

# The Next Step in Understanding Population Dynamics: Comprehensive Numerical Simulation

John C. Sanford<sup>1</sup> and Chase W. Nelson<sup>2</sup>

<sup>1</sup>Department of Horticulture, NYSAES, Cornell University, Geneva, NY,

<sup>2</sup>Rainbow Technologies, Inc., Waterloo, NY,  
USA

## 1. Introduction

Natural populations are always changing. Hardy-Weinberg assumptions are almost never realized because populations are seldom in equilibrium, and many random events (e.g., mutations, population size fluctuations, and environmental perturbations) irrevocably alter the genetic makeup of populations. Such genetic change can be either for the better (some populations adapt and expand) or for the worse (some populations shrink and become extinct). Change that occurs as the result of natural selection is termed *adaptive evolution* because natural selection favors the survival of organisms that are best adapted to their environments (i.e., have high fitness). On the other hand, *nonadaptive evolution* refers to change that occurs as the result of factors that act independently of organismal fitness (e.g., random genetic drift or mutation pressure). Because change within a population depends on so many variables and involves innumerable chance events, the study of population dynamics is both challenging and fascinating.

Geneticists study population-level change in three very different ways. The first approach is the empirical one. This involves the actual observation of living populations over time, and involves direct experimentation with the variables that affect population dynamics. One author (JCS) engaged in this approach the first 20 years of his career, first as a plant geneticist and breeder, and later as someone who was actively involved in plant genetic engineering. The empirical approach has the limitation that it is only possible to study changes that happen in a short amount of time, and which display effects that are highly visible or easily tracked. For example, when selecting for increased crop yield in plant breeding, it is only possible to reliably determine very substantial yield differences (roughly 10% or more) between genotypes, even with carefully replicated field trials. Such observations are very useful, but obviously miss most of what is happening at the genetic level (i.e., the innumerable subtle genetic changes that are happening within the population). While this empirical method has serious limitations, it has played a major role in the development of modern agriculture and modern medicine.

The second method is the historical/comparative approach. This involves observing differences (especially amino acid and DNA sequence differences) within existing

populations and trying to infer population histories (e.g., see Li & Durbin, 2011). Such inferences can include the degree of relatedness between populations, the time of their divergence, and what parts of their genomes were affected by either positive selection, negative selection, or no selection. This historical approach is inherently limited by its various underlying assumptions, such as the concept of constant-rate molecular clocks or the neutrality of synonymous mutations. Because many exceptions to each assumption exist (e.g., see Sauna & Kimchi-Sarfaty et al., 2011 on synonymous mutations), the reliability and relevance of the historical approach has been hotly debated (for example, see Wilson & Cann, 1992 and Thorne & Wolpoff, 1992 on the molecular clock), but there is no doubt that it can sometimes help us to correctly infer certain genetic events in the past. There are a variety of computational tools that are designed to facilitate this historical approach in population genetics. These tools fall under the umbrella of bioinformatics, and include software packages that align sequences, infer phylogenetic trees, and perform various statistical analyses with the given data.

The third method is the theoretical approach. This involves studying how hypothetical or idealized populations might behave in forward-time, starting from a specific state, based upon our knowledge of genetics and population biology. Ten years ago one author (JCS) shifted his research focus from the empirical approach to the theoretical approach, because the theoretical approach allows consideration of the bigger picture and bigger questions.

The field of theoretical genetics was established primarily by mathematicians (e.g., Fisher, Haldane, and Wright). These scientists realized that even though mutations arise and segregate randomly, and survival is influenced by many random elements, and mating is largely random, still, directional processes such as selection can play an important role in shaping the genetic makeup of a population over long periods of time.

The mathematical approach to population dynamics has been very fruitful in terms of understanding numerous specific aspects of population change, when each is considered in isolation. This is particularly true, for example, where selection for a single trait is mathematically modeled, or where numerous neutral mutations are drifting in a population. The main limitation of the mathematical modeling approach is that it invariably requires extreme simplification of the model (e.g. just considering one or a few loci or mutations, or just one or a few variables). Unfortunately, real biological populations are not at all simple, and so there arises the possibility that the results of simplified mathematical models may not correspond to biological reality. This is especially a concern where theoretical models have become highly abstract, such that common sense can no longer help us gauge whether or not theoretical predictions are reasonable.

For these reasons mathematical models need to be tested. One way to do this is by returning to the empirical approach: studying living biological populations through many generations to validate theory. However, this is usually not practical, especially for organisms with long generation times. As a practical alternative, mathematical models can be tested in virtual populations using numerical simulation. It was for this reason that simple numerical computer simulations were first developed by population geneticists - to test and validate specific mathematical models.

Numerical simulations can be seen as the empirical enactment of real processes, but in a virtual environment. Even though numerical simulation experiments happen in a computer

environment, numerical simulators can be used to conduct real experiments, and can illuminate processes happening in the real world. In terms of modeling fitness change over time, a good numerical simulation can act very much like an accountant's spreadsheet. Spreadsheets can be made to accurately and honestly reflect the true financial status of a corporate entity. Every dollar is tracked from beginning to end, as it comes and as it goes. In fact, in a large corporate entity, a spreadsheet is the only reliable way to see the big financial picture. Corporations and governments may not be able to trust their accountants, but they can at least trust the operation of their spreadsheets. Likewise, when numerical simulators are carefully designed to reflect the real world, they can be powerful and trustworthy tools. When used properly, numerical simulations can inform us about what is likely to be happening in the real world, even when direct observation is not feasible.

Population geneticists first used simple numerical simulation to validate a particular component of genetic systems. However, as computational power has grown, and as the science of numerical simulation has become more sophisticated, we have now reached the point where we can analyze population dynamics in a comprehensive, integrated and empirical manner within a virtual environment, independent of copious and often very abstract mathematical modeling. Such simulation should enable us to obtain a more biologically integrated picture of how real populations change.

Seven years ago one author (JCS) had the opportunity to oversee the development of a comprehensive numerical simulator for the aforementioned purposes. Since that time, a group of biologists and computer scientists have been collaborating to develop a numerical simulator that can simultaneously model all the major known factors that affect genetic change, as well as their relevant interactions, to better approximate what occurs in the real world. The resulting program, Mendel's Accountant (Mendel), appears to be the first program that has seriously endeavored to do this. Mendel has been described in previous publications (Sanford et al., 2007a, 2007b), and is now beginning to be used for both research and teaching. This tool should not be viewed as a replacement for previous tools already developed within this field, but it is clear that it represents a major step forward.

## 2. Mendel's Accountant

Mendel's Accountant simulates genetic change within a population as it moves forward through time. Mendel does this by establishing a virtual population of individuals, and then precisely simulates mutation, selection, and gene transmission through many generations, always in the most biologically realistic manner possible. Mendel is unique in that it attempts to treat all aspects of population dynamics simultaneously and comprehensively, thereby ushering in for the first time the prospect of simulating reasonable approximations of biological reality.

Mendel's Accountant is an apt name for this program because it is largely a "genetic accounting" program. Every generation, huge numbers of specific mutations are introduced into a population, spread over the genomes of many individuals. Through the ensuing generations, some of these mutations are lost, while others increase in frequency. Each mutation must be tracked through many individuals and through many generations, along with all data that apply to that mutation (each has an allelic ID, mutational fitness effect, degree of dominance, and chromosomal location). During a large run, Mendel can track

hundreds of millions of different mutations. Not only does Mendel do the genetic accounting associated with tracking individual mutations, it simultaneously does the genetic accounting associated with tracking: 1) linkage blocks as they recombine; 2) net fitnesses of each individual; 3) the distribution of the fitness effects of all the accumulating mutations; and 4) the resulting distribution of allele frequencies.

Genetic accounting via numerical simulation is possible because the underlying processes (Mendelian inheritance, random mutations, differential reproduction) are all relatively simple and mechanistic in nature and are therefore subject to straightforward accounting procedures. Furthermore, all the relevant biological variables are easily specified as parameters for use in simulation (e.g., population size, mutation rate, distribution of mutational fitness effects, heritability, and amount of selective elimination each generation). Like many high-performance numerical simulations, the core of the Mendel program is written in Fortran 90, allowing the execution of tasks that are extremely demanding computationally, making it possible to process huge amounts of genetic data.

To explain how Mendel works in the simplest way possible, it is useful to consider the series of decisions that an experimenter must make. Firstly, the experimenter must define the species and its reproductive structure. Is it haploid or diploid? How big is the genome, and what fraction is functional? Is its reproduction sexual or clonal? Does the species ever self-fertilize? All these biological factors can be modeled by Mendel, and must be specified by the user, because they have a substantial impact on population dynamics. These parameters determine how reproduction and gene transmission will occur within the virtual species.

Secondly, the experimenter must define the characteristics of a particular population within the species. How big is the population before and after selection? Are there sub-populations? How many generations do we wish to observe? These parameters define the actual scope and architecture of a particular experiment.

Thirdly, the user must specify reproductive details. The reproductive rate must always be high enough to create a population surplus each generation, such that this surplus can then be selectively removed each generation. For example, the default reproduction rate is 3. In this case the number of offspring generated each generation is always 3 times larger than the specified population size. This creates a surplus population large enough for selection to remove two of every three offspring in the next generation. If the population under study reproduces sexually, recombination will occur at this stage. The experimenter must specify the number of chromosomes (assuming two cross-overs per chromosome) and the number of linkage blocks (this affects segregation of linkage blocks during gamete production).

Fourthly, the experimenter needs to specify the mutations that will be added to the population. After creating a new virtual population of offspring, Mendel then begins to add new mutations to those individual offspring. Mendel assigns mutations to individuals randomly, following a Poisson distribution. The experimenter specifies a mutation rate appropriate for the species under study (or one that is of theoretical interest). Likewise, the experimenter must specify a distribution of mutational fitness effects. Typically this distribution will include deleterious, neutral, and beneficial mutations. The mutations that are added to the population are drawn randomly from a user-specified pool of potential mutations (usually having a Weibull distribution of fitness effects). Drawing from such a distribution, some mutations will have large effects, but most will have small (nearly-

neutral) effects (Kimura, 1983), as occurs in nature (Eyre-Walker & Keightley, 2007). Each new mutation has an identifier for tracking purposes, a fitness effect, a specified degree of dominance, and a chromosomal location (i.e., a designated linkage block).

Lastly, the experimenter needs to specify the nature of the selection process. Once Mendel has created a newly mutated population of offspring, it must implement selective removal. To do this Mendel first calculates the combined effect of all mutations in each individual (initial individuals containing zero mutations having a fitness of one, with beneficial mutations increasing fitness and deleterious mutations reducing fitness). Mutations can be combined either additively or multiplicatively (or in alternative ways, i.e., epistically). Once the fitness of each individual has been calculated, a certain fraction of the population is selectively eliminated based upon genetic fitness, usually eliminating the exact population surplus, so that the original population size is restored. Selective removal can be either by truncation selection, probability selection, or partial truncation. To add biological realism, the user can specify a heritability of less than one, such that fitness variations caused by environmental noise will be added to the genetic fitness to establish the fitness phenotype, which is then the basis for selection. The individuals that survive selection will then be ready to repeat the cycle of mutation, reproduction, and selection.

During a single experiment, Mendel can routinely simulate hundreds of millions of newly arising mutations. Each mutation is tracked through all generations, until it is either lost or goes to fixation, or until the experiment is complete. Throughout the whole run Mendel is continuously monitoring, recording, and plotting the average number of mutations per individual, individual and average fitness, population size history, the fitness distributions of accumulating mutations, selection threshold histories, linkage block net fitness values, and mutant allele frequencies.

### 3. Forward-time population genetic numerical simulations

There are numerous forward-time simulation tools currently in use within the field of population genetics. Detailed reviews on the subject are available elsewhere (e.g., see Carvajal-Rodríguez, 2008; Kim & Wiehe, 2008; Liu et al., 2008; Carvajal-Rodríguez, 2010). It is useful to provide a general overview of these programs to properly appreciate the types of problems that can be addressed with such simulations. Every forward-time simulation is designed with a particular application in mind, and each is best suited to study a certain class of scenarios.

#### 3.1 FPG

The FPG (forward population genetic) simulation is the most similar to Mendel in concept (Hey, 2009). The user is able to define a mutation rate per generation for deleterious, neutral, and beneficial mutations, a fitness model (i.e., whether mutations combine additively, multiplicatively, or epistically), a population size (i.e., number of genomes), and various other parameters. It is possible to track average fitness over time and perform analyses for linkage disequilibrium, fitness, and heterozygosity at the conclusion of an experiment. However, FPG is not readily accessible to most biologists. Running the program requires the user to understand and construct a string of input values at the command line level. Some of these values are not intuitive, e.g., a populational selection coefficient. FPG is also limited in

terms of genome and population sizes. Its distributed version allows only 1000 sequences, and each sequence is restricted to 32 polymorphic sites, limiting the total number of effective mutations to 32000. This may be sufficient to model some long-term dynamics of populations with very small genomes, but it is generally inadequate for eukaryotic organisms. Finally, when large numbers of mutations occur, FPG appears to ignore fixed mutations after the fitness exceeds what can be stored as a floating point. Thus this program appears to ignore fixed mutations and their fitness effects to save computational resources, sometimes leading to counterintuitive output. Because of these considerations, FPG may work well for simple illustrative case studies but simply cannot handle the population sizes and number of mutations necessary to realistically address most biological scenarios. Perhaps its greatest limitation is that it models all mutations within a class (e.g., deleterious) as having identical fitness effects.

### 3.2 SimuPOP

SimuPOP is another forward-time simulation well-suited for tackling problems that involve a small number of functional loci. This program is especially helpful when studying the evolutionary dynamics of disease predisposing alleles. SimuPOP also allows the user to define auxiliary information for individual organisms. This information can be used to group organisms into virtual subpopulations, potentiating assortative mating based on characteristics such as genotype, sex, or age. Though flexible, simuPOP is challenging to use. Perhaps one of its largest drawbacks is that the user must write a Python macro to run an experiment. This can be an arduous task for even modestly complex evolutionary scenarios, as it requires a deep understanding of simuPOP and how to utilize its various components.

### 3.3 FREGENE

FREGENE is most innovative not so much for its novel implementation and flexibility, but rather for its use of a rescaling technique to make large problems less computationally intensive. Specifically, both population size ( $N$ ) and number of generations are decreased by a factor of  $\lambda > 1$ , while all rate parameters (e.g., mutation, recombination, and migration rates) are increased by the same factor. This can be relaxed at the end of an experiment, such that (for example) the population size expands linearly from  $N/\lambda$  to  $N$ . Though this is clearly more computationally expedient than modeling full populations for full lengths of time, it is not ideal. For example, many processes depend on the absolute (not scaled) parameters, such as the fixation probability of beneficial mutations (Kim & Wiehe, 2008). Moreover, more advanced simulation software can usually handle the true population sizes and rates of interest, so there is often no need for such rescaling. FREGENE also implements an uncommon distribution of mutational fitness effects, drawing selection coefficients from two normal distributions (one each for deleterious and beneficial mutations) with user-specified means and variances. This is also sub-optimal, as it has long been agreed that the distribution of mutational fitness effects is approximately exponential, such that the majority of effective mutations are relatively low-impact (Eyre-Walker & Keightley, 2007).

### 3.4 Forwsim

Forwsim is another tool that implements a novel technique to save computational resources. To do this, the user can ask the simulation to look  $k$  generations ahead in order to determine

which chromosomes will be passed to future generations. Once it is determined which chromosomes cannot contribute to future generations, those chromosomes are no longer simulated. Though there is a computational trade-off between looking  $k$  generations ahead and precluding the copying of unnecessary chromosomes, this process does serve to make many evolutionary scenarios more manageable. However, as previously stated, such techniques are often no longer necessary.

### 3.5 Avida

Finally, it is useful to contrast forward-time numerical simulations with digital life programs, especially the Avida simulation (Lenski et al., 2003; Ofria & Wilke, 2004), an elaboration of Tierra (Ray, 1991). There are important differences between the digital life approach and the numerical simulation approach described here. Forward-time numerical simulations attempt to simulate biological processes primarily by tracking numerical values (e.g., fitness) that change based on user-specified conditions. Using values measured in biological research (e.g., mutational fitness effects), the goal of numerical simulation is to accurately predict various populational dynamics under those conditions. A distribution of mutational fitness effects is specified, mutations are assigned certain locations on chromosomes, and their fitness effects simply increment or decrement an organism's fitness according to the selection model and gene interaction. Digital life, on the other hand, attempts to instantiate a model genome itself in the form of self-replicating computer code. Fitness then becomes an emergent value according to that program's behavior in the software environment. A full distribution of mutational fitness effects cannot be specified, only observed. In the case of Tierra, whatever changes allow the program to replicate faster in the simulated environment will be beneficial. Programs tend to shrink, allowing quicker self-replication, and parasitic behavior also emerges.

Avida builds on the concept of Tierra by introducing an external fitness function. In practice, this means that the Avida environment continually examines the population of digital organisms for certain computational operations. When these operations arise, the lucky organism can be rewarded with the ability to execute additional genomic instructions, allowing it to execute its code, and thus replicate, faster. The size of a reward is decided by the user, and experimental results depend critically on these values (Nelson & Sanford, 2011). It should be kept in mind that these operations are arbitrary, i.e., they only increase fitness because the programmer has imposed an arbitrary rule with fitness rewards. In other words, while genome shrinkage is a genuine way to increase replication speed, the fitness rewards based on certain computational operations can occur only because the programmer has altered the software environment to implement such a scheme.

Experiments with digital life systems are usually conducted with the goal of shedding light on general principles that are relevant to all self-replicating systems. However, though digital life research has produced a large number of publications in the biological literature, it appears to lack the ability to address the real issues in the genomics era. For example, genomes in Avida – only 50 to 100 monomers – are many orders of magnitude smaller than real biological genomes, and each Avida “mutation” which introduces a complete computational operation is assigned an unreasonably large fitness effect (i.e., 1.0 – 31.0; see Nelson & Sanford, 2011). Because mutations have an essential role in terms of introducing novel genetic variation, it is critical to simulate mutations realistically, and to examine

realistically large genomes in which the fate of multiple low-impact mutations can be studied. The majority of fixations over the course of evolution do not involve highly beneficial mutations, but rather primarily involve nearly-neutral mutations (e.g., see Kimura, 1983; Hughes, 2008), so the biological relevance of digital life appears extremely limited.

It should now be clear that population genetics is more than an academic exercise. Many real-world problems that are informed by population genetics need resolution and demand biological models that honestly reflect nature. For example, there is very strong evidence that the human population is experiencing a marked decline in fitness due to the accumulation of very slightly deleterious mutations in both the mitochondrial and nuclear genomes (e.g., see Muller, 1950; Kondrashov, 1995; Loewe, 2006; Lynch, 2010). Thus there is a genuine need for serious and honest numerical simulations that can enable us to study these types of real-world problems. Specifically, it is imperative that we have the ability to model and calculate the net fitness effects of large numbers of low-impact mutations in large genomes.

Tool	Organisms & Populations	Mutation	Fitness & Selection
Avida (C++) Lenski et al., 2003; Ofria & Wilke, 2004; Nelson & Sanford, 2011	Organisms (typically asexual and haploid) contain genomes comprised of machine code instructions, as well as stacks and registers for storing and manipulating numbers. Populations exist on a two dimensional grid.	Mutations randomly substitute, insert, or delete one of 26 machine code instructions in the genome. The user specifies the various mutation rates.	Organisms replace each other when they produce a daughter cell. Speed of replication determines fitness, and this can be increased by either genome shrinkage or rewards for computational tasks.
Forwsim (C++) Padhukasahasram et al., 2008	Diploid organisms contain chromosome arrays that store the locations of mutations. The user specifies a probability of selfing. Populations of constant size reproduce with non-overlapping generations.	Mutations occur at a Poisson-distributed rate and insert new integers into chromosome arrays, which may undergo recombination. The number of sites is finite, but mutations occur only at non-polymorphic sites. Locations which are no longer polymorphic are removed.	If using natural selection, the evolution of selected and neutral sites is carried out separately, with selected sites considered first. The program excels at simulating genetic drift in a standard Wright-Fisher process.
FPG (C) Hey, 2009	Diploid organisms contain a user-defined number of chromosomes with a user-specified number of segments. Each segment can hold 32 mutations. It is possible to specify subpopulations. Populations of constant size reproduce with non-overlapping generations.	The user specifies one mutational fitness effect; deleterious and beneficial mutations have equal but opposite effects. Dominance and epistasis are user-specified. No segregating site in the population can receive another mutation. A maximum of 32,000 mutations can be tracked.	The user specifies a fitness model (additive, multiplicative, or epistatic). If fitness falls below a certain point, fixed mutations are ignored thereafter.

Tool	Organisms & Populations	Mutation	Fitness & Selection
FREGENE (C++) Hoggart et al., 2007; Chadeau-Hyam et al., 2008	Sexual diploid organisms (specified selfing probability). Genomes are single linear sequences represented as lists of sites at which a minor allele is present. Genomes recombine. Allows subpopulations. Populations of constant size reproduce with non-overlapping generations.	Deleterious and beneficial mutational fitness effects are modeled as two user-specified normal distributions. Two allele, finite sites model. Dominance, heritability, and recombination are also user-specified. Mutations occur independently at a constant rate. After a derived allele reaches fixation, it no longer affects fitness.	The user defines loci to be under selection, and can restrict such selection to subpopulations or time periods. Individual fitness is equal to unity plus the sum of contributions of selected sites.
Mendel's Accountant (Fortran 90) Sanford et al., 2007	Organisms (typically sexual diploid; selfing is user-specified) contain a user-defined number of chromosomes with a user-specified number of linkage blocks. Allows subpopulations. Populations of constant size reproduce with non-overlapping generations.	Mutations occur at a Poisson-distributed rate with Weibull-distributed fitness effects. Proportions of deleterious and beneficial mutations are user-specified. Allows heritability, dominance, and recombination. Infinite sites, infinite allele model, though back mutations may be specified.	The user selects a fitness model (additive, multiplicative, or epistatic). The user also sets the reproduction rate, with a higher reproduction rate allowing more intense selection (population size must be maintained). All loci may incur mutations affecting fitness.
SimuPOP (Python)	User defines ploidy, number of chromosomes, and loci under selection. The user chooses a mating scheme (including random mating and selfing). Allows subpopulations. Individual information can be specified, e.g., age, allowing non-random mating.	User-defined loci are used to store mutations, represented in a genome by sequential non-negative numbers. Implements a finite sites k-allele model, but allows other models to be specified.	The user specifies a fitness model (multiplicative, additive, or heterogenous). Fitness is assigned to individuals before mating, and parents to reproduce are chosen with probabilities proportional to their fitness.

Table 1. Comparison of several available population genetic simulations.

Readers are encouraged to examine these and other programs to determine which is best suited for particular applications (Carvajal-Rodríguez, 2008; Kim & Wiehe, 2008; Liu et al., 2008; Carvajal-Rodríguez, 2010). Among these forward-time simulations, Mendel appears to be unique in that it is the first comprehensive (and hence most biologically realistic) population genetics numerical simulator. Mendel can simultaneously consider nearly all of the major factors that are recognized to be operational in a real population, and yields a multi-dimensional view of how populations really change. Moreover, the user has an intuitive and user-friendly interface, such that the user need only specify desired values for all available parameters. Mendel was designed and implemented in Fortran 90 to optimize use of computer resources, which allows the user to use an ordinary laptop computer to

track many millions of mutations – orders of magnitude more than would be possible with any other application currently available. Very significantly, Mendel for the first time gives us the ability to model the net fitness effects of large numbers of low-impact mutations in large genomes.

Although Mendel is undergoing continuous enhancement, it has already demonstrated its ability to address a wide range of biological questions. Mendel is open-source code, and researchers are welcome to use it as a spring-board for further improvement. Mendel would seem to provide the most logical platform for building even more advanced simulators, which may eventually enable us to test essentially any biological scenario.

## 4. Applications

We briefly summarize below some basic findings already observed in various comprehensive numerical simulation applications. Some of these findings were exactly as would have been expected, while other findings seemed very surprising (although, upon reflection, they are clearly logical and correct).

### 4.1 Deleterious mutation accumulation

Mendel keeps a tally of how many deleterious mutations have accumulated in each individual. Mendel very consistently shows us that the mean deleterious mutation count per individual increases at an approximately constant rate over time (Sanford et al., 2007b; Gibson et al., 2012). This appears to be a very fundamental phenomenon (Figure 1). In fact, we can only simulate a substantially non-linear accumulation of deleterious mutation count per individual by using highly artificial parameters (Figure 2; see Brewer et al., 2012). Specifically, to cause mutation count per individual to plateau requires all deleterious mutations to have approximately equal affects on fitness, full or partial truncation selection, and sexual recombination. This combination of conditions is highly improbable under most natural circumstances. For example, organisms with very small genomes such as viruses should have a relatively narrow range of mutational fitness effects, but such organisms generally lack any type of regular sexual recombination. The general problem of ever-increasing genetic load within natural populations represents a widely recognized evolutionary paradox (Kondrashov, 1995; Crow, 1997; Sanford et al., 2007b; Gibson et al., 2012) and requires more research. Biologically realistic numerical simulations are the only practical means to further elucidate this problem, because the problem involves high numbers of very low-impact mutations, biological noise, and selection interference.

### 4.2 Beneficial mutation accumulation

Mendel also keeps a tally of how many beneficial mutations have accumulated in each individual. Like deleterious mutations, the number of beneficial mutations per individual tends to increase at a relatively constant rate, except for a very small class of beneficial mutations that have relatively large effects on fitness. Above a certain fitness effect, beneficial mutations are strongly amplified, leading to a period of accelerated mutation accumulation for that set of mutations and any mutations linked to them. The rapid amplification of high-impact beneficial mutations is as would be as expected, but it is striking to see that the large majority of beneficial mutations are too subtle to respond to

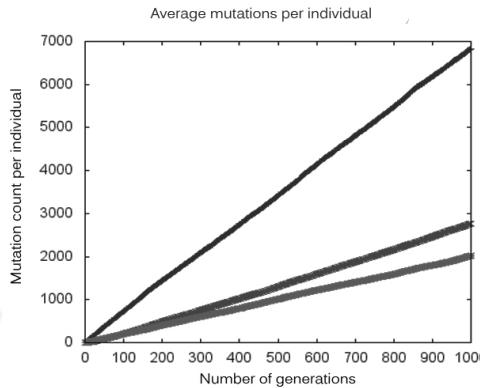


Fig. 1. Mutation accumulation in Mendel. Comprehensive numerical simulation reveals that mutation accumulation over time is largely linear. Mean mutation count per individual was tracked during the course of a Mendel experiment using default settings, except that selection efficiency was optimized (fitness heritability = 1, truncation selection). Input mutations were 10% beneficial, 20% neutral, and 70% deleterious. In this experiment, mean fitness increased by 210%. As can be seen, all three classes of mutation accumulated essentially linearly, but differed relative to their rate of accumulation (slope). Neutral mutations accumulated as expected, just as if there were no selection (bottom line). The deleterious mutations accumulated slightly slower than would be expected if there were no selection (upper line). The beneficial mutations accumulated almost 3 times faster than would be expected if there were no selection (middle line).

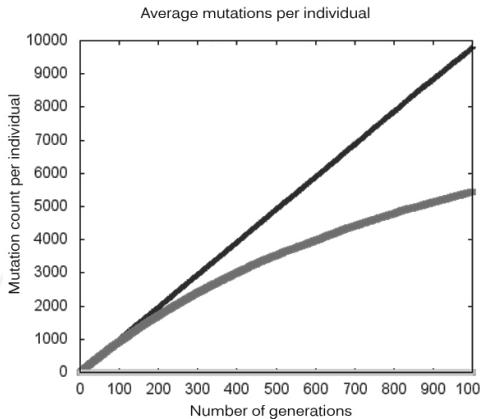


Fig. 2. The effects of uniform fitness effects on deleterious mutation accumulation. Only highly unrealistic conditions cause deleterious mutation accumulation to significantly diverge from linearity (such that deleterious mutation count per individual begins to plateau). The straight line reflects an experiment employing the basic Mendel default settings but with all mutations being deleterious (a broad distribution of mutation fitness effects, fitness heritability = 0.2, probability selection). The curved line reflects the same run but where all mutations had an equal fitness effect (-.0001) and truncation selection was employed.

selection (Sanford et al., 2012). Except for those few high-impact beneficial mutations which are strongly amplified, the ratio of beneficial versus deleterious mutations does not change dramatically in response to selection (see Figures 1 and 3). Since it is well known that deleterious mutations arise much more frequently than do beneficial mutations, this means that many more functional nucleotide sites are being disrupted than are being established, even with intense selection. This suggests there should be a strong natural tendency toward net loss of genetic information over time, even while a limited number of beneficial mutations are being strongly amplified. This represents a second major evolutionary paradox that demands serious attention by researchers. Again, it seems clear that this problem can best be understood by further numerical simulation experiments.

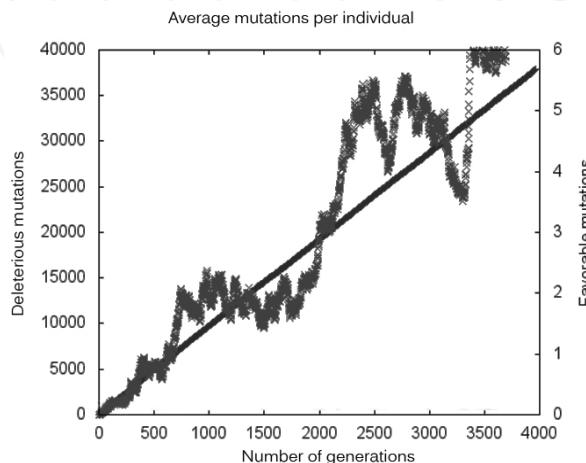


Fig. 3. A comparison of beneficial and deleterious mutation accumulation. Beneficial mutations accumulate essentially linearly, but their dynamics are quite erratic due to their rare occurrence. It is generally understood that beneficial mutations are rare, which makes their study problematic. The Mendel default setting for the relative rate of beneficial mutation is one in 10,000 mutations. Given the Mendel default setting (zero neutral mutations, one in 10,000 mutations beneficial), beneficial mutation accumulation (jagged line, scale on right) is dwarfed by deleterious mutation accumulation (straight line, scale on left). For this reason it is usually necessary to employ rates of beneficial mutation which are exaggerated by several orders of magnitude in order to study the behavior of this class of mutation in detail.

#### 4.3 Change in mean fitness over time

Mendel continuously computes the mean fitness of the population based upon the mutation content of each individual. Mendel reveals that populations tend to decline in fitness (due to the continuous accumulation of deleterious mutations), except when there is a sufficiently high rate of beneficial mutations with sufficiently high fitness effects (Figure 4). There is a critical point where beneficial mutations are both frequent enough and have a strong enough fitness impact to allow stabilization of population fitness. Above this critical point, mean fitness can then increase very rapidly. This raises a variety of interesting research

problems. For example, what conditions are required to reach this critical point needed for fitness stabilization? It appears that when mean fitness is increasing due to just a few high-impact mutations at a few chromosomal locations, a much larger number of functional nucleotides are being disrupted due to relatively low-impact deleterious mutations. The latter are genetically linked to the former in the vast majority of cases. Does this mean that the functional genome size is continuously shrinking? How might we simulate selection for traits which require many beneficial mutations, but none of which are beneficial or selectable apart from the others? How might the multitude of functional, but low-impact, nucleotides in a genome arise? All these questions can best be addressed using comprehensive numerical simulation.

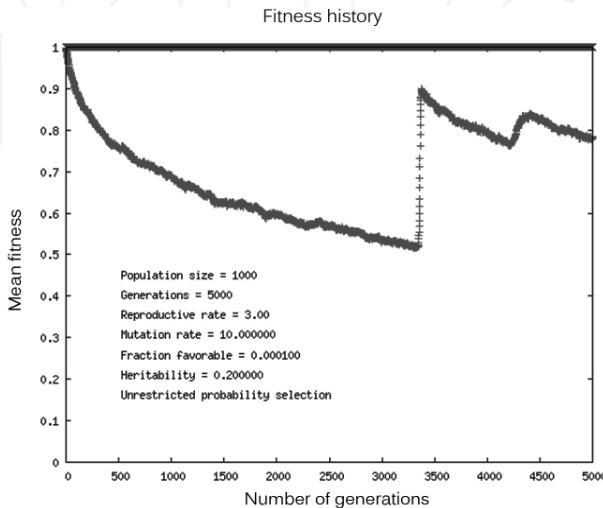


Fig. 4. Fitness trajectory and high-impact beneficial mutations. Change in a population's mean fitness is determined by the net effect of a large number of low-impact deleterious mutations and a small number of relatively high-impact beneficial mutations. In this experiment Mendel's default settings were used (beneficial mutation rate = 0.0001), except that beneficial fitness effects were allowed to range up to 1.0 (one such mutation would double fitness). In 5000 generations the mean deleterious mutation count per individual was 47,648, while the mean beneficial mutation count per individual was 9.8. As can be seen, just two high-impact beneficial mutations largely compensated for over 40,000 deleterious mutations.

#### 4.4 Selection threshold

Mendel's Accountant enables the empirical determination of the "selection threshold" of a given population, which is quantified using a newly proposed statistic, ST. The selection threshold concept is key to understanding the big picture regarding the actual capabilities and limitations of natural selection within a given population. A population's selection threshold is an emergent property of a population and its exact circumstances. It is the consequence of the net effect of all those variables that enhance or interfere with selection efficacy. One of the primary variables that limits selection efficacy is the phenomenon of selection interference, wherein selection for one mutation interferes with selection for other

mutations. This phenomenon has been recognized for a long time, but has until now eluded quantification.

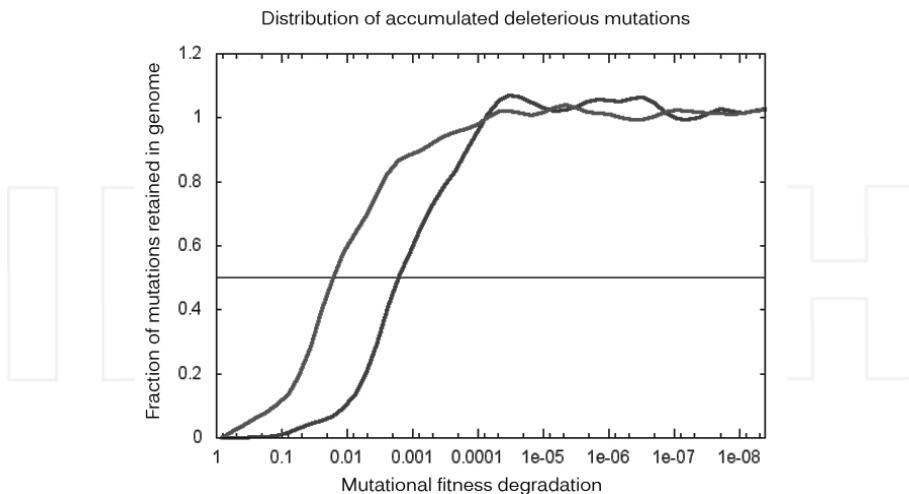


Fig. 5. The effects of fitness effect on mutation accumulation. Selection breaks down for most low-impact deleterious mutations. In this Mendel experiment, 80% of all mutations were made recessive (with 5% expression in the heterozygotic state) and 20% were made dominant (95% expression in the heterozygotic state). The rate of deleterious mutation accumulation (y axis) ranged from zero (no accumulation) to one (accumulation as if no selection). Mutational fitness effect is shown on the x-axis (log scale). As can be seen, deleterious mutations with very high impacts are selectively eliminated very effectively, but mutations with very low impacts are not affected by selection at all. The selection threshold (where the accumulation curve intersects 0.5, shown with a straight line) is where mutations are accumulating at half the rate they would in the absence of selection. In this experiment, the selection threshold is about one order of magnitude higher (curve to left) for recessive than for dominant mutations (curve to right). A similar selection threshold can be plotted for beneficial mutations.

The selection threshold value for deleterious mutations ( $ST_d$ ) is defined as the mutational fitness effect at which mutations accumulate at exactly half the rate as would occur if there were no selection (Figure 5). By parity of reasoning, the selection threshold value for beneficial mutations ( $ST_b$ ) is defined as the mutational fitness effect at which mutations accumulate at twice the rate as would occur if there were no selection (not shown). Mendel continuously monitors these selection threshold values during an experiment. We observe that the selection threshold is initially very high in all experiments, but drops dramatically in the first several hundred generations, and eventually approaches a (minimum) equilibrium value. The amount of time required to reach this “selection equilibrium” is strongly affected by population size, with large populations requiring deep time to reach their full selection potential (minimal selection threshold).

Selection threshold values are, to our knowledge, the only available diagnostic of how effectively selection can operate under a specific set of circumstances, in terms of eliminating

bad mutations and amplifying good mutations. It is worth noting that Kimura's (1983) well-known inequality,  $|s| \leq 1/(2N_e)$ , is an attempt to estimate the selection threshold based upon the random noise inherent in finite population sizes alone. However, many other parameters affect the efficacy of selection, including the mutation rate, the distribution of mutational fitness effects, environmental noise, mode of selection, and others. The final outcome of an experiment largely hinges on the emergent selection threshold values. This statistic is the best means to bring together all the "pieces of the population puzzle," enabling researchers to gauge the long-term genetic health of a population. A low threshold value should reflect a healthy population, allowing selection to "see" more low-impact mutations, while a high threshold value will reflect a population that is at risk of on-going genetic deterioration due to "selection breakdown" for most (low-impact) nucleotide positions in the genome.

#### 4.5 Net effect of linkage blocks

We know that chromosomes do not recombine uniformly, but have recombinational hotspots, which sub-divide each chromosome into numerous "linkage blocks". There is little recombination within a linkage block, so mutations that arise within the same linkage block will tend to be transmitted linked together indefinitely, from generation to generation. Mendel models this in a biologically realistic way, so that the effects of the linked mutations can be studied. This opens the way to study the phenomenon of "Muller's Ratchet" as it applies to individual linkage blocks. Mendel also allows the examination of the relative abundance of linkage blocks which have a net fitness gain versus a net fitness loss (Figure 6).

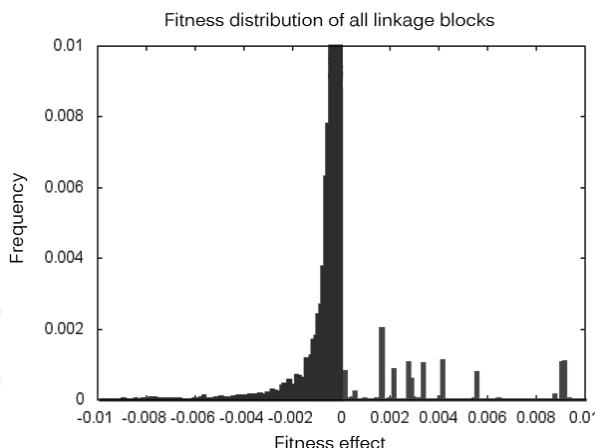


Fig. 6. The net fitness effects of linkage blocks reflect the distribution of accumulating mutations. This Mendel experiment employed the default parameters, except 20% of mutations were neutral, 1% were beneficial, and the maximal beneficial fitness effect was .01. Fitness was nearly stable and rising gradually. Linkage blocks with a net deleterious fitness effect are shown left of center, while linkage blocks with a net beneficial effect are shown right of center. As can be seen, almost all linkage blocks had a net effect which was modestly deleterious. Only 1.7% of all linkage blocks had a net beneficial effect, but that net beneficial effect was usually substantial.

#### 4.6 Population bottlenecks

In nature, populations routinely go through bottlenecks in population size. It has been proposed that this might be biologically useful in terms of genetically homogenizing a population. It has even been proposed that regular downward fluctuations in population size might help “pump out” or purge a population’s deleterious mutations. We have used Mendel to study episodes of population size contraction in mature populations, and we have observed that any episode that even marginally reduces total genetic variation simultaneously causes irreversible genetic damage. This is seen as a substantial fitness decline that does not fully recover when population size returns to normal, and which corresponds to an increased rate of fixation of deleterious mutations and an elevated deleterious selection threshold (Figure 7). It appears that it is very problematic to achieve homogenization of a mature out-crossing population by bottlenecking without risking population extinction. We have also used Mendel to examine cyclic bottlenecking. We observe that this does not “pump out” deleterious mutations, but rather “pumps in” such mutations due to elevated selection thresholds during each population contraction and a correspondingly higher rate of fixation of deleterious alleles.

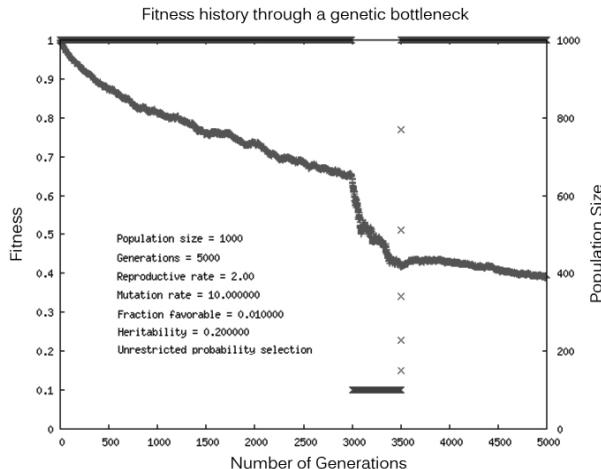


Fig. 7. The effects of population bottlenecks on average fitness. Population bottlenecks sufficient to have any noticeable impact cause irreversible genetic damage in out-crossing species. In this Mendel experiment, 1% of all mutations were beneficial, with the maximal beneficial fitness effect being .01. Eighty percent of mutations were recessive. After 3000 generations, population size was reduced from 1000 to 100 for 500 generations, after which population size was allowed to expand, restoring the population size (top line) to 1000 (scale on right). As can be seen, during the bottleneck fitness declined roughly 30% and failed to recover substantially when population size was restored.

#### 4.7 Allele frequencies

Mendel tracks allele frequencies. This allows the study of the rate of polymorphism (alleles with a frequency of more than 1%), and the rate of fixation (alleles with frequencies over

99%). As expected, Mendel shows that the vast majority of new alleles are lost by drift while they are still very rare. We see that the maximal number of polymorphic alleles is primarily limited by population size and population sub-structure (i.e., sub-populations that seldom inter-mate). Likewise, rate of fixation is profoundly affected by population size and population sub-structure. The rate of fixation is extremely slow except for relatively high-impact beneficial mutations, which can fix quite rapidly – in just hundreds of generations. Surprisingly, we routinely see that under the most biologically realistic conditions, many more deleterious mutations go to fixation than do beneficial mutations (because they arise at a much higher rate; Figure 8). Much more work remains to be done to better understand the determinants of polymorphism frequencies and rates of fixation.

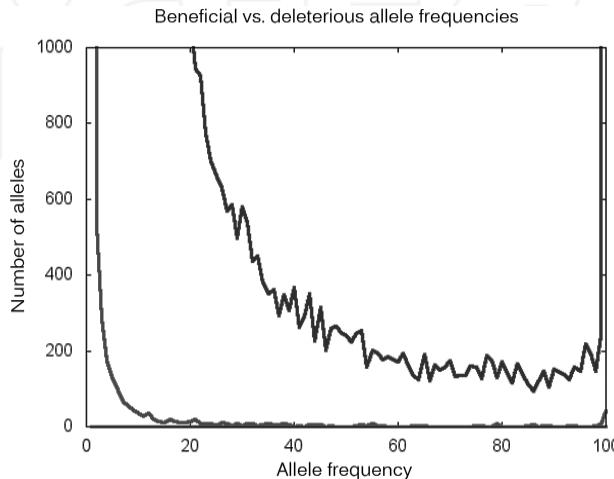


Fig. 8. Relative abundance of allele frequencies. When beneficial mutations are rare, the vast majority of fixation events tend to involve deleterious mutations (distribution to right).

Drawing from the same bottleneck experiment shown in Figure 7, it can be seen that allele frequencies are strongly skewed to the left, meaning that the vast majority of alleles present in the population were rare, as is consistently seen. This was particularly true in the case of beneficial mutations (distribution to left), which arise infrequently. Although a generous 1% of all new mutations in this experiment were beneficial, only 41 beneficial alleles went to fixation, while 9327 deleterious mutations went to fixation.

#### 4.8 Sexual reproduction

Much has been said about the biological importance of sexual recombination. This can most clearly be seen by using numerical simulation to contrast deleterious mutation accumulation in a normal sexual population and an identical population that reproduces asexually (Figure 9). The difference is very dramatic; the asexual population undergoes genetic degeneration very rapidly and the decline in fitness is distinctly linear, while in the sexual population, the fitness decline is modest and approximates exponential decay. These findings confirm expectation, because it has long been known that the absence of recombination in asexual genomes causes a gravely deterministic decay in fitness known as Muller's ratchet (e.g., see Loewe, 2006).

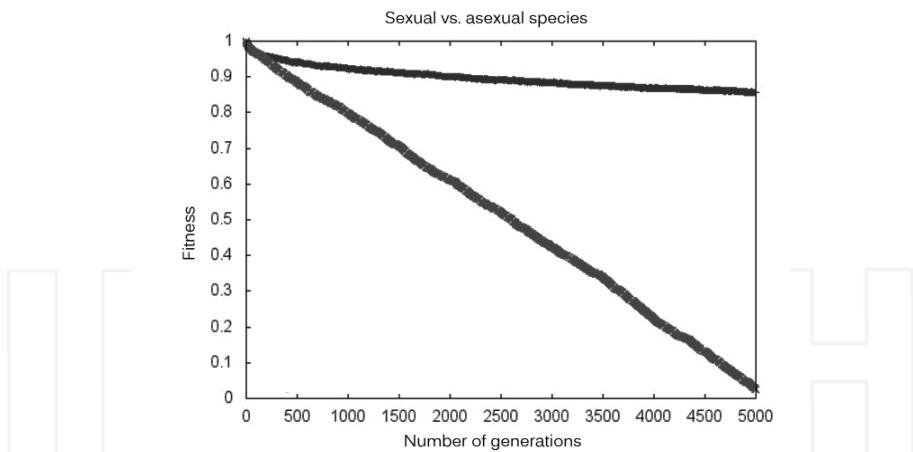


Fig. 9. The effects of sexual reproduction on fitness decline. Asexual populations are subject to disastrous fitness decline due to Muller's Ratchet. A Mendel experiment with default settings (except truncation selection was employed) was compared to the same run where reproduction was clonal, i.e., without sexual recombination. The upper line represents the first run using the default settings, resulting in a relatively modest level of fitness decline. The lower line represents the second run where sexual recombination was turned off, resulting in a rapid and very linear decline in fitness, causing extinction in just over 5000 generations.

#### 4.9 Future developments

In addition to the features and applications described above, other Mendel features recently developed or still under development include simulation of: a) synergistic epistasis; b) group selection; c) selection for altruistic traits; and d) analysis of specific sets of mutations that any user may upload into a population prior to simulation. Mendel is being modified to be compatible with most computer environments. Hopefully many researchers will use this program as a platform to develop far superior numerical simulations.

#### 5. Conclusion

It is clear that there is great utility in comprehensive numerical simulations, because they alone allow us to examine – simultaneously – the many elements of a given population's dynamics. Not only does this mean we can finally get an integrated “big picture” view of how a population changes, but we can use the same program to examine the same population in great detail from many specific vantage points. In light of the examples summarized above, it is clear there are numerous research problems which cannot be adequately addressed without comprehensive numerical simulation tools.

#### 6. Acknowledgments

This work was supported in part by the FMS Foundation and by Rainbow Technologies, Inc.

## 7. References

Brewer, W.; Baumgardner, J. & Sanford, J. (2012). Using Numerical Simulation to Test the "Mutation-Count" Hypothesis, In: *Biological Information - New Perspectives* (in press).

Carvajal-Rodríguez, A. (2008). Simulation of Genomes: A Review. *Current Genomics*, Vol. 9, (May 2008), pp. 155-159.

Carvajal-Rodríguez, A. (2010). Simulation of Genes and Genomes Forward in Time. *Current Genomics*, Vol. 11, (March 2010), pp. 58-61.

Chadeau-Hyam, M.; Hoggart, C.J.; O'Reilly, P.F.; Whittaker, J.C.; De Iorio, M. & Balding, D.J. (2008). Fregene: Simulation of Realistic Sequence-Level Data in Populations and Ascertained Samples. *BMC Bioinformatics*, Vol. 9, (September 2008), p. 364.

Crow, J.F. (1997). The High Spontaneous Mutation Rate: Is it Health Risk? *Proceedings of the National Academy of Sciences USA*, Vol. 94, (August 1997), pp. 8380-8386.

Eyre-Walker, A. & Keightley, P.D. (2007). The Distribution of Fitness Effects of New Mutations. *Nature Reviews Genetics*, Vol. 8, (August 2007), pp. 610-618.

Gibson, P.; Baumgardner, J.; Brewer, W. & Sanford, J. (2012). Can Biological Information Be Sustained By Purifying Natural Selection? In: *Biological Information - New Perspectives* (in press).

Hey, J. (2009). FPG - A Computer Program For Forward Population Genetic Simulation, 20.07.2011, Available from [http://lifesci.rutgers.edu/~heylab/ProgramsandData/Programs/FPG/FPG\\_Documentation.htm](http://lifesci.rutgers.edu/~heylab/ProgramsandData/Programs/FPG/FPG_Documentation.htm)

Hoggart, C.J.; Chadeau-Hyam, M.; Clark T.G.; Lamariello, R.; Whittaker, J.C.; De Iorio, M. & Blading, D.J. (2007). Sequence-Level Population Simulations Over Large Genomic Regions. *Genetics*, Vol. 177, (November 2007), pp. 1725-1731.

Hughes, A.L. (2008). Near Neutrality: Leading Edge of the Neutral Theory of Molecular Evolution. *Annals of the New York Academy of Sciences*, Vol. 1133, (2008), pp. 162-179.

Kim, Y. & Wiehe, T. (2009). Simulation of DNA Sequence Evolution Under Models of Recent Directional Selection. *Briefings in Bioinformatics*, Vol. 10, (January 2009), pp. 84-96.

Kimura, M. (1983). *The Neutral Theory of Molecular Evolution*, Cambridge University Press, ISBN 0-521-31793-2, Cambridge.

Kondrashov, A.S. (1995). Contamination of the Genome by Very Slightly Deleterious Mutations: Why Have We Not Died 100 Times Over? *Journal of Theoretical Biology*, Vol. 175, (August 1995), pp. 583-594.

Lenski, R.E.; Ofria, C.; Pennock, R.T. & Adami, C. (2003). The Evolutionary Origin of Complex Features. *Nature*, Vol. 423, (May 2003), pp. 139-144.

Li, H. & Durbin, R. (2011). Inference of Human Population History From Individual Whole-Genome Sequences. *Nature*, Vol. 475, (July 2011), pp. 493-496.

Liu, Y.; Athanasiadis, G. & Weale, M.E. (2008). A Survey of Genetic Simulation Software For Populations and Epidemiological Studies. *Human Genomics*, Vol. 3, (September 2008), pp. 79-86.

Loewe, L. (2002). Evolution@home: Global Computing Quantifies Evolution Due To Muller's Ratchet. *BMC Bioinformatics*, Vol. 6, Suppl. 3, (September 2005), p. P18.

Loewe, L. (2006). Quantifying the Genomic Decay Paradox Due To Muller's Ratchet in Human Mitochondrial DNA. *Genetical Research*, Vol. 87, (April 2006), pp. 133-159.

Lynch, M. (2010). Rate, Molecular Spectrum, and Consequences of Human Mutation. *Proceedings of the National Academy of Sciences USA*, Vol. 107, No. 3, pp. 961-968.

Muller, H.J. (1950). Our Load of Mutations. *American Journal of Human Genetics*, Vol. 2, (1950), pp. 111-176.

Nelson, C.W. & Sanford, J.C. (2011). The Effects of Low-Impact Mutations in Digital Organisms. *Theoretical Biology and Medical Modelling*, Vol. 8, (April 2011), p. 9.

Ofria, C. & Wilke, C.O. (2004). Avida: A Software Platform for Research in Computational Evolutionary Biology. *Artificial Life*, Vol. 10, No. 2, (2004), pp. 191-229.

Padhukasahasram, B.; Marjoram, P.; Wall, J.D.; Bustamante, C.D. & Nordborg, M. (2008). Exploring Population Genetic Models With Recombination Using Efficient Forward-Time Simulations. *Genetics*, Vol. 178, (April 2008), pp. 2417-2427.

Peng, B. & Amos, C.I. (2008). Forward-Time Simulations of Nonrandom Mating Populations Using simuPOP. *Bioinformatics*, Vol. 24, (June 2008), pp. 1408-1409.

Peng, B. & Amos, C.I. (2010). Forward-Time Simulations of Realistic Samples For Genome-Wide Association Studies. *BMC Bioinformatics*, Vol. 11, (September 2010), p. 442.

Peng, B. & Kimmel, M. (2005). simuPOP: A Forward-Time Population Genetics Simulation Environment. *Bioinformatics*, Vol. 21, (September 2005), pp. 3686-3687.

Peng, B. & Liu, X. (2010). Simuating Sequences of the Human Genome with Rare Variants. *Human Heredity*, Vol. 70, (January 2010), pp. 287-291.

Ray, T.S. (1991). An Approach to the Synthesis of Life, In: *Artificial Life II, SFI Studies in the Sciences of Complexity*, Vol. X: 1991, Langton, C.G.; Taylor, C.; Farmer, J.D. & Rasmussen, S., (Eds.), pp. 371-408, Addison-Wesley, Redwood City, CA.

Sanford, J.; Baumgardner, J.; Brewer, W.; Gibson, P. & ReMine, W. (2007a). Mendel's Accountant: A Biologically Realistic Forward-Time Population Genetics Program. *Scalable Computing: Practice and Experience*, Vol. 8, No. 2, (2007), pp. 147-165.

Sanford, J.; Baumgardner, J.; Gibson, P.; Brewer, W. & ReMine, W. (2007b). Using Computer Simulation to Understand Mutation Accumulation Dynamics and Genetic Load, In: *7th International Conference on Computational Science, Proceedings, Part II, LCNS 4488: 2007*, Shi, Y.; van Albada, G.D.; Dongarra, J. & Sloot, P.M.A., (Eds.), pp. 386-392, Springer-Verlag, Beijing, China.

Sanford, J.; Baumgardner, J. & Brewer, W. (2012). Selection Threshold Severely Constrains Capture of Beneficial Mutations, In: *Biological Information – New Perspectives* (in press).

Sauna, Z.E. & Kimchi-Sarfaty, C. (2011). Understanding the Contribution of Synonymous Mutations to Human Disease. *Nature Reviews Genetics*, Vol. 12, (October 2011), pp. 683-691.

Thorne, A.G. & Wolpoff, M.H. (1992). The Multiregional Evolution of Humans. *Scientific American*, Vol. 266, No. 4, (April 1992), pp. 76-83.

Wilson, A.C. & Cann, R.L. (1992). The Recent African Genesis of Humans. *Scientific American*, Vol. 266, No. 4, (April 1992), pp. 68-73.