and AVISE, J. C. 2007. Rapid concerted evolution in animal mitochondrial DNA. *Proceedings of the Royal Society of London B: Biological Sciences*, **274**, 1795–1798.
- TEMPLETON, A. R. 2006. *Population Genetics and Microevolutionary Theory*. John Wiley & Sons, Inc., Hoboken, NJ, USA.
- TURCOTTE, D. and SCHUBERT, G. 2014. *Geodynamics*. Cambridge University Press, Cambridge, United Kingdom.
- VOGWILL, T., PHILLIPS, R. L., GIFFORD, D. R. and MACLEAN, R. C. 2016. Divergent evolution peaks under intermediate population bottlenecks during bacterial experimental evolution. *Proc. R. Soc. B*, **283**, 20160749.
- WHEELER, Q. and MEIER, R. 2000. *Species Concepts and Phylogenetic Theory: A Debate*. Columbia University Press.
- WILBUR, H. M. 1980. Complex Life Cycles. *Annual Review of Ecology and Systematics*, **11**, 67–93.
- WILKE, C. O., WANG, J. L., OFRIA, C., LENSKE, R. E. and ADAMI, C. 2001. Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature*, **412**, 331–333.
- YAMAGUCHI, R. and IWASA, Y. 2013. First passage time to allopatric speciation. *Interface Focus*, **3**, 20130026.

Figure captions and table

FIGURE 1. The software package REvoSim, described for the first time herein. The panel on the left shows a colour derived from the coding genome of the modal digital organism in each pixel. That on the right shows the environment in which they are living.

FIGURE 2. The process for generating ‘masks’ for the first four levels of red, green and blue. Grey shading indicates a bit flipped from the value above.

FIGURE 3. Calculation of f using an environment of RGB levels (0,3,2), the mask tables of figure 2, the arbitrarily chosen coding genome 10111011100101101101101100011010, $t=66$ and $s=15$.

FIGURE 4. Eight 8-bit genomes for species comparison example (left). Disimilarities (right) between the genomes

shown on the left, calculated as the number of different bits for each pair.

FIGURE 5. Connections between the genomes of Figure 4 that have a dissimilarity of one (left), compared with (right) those that have a dissimilarity of one (solid lines) or two (dashed lines).

FIGURE 6. A summary flowchart showing the REvoSim algorithm.

FIGURE 7. A graph showing the mean fitness, probability of breeding (i.e. the number of entries on the breed list divided by the total population), and number of breeding individuals (entries on the breed list) for the whole REvoSim grid over the first 2500 iterations of a run. This uses the default, static environment and default settings. We note that not all entries on the breedlist will successfully breed in a run with sexual breeding. For clarity, breed probability and breeding individuals are smoothed using an equally weighted moving mean with a window of 3.

FIGURE 8. Fitness trajectories in evolving *E. coli* communities (top) and REvoSim (bottom). Left redrawn from (Lenski 2017) and shows the trajectory for the mean relative fitness across all populations in the LTEE over 50 000 generations. Error bars are 95% confidence limits based on replicate populations. Right is based on the mean relative fitness from 478 REvoSim runs (a batch of 500, those in which the population went extinct were discarded). Each run was reseeded with an individual of fitness nine to approximate the process by which the LTEE curve was derived (REvoSim is not limited to replicates using the same individual). Error bars mean absolute deviation, chosen to demonstrate the spread of the data. SI Figure 1 (Garwood et al. 2019) is a copy of this figure with 95% confidence limits based on 478 repeats for direct comparability with the left graph (we note that the width of a confidence interval decreases with increasing repeats, such as those possible in REvoSim, and are thus potentially less appropriate for the latter than the LTEE).

FIGURE 9. A graph showing species counts over the course of the exemplar REvoSim run presented herein (polled every 150 iterations). An animated version demonstrating this whilst the simulation is occurring is provided in the SI movie. Seconds on x-axis represent time from start of movie.

<i>F</i>	Count	Proportion of genomes
0	3,261,775,269	75.944%
1	251,349,612	5.852%
2	210,471,648	4.900%
3	169,066,312	3.936%
4	130,105,296	3.029%
5	95,746,779	2.229%
6	67,229,976	1.565%
7	44,923,140	1.046%
8	28,473,840	0.663%
9	17,046,291	0.397%
10	9,585,504	0.223%
11	5,029,236	0.117%
12	2,443,032	0.057%
13	1,094,049	0.025%
14	476,496	0.011%
15	150,816	0.004%

TABLE 1. Fitness values *f* computed for all possible genomes for a single environment (128,128,128) using *t*=66, *s*=15, and a random set of 'masks' generated using the process of Figure 3.