Geonomics: Forward-Time, Spatially Explicit, and Arbitrarily Complex Landscape Genomic Simulations

Drew E. Terasaki Hart ,* Anusha P. Bishop , and Ian J. Wang

Department of Environmental Science, Policy, and Management, College of Natural Resources, University of California, Berkeley, CA, USA

***Corresponding author:** E-mail: drew.hart@berkeley.edu.

Associate editor: Melissa Wilson

Abstract

Understanding the drivers of spatial patterns of genomic diversity has emerged as a major goal of evolutionary genetics. The flexibility of forward-time simulation makes it especially valuable for these efforts, allowing for the simulation of arbitrarily complex scenarios in a way that mimics how real populations evolve. Here, we present Geonomics, a Python package for performing complex, spatially explicit, landscape genomic simulations with full spatial pedigrees that dramatically reduces user workload yet remains customizable and extensible because it is embedded within a popular, general-purpose language. We show that Geonomics results are consistent with expectations for a variety of validation tests based on classic models in population genetics and then demonstrate its utility and flexibility with a trio of more complex simulation scenarios that feature polygenic selection, selection on multiple traits, simulation on complex landscapes, and nonstationary environmental change. We then discuss runtime, which is primarily sensitive to landscape raster size, memory usage, which is primarily sensitive to maximum population size and recombination rate, and other caveats related to the model's methods for approximating recombination and movement. Taken together, our tests and demonstrations show that Geonomics provides an efficient and robust platform for population genomic simulations that capture complex spatial and evolutionary dynamics.

Key words: **landscape ecology, evolutionary genetics, population dynamics, environmental change, spatial modeling, Python.**

Introduction

Spatial patterns of genomic diversity result from the complex interplay of many underlying ecological and evolutionary processes and are shaped by a wide variety of geographic and environmental factors. Understanding how these patterns develop in natural systems has emerged as a primary goal of modern evolutionary genetics. These systems often occupy complex and potentially changing landscapes and might include populations that are not at demographic equilibrium. They may undergo neutral evolution as well as natural selection, sometimes on multiple traits of variable genetic architecture. The study of complex natural systems is crucial for developing evolutionary and ecological theory (Epperson et al. 2010; Barrett et al. 2019; Pelletier 2019), understanding the forces governing the evolution and maintenance of genetic diversity (Manel et al. 2003; Schoville et al. 2012), anticipating ecological futures in the Anthropocene (Bay et al. 2018; Capblancq et al. 2020), and informing conservation and management (Crossley et al. 2017; Lind et al. 2017). The complex genomics of many such systems are beyond the reach of analytical population genetics, and their spatial complexity and evolutionary dynamics make them intractable for coalescent simulation (Hoban et al. 2012). This hinders not only our understanding of many empirical systems but also our ability to predict their dynamics and, thus, to manage them. Hence, in population and landscape genomics, as in many other fields, forward-time simulation is a crucial tool for dissecting complex study systems.

However, the current suite of forward-time genomic simulators, although numerous, is still of limited utility for such work. Most available software is restricted, either genomically or geospatially, in the complexity, it can model. Many programs can model systems of considerable genomic complexity (e.g., simuPOP, Peng and Kimmel 2005; NEMO, Guillaume and Rougemont 2006; QuantiNemo, Neuenschwander et al. 2008) yet incorporate only rudimentary spatial components or none at all. Other programs are designed specifically for landscape genetic simulations (e.g., CDPOP, Landguth and Cushman 2010; CDMetaPOP, Landguth et al. 2017; SimAdapt, Rebaudo et al. 2013) but are limited in their genomic complexity. For instance, many programs are unable to model simultaneous selection on multiple, polygenic traits. To our knowledge, SLiM (Messer 2013; Haller and Messer 2017; Haller and Messer 2019) is the only package currently capable of simulating scenarios that are sufficiently complex, both genomically and geospatially, to model population genomic patterns emerging under dynamic evolutionary processes (according to a search of the National Cancer Institute's Genetic Simulation Resources website; Peng et al. 2013), and its extreme generalizability and complexity allow it to be used for landscape genomics simulation. Furthermore, many species are distributed continuously in space, and examining continuous fields of genetic variation can require distinct methods and assumptions (Bradburd and Ralph 2019), yet most population genomic simulation packages, aside from SLiM, are population-based. Such software requires individuals to be assigned to discrete subpopulations, which can at best be arranged on a high-resolution, regular grid in order to approximate continuously distributed populations.

Here, we present Geonomics, a Python package for forward-time, individual-based, continuous-space, population genomic simulations on complex landscapes. Geonomics models are parameterized by way of an informatively annotated parameters file that provides the user a straightforward means of building models of arbitrary complexity while offering reasonable default settings and "off switches" for parameters and components unrelated to the user's interests. Models consist of 1) a landscape with one or more environmental layers, each of which can undergo arbitrarily complex environmental change events and 2) one or more species having genomes with realistic architecture and any number of associated phenotypes. Species undergo non-Wright-Fisher evolution in continuous space, with localized mating and mortality, such that species-level phenomena and simulation dynamics are emergent properties of a model's parameterization. Evolution is comprehensively tracked by way of recently developed data structures that record the complete spatial pedigree (Kelleher et al. 2018), providing for the customizable output of rich, 3D data sets in a variety of common formats, including VCF and FASTA for genomic data, GeoTiff for landscape data, and CSV, Shapefile, and GeoJSON for individuals' nongenomic data (location, environmental values, phenotypes, age, and sex). All of this allows Geonomics to produce realistic landscape genomic results useful for a wide variety of theoretical and empirical purposes.

New Approaches

Model Design: Overview

A Geonomics model consists of two core components: the species and the landscape. The species is composed of a set of individuals and a wide variety of demographic and life-history parameters, including an intrinsic growth rate, mate-search radius, the mean number of offspring per mating event, reproductive age, and maximum age, among others. A species can undergo any number of change events, including changes to demographic and life history parameters and various types of population size changes. Each individual in the species has an *x*, *y* location, a sex, an age (or life-history stage), a set of phenotypes and a diploid genome consisting of any number of diallelic loci, which can represent either a contiguous haplotype block or a set of distinct loci. Loci can exhibit different types of dominance, and recombination rates can be heterozygous across the genome.

Phenotypic traits are continuous and quantitative and can be monogenic or multigenic. Each trait is defined by the loci that comprise its genetic basis, the effect sizes of those loci, and a phenotypic selection coefficient, which can be made heterogeneous in both space and time, allowing for spatially complex selection scenarios. While the strength of selection is determined by that coefficient, the force of selection is represented by the environmental raster layer to which the trait is adapting. Loci can have separate mutation rates for three types of mutations: neutral, deleterious, and trait-affecting mutations. Neutral mutations do not affect fitness, and deleterious mutations decrease fitness without affecting simulated phenotypic traits. Trait-affecting mutations, on the other hand, introduce mutations at previously unmutated loci mapped to a trait. This adds to the genetic variation affecting a trait, thus generating phenotypic variance upon which natural selection can operate. Mutation rates can be defined separately for each trait.

The other core component of a Geonomics model, the landscape, is a stack of raster layers. Each layer can be set to serve as one or more of 1) a resistance raster, which controls individual movement or offspring dispersal, 2) a carrying-capacity raster, which controls population density, and 3) a fitness raster for a trait, which governs natural selection. A key feature of Geonomics is that each layer can undergo any number of arbitrarily complex environmental change events which, as they unfold, influence the dynamics of any species whose carrying capacity, movement and dispersal, or fitness depend on the corresponding layer.

Model Operation: Overview

A Geonomics run begins with a burn-in stage during which individuals move and reproduce, without genomes or selection, until a series of statistical tests is passed. These tests include a time-lagged *t*-test and an augmented Dickey–Fuller test, which are run as a pair for 1) the total population size, serving as a test of temporal demographic stability, and 2) both the mean and the standard deviation of timestep-differenced cell-wise counts of individuals, serving as a test of spatial demographic stability. This burn-in period results in a stationary spatial distribution of the species on the landscape. Following burn-in, each individual has its genome randomly assigned according to the genomic architecture parameters, such that the main phase of each run begins with no pedigree and, thus, without population structure. Each time step in the main phase is a series of four operations, some optional (fig. 1):

1. movement (optional);
2. mating (requisite), which includes mate search, mate choice, offspring creation, and offspring dispersal;
3. mortality, which is due to density-dependence (requisite) and natural selection (optional);
4. change events (optional), including both environmental and demographic changes.

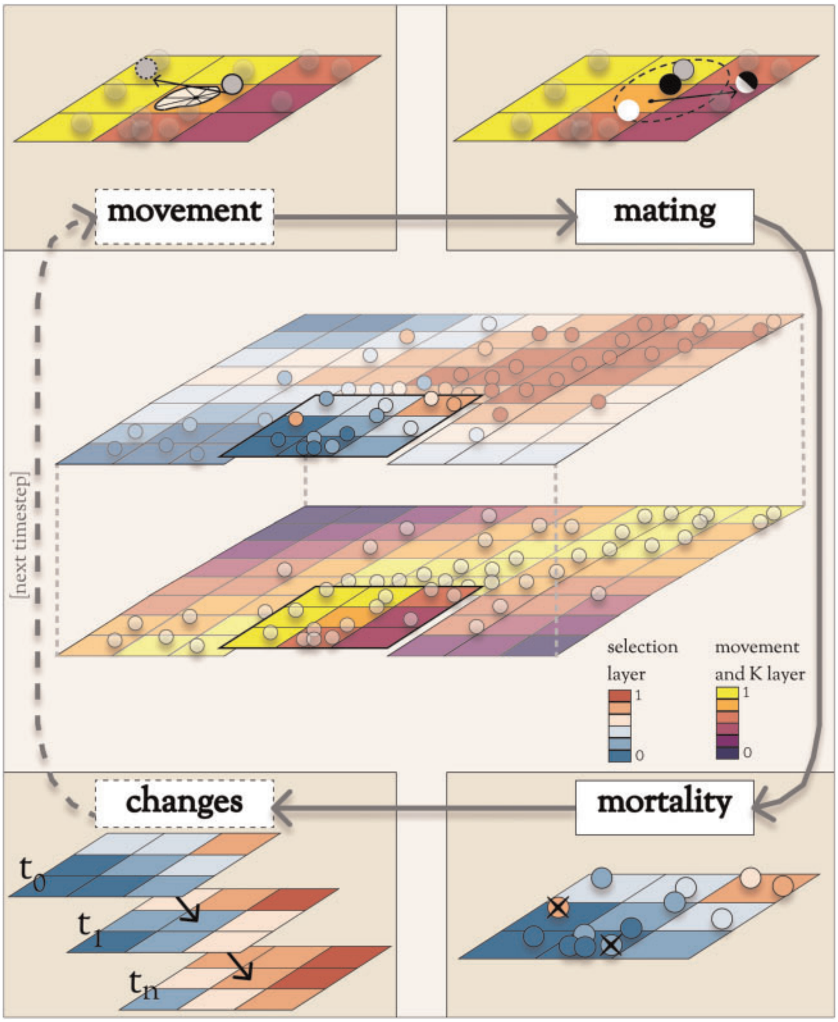

Fig. 1. Operations during the main phase of a Geonomics model run. In the center is a species on a multilayer landscape that includes a selection layer (above) and a layer for movement and carrying capacity (below). Surrounding the landscape is a flow-diagram of the major operations during a time step. Operations in dashed boxes are optional. During the movement stages (top-left), individuals move along movement vectors drawn from various distribution options (shown is an example of a cell-specific von Mises mixture distribution). During the mating stage (top-right), each mating individual (black circle) randomly chooses a mate (white circle) from all potential mates within its mating radius (dashed circle). The resulting offspring (half-black, half-white circle) disperses from its parents' midpoint along a randomly drawn dispersal vector. During the mortality stage (bottom-right), deaths are modeled as a Bernoulli process, with the probability of mortality a product of density-dependence and selection on all traits. During the changes stage (bottom-left), environmental and demographic change events, which can be represented by a series of change rasters corresponding to scheduled time steps, take place.

Model Operation: Movement

Movement takes place in continuous space—individuals have *x*, *y* coordinates, on either real or simulated landscapes, rather than being arbitrarily restricted to grid cells or bounded populations. Each individual moves along a vector, composed of a distance drawn from a Wald distribution and a direction drawn either from a uniform distribution on the unit circle or from a movement surface—an array of unimodal or multimodal von Mises distributions derived from a landscape layer that serves as a resistance surface (sensu McRae 2006; Spear et al. 2010). On a unimodal movement surface, each cell is assigned a single von Mises distribution, with mode parameter μ set to the direction of the highest-valued cell in the 8-cell neighborhood. On a multimodal surface, each cell's mixture distribution is a weighted sum of eight such unimodal distributions, one pointing toward the center of each cell in the 8-cell neighborhood and with normalized weights equal to the values of the neighboring cells. This approach to simulating movement generates realistic, anisotropic movement across a heterogeneous landscape (fig. 2) while avoiding time-consuming computational steps, such as repeated searches for minimum-resistance neighboring cells.

Model Operation: Mating

Potential mating pairs are randomly drawn from among all eligible pairs of individuals within the mate-search radius (unless strict nearest-neighbor mating is chosen), with pairing probabilities either uniform or inverse-distance weighted within the mating radius, and with eligibility based on both sex and age. From among those pairs, actual mating-event decisions are Bernoulli distributed, with probability equal to the intrinsic birth rate. Each mating pair produces a number

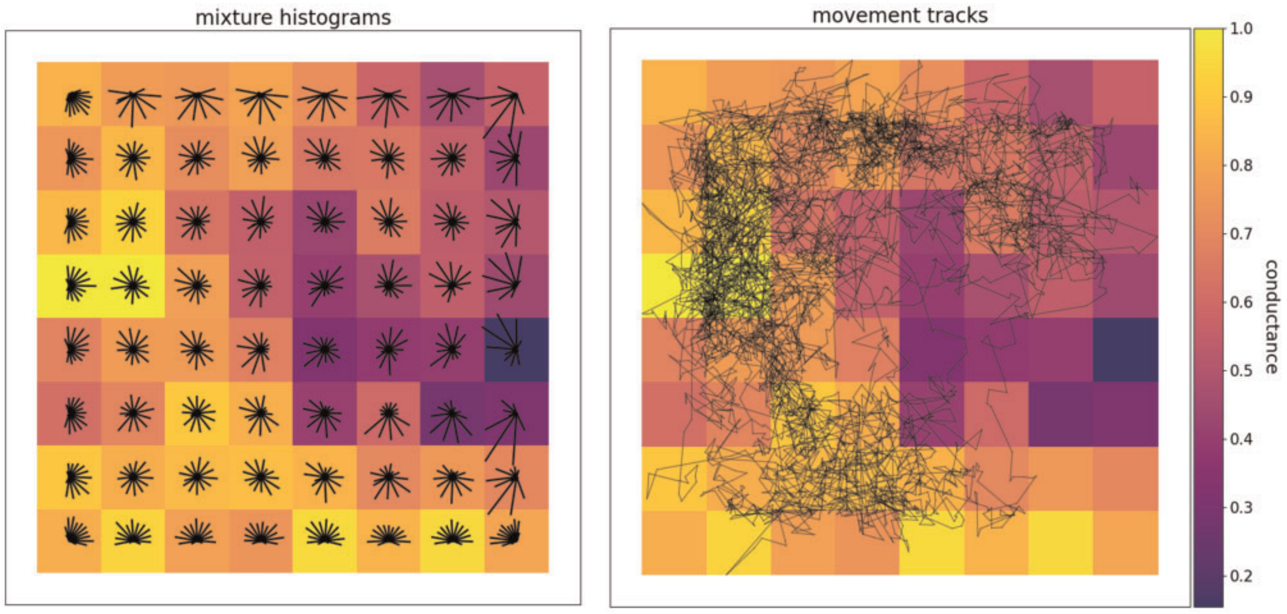

Fig. 2. A raster layer representing a movement surface with example movement histograms for each cell (left) and a movement track for a sample individual (right). The circularized histograms represent the movement directions that could be drawn from the von Mises mixture distribution approximations within each cell. Longer bars in a histogram indicate higher probability of movement in their direction. The movement track, plotted with the *gnx.help.param_help.plot_movement* function in Geonomics, is 5,000 steps long. Both preferential movement toward higher-suitability regions of the landscape (i.e., cells closer to 1 in value) and occasional long-distance movements between relatively isolated portions of the landscape are evident.

of offspring according to a fixed value or drawn from a Poisson distribution with λ equal to the mean number of offspring. Each parent produces a gamete for each of its offspring, using realistic recombination and Mendelian segregation. Gametes are united to create offspring individuals, which then disperse to new locations. As with movement, dispersal vectors can be drawn isotropically or anisotropically based on a resistance surface.

Model Operation: Mortality

Mortality is modeled as a Bernoulli process with the probability of an individual death a combination of the probabilities of death due to density-dependence (using a logistic-growth model) and due to natural selection (based on the cumulative fitness for all traits), calculated as:

$$P(d_i) = 1 - \left(1 - P\left(d_{x,y}\right)\right) \prod_{p=1}^{m} \omega_{i,p}, \tag{1}$$

where $P\left(d_{x,y}\right)$ is the probability of death due to density-dependence for individual i, m is the number of traits, and $\omega_{i,p}$ is the fitness of individual i for trait p. The probability of density-dependent death at location x,y is calculated as:

$$P\left(d_{x,y}\right) = \frac{E\left[N_{d;x,y}\right]}{N_{x,y}} = \frac{E\left[N_{b;x,y}\right] - \frac{dN_{x,y}}{dt}}{N_{x,y}}, \tag{2}$$

where, for location x,y, $E\left[N_{d;x,y}\right]$ is the expected number of deaths,$N_{x,y}$ is the population density, $E\left[N_{b;x,y}\right]$ is the expected number of births, and $\frac{dN_{x,y}}{dt}$ is the population logistic growth rate. The fitness of individual i for trait p is calculated as:

$$\omega_{i,p} = 1 - \phi_{p;x,y}\left(\left|e_{p;x,y} - z_{i;p}\right|\right)^{\gamma_p}, \tag{3}$$

where $\phi_{p;x,y}$ is the phenotypic selection coefficient on trait p at location x,y, $e_{p;x,y}$ is the value of the selection layer for trait p at location x,y, $z_{i;p}$ is the phenotype of individual i for trait p, and γ_p defines the curvature of the fitness function for trait p. The phenotype is a result of the additive effects of that individual's genotypes at all underlying loci, and is calculated as:

$$z_{i;p} = \sum_{l=0}^{n} \alpha_{p,l} g_{i,l} + g_0 \tag{4}$$

where n is the number of loci, $\alpha_{p,l}$ is the effect size of locus l on trait p, $g_{i,l}$ is the genotype at locus l for individual i, and g_0, the baseline genotype, equals 0 for monogenic traits or 0.5 for polygenic traits.

Model Operation: Change Events

Each demographic or environmental change event unfolds as a series of incremental changes that occur at the ends of scheduled time steps. A demographic change event can be exponential, random, or cyclical, or it can follow an arbitrarily complex, custom trajectory. Each event is parameterized by defining the time steps at which its changes take place and the factor by which the carrying capacity raster is multiplied.

Simple environmental change events are defined by a terminal raster for the final environmental state and a list of time steps at which incremental changes occur (based on cell-wise linear interpolation between the beginning and terminal states). More complex, custom events can be simulated by providing a series of environmental rasters labeled with the time step at which each will be applied. This option makes it easy to simulate evolution on real-world landscapes undergoing nonlinear, spatially heterogeneous environmental change.

Results

Validation

To validate the performance of Geonomics, we ran a series of simulations based on classical population genetic models, covering both neutral and nonneutral evolutionary scenarios. Because the classical models are simpler than the individual-based, spatially explicit, continuous-movement models built by Geonomics, we parameterized the simulations so as to accurately emulate these models while minimizing artifacts (see Validation Testing, Supplementary Material online). Our goal was to statistically and heuristically validate Geonomics' full range of functionality and to ensure that it accurately models neutral and nonneutral evolutionary processes.

To verify that Geonomics effectively models neutral evolution, we first examined the average time to fixation for a neutral allele in a finite population using simulations approximating a Wright-Fisher model (Fisher 1923; Wright 1930). We simulated allele-frequency trajectories for 250 independent loci (25 of which are plotted in supplementary fig. S1, Supplementary Material online), and we found that that fixation time did, indeed, increase with population size and was proportional to $4N_e$, as expected (Kimura and Ohta 1969), in our simulations (supplementary fig. S2, Supplementary Material online). We then tested for changes in the rate of drift surrounding a population bottleneck event by forcing a population to undergo a 70% reduction in size for 50 of 300 timesteps. We found that rates of allele frequency change increased during the bottleneck, then returned to prior levels shortly thereafter (supplementary fig. S3, Supplementary Material online). Finally, we quantified the accumulation of genetic structure under a stepping-stone model (Kimura 1953) to certify that genetic covariance decreases with distance (Kimura and Weiss 1964). As expected, we saw that migration rates decreased as a function of inter-island distance (supplementary fig. S4, Supplementary Material online), whereas F_{ST}, calculated from both heterozygosity data and genetic variance data, increased as a function of inter-island distance (and, therefore, decreased with pairwise migration rate) and as a function of time (supplementary figs. S4 and S5, Supplementary Material online). We also performed a discriminant analysis of principal components (DAPC) using the R package *adegenet* (Jombart et al. 2010) to confirm the expected population structure of six island clusters (supplementary fig. S6, Supplementary Material online).

To validate the performance of Geonomics for modeling nonneutral evolution, we first performed simulations under a simple scenario of divergent selection between two discrete habitats. As expected, simulations on a landscape evenly divided by two habitat blocks led to local adaptation, producing a significant pattern of phenotype-habitat matching (supplementary fig. S7, Supplementary Material online), with mismatches concentrated along the border between habitats. Additionally, over time, the species reached migration–selection equilibrium—the frequencies of the beneficial alleles in each habitat increased up to a stationary level, with that level being positively correlated with the strength of selection (supplementary fig. S8, Supplementary Material online). A plot of the mean difference between individuals' phenotypic and environmental values shows a strong decline over model time, with the rate and level of decline increasing as a function of increasing strength of selection (supplementary fig. S9, Supplementary Material online). Finally, logistic regressions show no significant relationships between phenotypic and environmental values at the outset (pseudo-R^2s $\approx$ 0.0, P values > 0.1), but show highly significant relationships at the ends of the simulations ($P < 0.0001$ for all values of Φ), with the amounts of variation explained increasing as a function of selection strength (pseudo-$R^2 = 0.327$ for $\Phi = 0.01$, 0.376 for $\Phi = 0.05$, and 0.406 for $\Phi = 0.1$).

We next tested the ability of Geonomics to recreate the genetic structure expected under local adaptation along an environmental cline: monotonic change in the allele frequency of a nonneutral locus across the cline. On a landscape with a symmetric environmental selection gradient, Geonomics again produced the expected spatial pattern of local adaptation (supplementary fig. S10, Supplementary Material online), and when we fitted sigmoid *tanh* clines (Szymura and Barton 1986; Porter 2013) for all loci, the locus underlying the monogenic trait was the only one to exhibit clinal variation (supplementary fig. S11, Supplementary Material online). In a family of genotype-environment analyses using Bonferroni-corrected, locus-wise logistic regressions, this locus was also the most significantly correlated with the environmental variable based on locus-wise logistic regressions ($P < 0.0001$). A plot of the mean difference between individuals' phenotypic and environmental values shows a strong decline over model time (supplementary fig. S12, Supplementary Material online), and logistic regressions show no significant relationship between phenotypic and environmental values at the outset (pseudo-$R^2 = 0$, P value $= 0.812$) but a significant relationship at the end of the simulation (pseudo-$R^2 = 0.169$, P value < 0.0001).

Finally, we verified that Geonomics can effectively model genomic data with physical linkage by simulating a selective sweep, introducing a beneficial mutation at the center of a 101-locus block of otherwise neutral loci. The results exhibited the classic genomic signal of a selective sweep (Przeworski 2002; Kim and Nielsen 2004), with a region of reduced nucleotide diversity surrounding the locus under selection and that region gradually eroding over time (supplementary fig. S12, Supplementary Material online). During these simulations, as the beneficial mutant spread through the population the population's mean fitness increased from $1 - \Phi$ (where Φ is the strength of selection) to a saturating value of 1 (supplementary fig. S14, Supplementary Material online), confirming that the population dynamics of the selective sweep played out as expected. To further support these results, we also validated Geonomics' method of recombination by examining effective recombination rates observed in a Geonomics model to those produced by an msprime simulation using the same randomly drawn, heterogeneous recombination map (see Recombination Test, Supplementary Material online); the resulting genome-wide pattern of recombination breakpoint densities recapitulates the one produced by msprime and the true recombination map (supplementary fig. S15, Supplementary Material online).

Example Applications: Overview

To demonstrate the broad utility of Geonomics for modeling complex evolutionary scenarios, we performed a series of simulations covering a range of potential applications. These demonstrations highlight scenarios for which Geonomics is particularly well suited, including spatially explicit simulations on highly heterogeneous landscapes, selection on multiple traits with complex genomic architecture, and microevolutionary responses to nonstationary environmental change.

Example 1: Isolation by Distance and by Environment

Genetic covariances between individuals or populations are often inversely correlated with linear or resistance-based geographic distance—a pattern known as isolation by distance (IBD; Wright 1943) or isolation by resistance (IBR; McRae 2006; McRae et al. 2008)—or with environmental distance—a pattern known as isolation by environment (IBE; Wang and Bradburd 2014). Understanding the landscape factors and population processes generating these patterns has emerged as a major focus of landscape genetics (Sexton et al. 2014; Wang and Bradburd 2014).

To demonstrate how Geonomics can simultaneously generate patterns of IBD and IBE, we built a simulation that uses a heterogeneous resistance layer as a movement surface and models selection for a 10-gene trait on a heterogeneous environmental layer ($\Phi = 0.05$). The model features a species with a stationary population size (roughly 2,450 individuals), experiencing both selection and neutral evolution. The resistance layer consists of a central barrier separating equal-area sides—the barrier has a high resistance to movement, but the movement is unconstrained on either side. This layer was also used as the carrying-capacity layer, yielding homogeneous population density on the two sides and zero density within the barrier region. The selection layer consists of two environmental gradients running in opposite directions on either side of the barrier, such that the landscape contains pairs of locations representing a range of combinations of geographic and environmental distances.

To observe the development of population structure, we collected data sets consisting of the genomes for all individuals at timesteps 0 and 1,000. We then used principal component analysis (PCA) to calculate pairwise genetic distances between all individuals for each data set. To visualize population structure, we extracted the first three principal components (PCs) and used them as the red, green, and blue (RGB) color values for mapping individuals on the landscape. To visualize the outcomes of selection, we produced paired maps of the same individuals colored by their phenotypes for the trait under selection (using Geonomics' "model.plot_phenotype(. . .)" method), and also created a set of the population–structure plots using DAPC. To visualize the time course of the simulation, we plotted the mean phenotype–environment mismatch (i.e., the mean of $|e\text{-}z|$, the driving force of selection) and mean fitness. We visualized signals of IBD and IBE in the final data set using a 3D scatterplot of Euclidean pairwise genetic distance against Euclidean pairwise geographic and environmental distances, colored by pairwise phenotypic distances. We tested the significance of the relationship between genetic distance and environmental distance, controlling for geographic distance, using paired partial Mantel tests with the *vegan* package (version 2.5-6; Oksanen et al. 2019) and using multiple matrix regression (MMRR; Wang 2013) in R version 4.0.2 (R Core Team 2020). Finally, because Geonomics models do not use defined landscape-resistance values, we quantified the barrier's increased landscape resistance by tracking all barrier-crossing events, using them to calculate the per time step crossing rate, then comparing that to the equivalent crossing rate of the same landscape zone in an otherwise identical model that omitted the barrier.

The RGB and phenotype plots of the initial population, with randomly assigned genomes, showed a clear lack of both spatial structure and local adaptation (fig. 3, top left). However, as expected, spatial structure developed over time, and the species showed signs of local adaptation over the course of the simulation (fig. 3, top right), as well as a corresponding, hierarchical population structure (supplementary fig. S15, Supplementary Material online). Average phenotype-environment mismatch decreased and average fitness increased over time (fig. 3, top middle). At the end of the simulation, the species demonstrated significant signals of both IBD (partial Mantel test: $r = 0.560$, $P \leq 0.001$; MMRR: $P \leq 0.001$) and IBE (partial Mantel test: $r = 0.121$, $P \leq 0.001$; MMRR: $P \leq 0.001$; MMRR full model $R^2 = 0.354$), as evidenced by the positive slopes on both horizontal axes of the 3D scatterplot (fig. 3, bottom). The colors of the points in the 3D scatter plot also indicate a clear pattern of increasing phenotypic differences between individuals with increasing environmental distance (fig. 3, bottom right) but not between individuals separated by increasing geographic distances (fig. 3, bottom left). Finally, the barrier zone had an observed crossing rate of 0.004 individuals per time step in this model, 13 times lower than the rate of 0.052 individuals per time step observed in the barrier-less but otherwise identical model. These results show that Geonomics effectively models IBD and IBE, driven by divergent natural selection, using two simple raster layers. More complex layers could be used to simulate IBD and IBE under a wide range of scenarios, and empirical layers could be used to simulate patterns of spatial genetic variation on real-world landscapes.

Example 2: Simultaneous Election

One of the most powerful features of Geonomics is that it can simulate selection on numerous traits simultaneously, each responding to a separate selection layer. Thus, a simulated species can experience multiple spatial selection regimes. Many natural systems are locally adapted to multiple

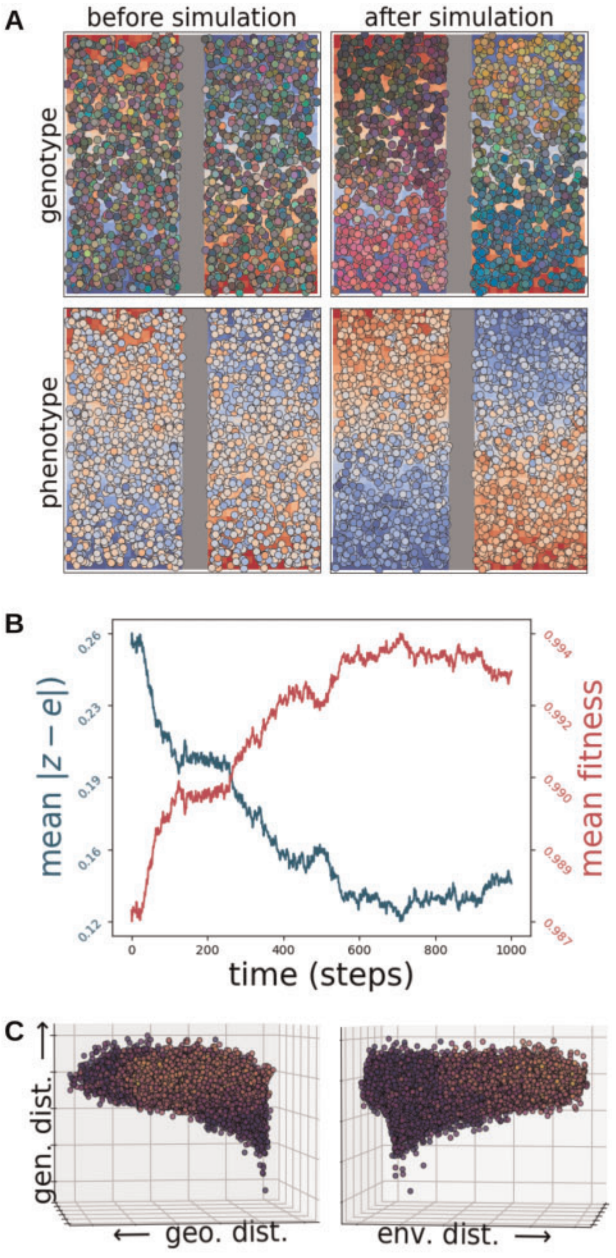

Fig. 3. Results of simulations for the isolation by distance (IBD) and isolation by environment (IBE) example application, in which a species evolved on a landscape with a barrier layer that served as the movement surface (displayed as a vertical gray band down the landscape) and an environmental layer that served as the selective surface for a 10-locus trait (displayed as the red to blue gradient on the landscape). (*A*) The population before the simulation (left column) and after it (right column), colored by genetic distance (top row), with colors derived from scores on the first three PCs of a genetic PCA used to assign RGB values, and by phenotype (bottom row). The most-fit individuals are those whose phenotypic colors perfectly match the cells on which they are located. (*B*) The time courses of the mean difference between individuals' phenotypes and their environmental values (blue) and of mean fitness values (red). (*C*) Two views of a 3D scatter plot of pairwise genetic distance as a function of Euclidean geographic distance (left) and Euclidean environmental distance (right), with points colored by phenotypic distance.

environmental variables (Fournier-Level et al. 2011; Lasky et al. 2012; Manel et al. 2012), so simulating these scenarios could be broadly valuable for investigating the nature of local adaptation in real environments.

To demonstrate how Geonomics can model simultaneous selection, we simulated a scenario in which a species undergoes natural selection along two orthogonal environmental gradients, each driving selection for a separate trait ($\Phi = 0.05$). Each trait had values ranging from 0 to 1, determined by 10 loci, all with equal effect sizes. Individuals had a mean movement distance of 0.5 cell widths on a 50 × 50-cell landscape, chosen to limit gene flow and allow for the development of strong spatial structure and, thus, the potential for local adaptation. We let the system evolve for 1,000 time steps and then mapped the species on each of the environmental layers, with individuals colored by phenotype in order to visually evaluate whether individual phenotypes matched their environmental backgrounds. The results showed clear patterns of phenotype–environment matching along both independent gradients (fig. 4) that evolved steadily through time (supplementary fig. S16, Supplementary Material online; compare to supplementary figs. S17 and S18, Supplementary Material online, with $\Phi = 0$), indicating strong evidence for simultaneous selection across the simulated landscape.

Example 3: Polygenic Adaptation to Climate Change in the Yosemite Region

Better understanding evolutionary responses to changing environments is essential for predicting species outcomes and preserving biodiversity under ongoing climate change (Hoffmann and Sgrò 2011; Franks and Hoffmann 2012; Bay et al. 2018; Capblancq et al. 2020). In many regions, climate shifts are projected to be spatially heterogeneous, including in montane regions where cooler, higher-altitude areas are warming more quickly than warmer, low-altitude regions (Rangwala et al. 2013; Mountain Research Initiative EDW Working Group 2015; but see Oyler et al. 2015). Of particular interest under these scenarios is the ability of species to adapt to changing local conditions (Franks and Hoffmann 2012).

To demonstrate the utility of Geonomics for studying microevolutionary responses to climate change, we simulated the response to projected climate change of a continuously distributed, locally adapted species, using the sagebrush lizard (*Sceloporus graciosus*) in the topographically complex Yosemite National Park region of California (USA) as an empirical model. To model climate change, we assembled time series raster stacks of projected mean annual temperature, annual precipitation, and habitat suitability for 19 even time steps from the present through the year 2100. For present temperature and precipitation, we used PRISM data (Daly et al. 2008), calculated as 30-year normals for 1981–2010 at 800 m resolution. For future years, we used means at a set of 5-year intervals (2015–2100), downscaled to 6 km resolution using the localized constructed analogs downscaling technique (LOCA; Pierce et al. 2014), from the Cal-Adapt database

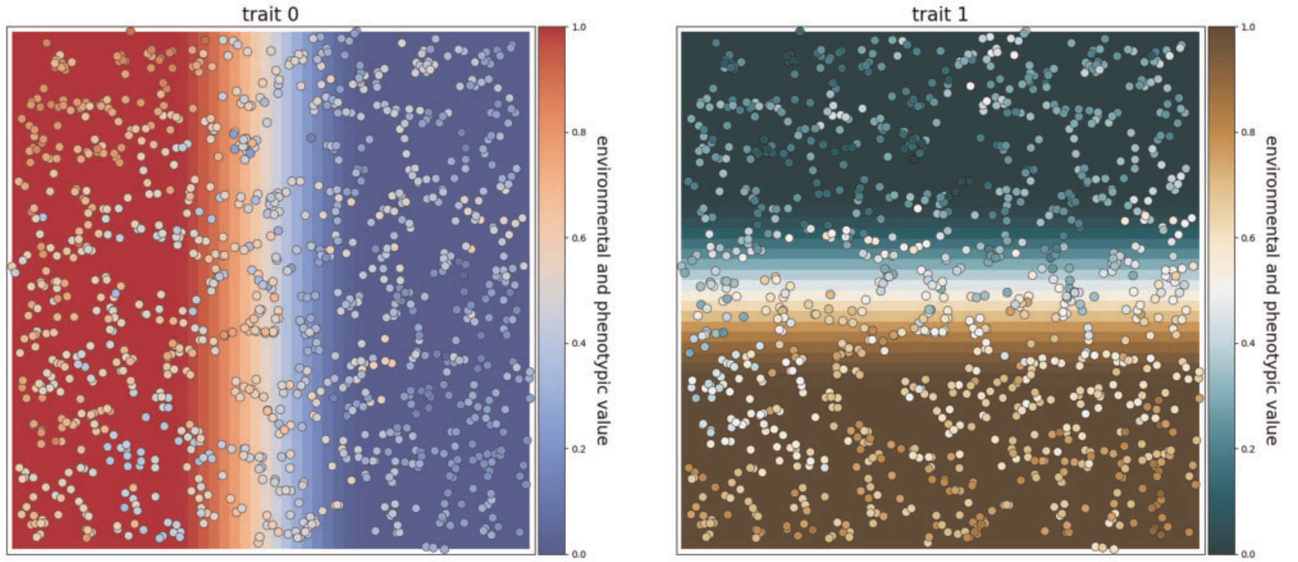

Fig. 4. Results of simultaneous selection on two traits with spatially distinct selective regimes. Each trait is controlled by 10 unlinked loci and has a selection coefficient of $\Phi = 0.05$. Individuals are colored by phenotype for the trait under selection on each layer.

(https://cal-adapt.org/). We calculated means of both variables from their minima and maxima observed across 32 global climate models, using a conservative representative concentration pathway (RCP 4.5). We developed time series of future temperature and precipitation layers at 800-m resolution by 1) calculating the raster difference between the first projected year and the current data, aggregated to the projected data's resolution; 2) adding that difference to the current data, such that each cell in the current data received the difference of the coarser, projected cell within which it lay; and 3) repeating that process for all remaining years. All data preparation was done using custom scripts (Supplementary Material online) in R (R Core Team 2020).

For the habitat suitability rasters, we constructed a species distribution model (SDM) using the present-day temperature and precipitation variables. We downloaded all georeferenced *S. graciosus* occurrence data from the Global Biodiversity Information Facility database (www.gbif.org), using the *gbif* function in the *dismo* R package (Hijmans et al. 2017). We clipped the points to California and Nevada, then subsampled the full data set to remove multiple points within the same raster cells. We generated pseudoabsence data by drawing random points from all cells in the California-Nevada region where the species was not observed (following the recommendations of Barbet-Massin et al. 2012). We extracted the current temperature and precipitation data at these points and used them as predictor variables in a binomial generalized linear model (GLM) with a logit link. We then projected that GLM onto the current and future temperature and precipitation rasters for our study region, producing a time series of predicted habitat suitability.

We generated the simulation's parameters file using the code provided in Code Sample S1, then edited the parameter values therein as needed. To simulate the nonneutral evolution of a polygenic, quantitative trait, we set the trait to be underlain by 100 loci randomly distributed across a genome of 1,000 loci and set a strength of selection of $\Phi = 0.5$. We set other life-history and demographic parameters (carrying capacity, age at reproductive maturity, number of offspring per individual, and maximum age) to reasonable values based on *S. graciosus* natural history (Stebbins 1948; Tinkle 1973; Rose 1976; Ruth 1978; Tinkle et al. 1993; Supplementary Material online).

We ran the main phase for 500 time steps without climate change (to develop a pattern of local adaptation), then ran an additional 100 time steps (years) with changing climate (see Code Sample S2). At time steps 500 (before the initiation of climate change events) and 600 (after completion of climate change events), we plotted the current temperature and habitat-suitability landscape layers along with a kriged surface of the current population's phenotypes and a kernel density map of the current population's density, two key emergent properties of the model that should be driven by temperature and habitat suitability, respectively. We then ran the model for an additional 50 timesteps to be able to more clearly visualize the effect of climate change on population size.

The model generated a clear and realistic pattern of adaptation to the spatial temperature gradient in the Yosemite region after the 500 iterations following burn-in, and that pattern demonstrated a spatial shift in phenotypes that aligns clearly with the spatial shift in temperature under the simulated climate change scenario (fig. 5A, rows 1 and 2; Video 1, Supplementary Material online). The model also generated a spatial pattern of population density that clearly aligns with spatial variation in habitat suitability prior to the onset of climate change that likewise shifted as expected in response to the climate change-induced shift in habitat suitability (fig. 5A, rows 3 and 4). We observed demographic changes in response to climate change over the course of the simulation as well. After climate change, mean population size was reduced by roughly 16.8% (from about 255,500 to 212,500 individuals), in line with the 17.9% reduction in the carrying capacity layer derived from the habitat suitability rasters (from 337,089.0 to 276,742.8 individuals, according to sums of the pre- and post-change carrying capacity layers). The population also exhibited sizable fluctuations during the climate change period, with oscillations exceeding 60,000 individuals (roughly 23.5% of the pre-change mean population size; fig. 5B). We interpret this as a result of the stepwise environmental changes comprising the climate change event. Each change causes a shift in the optimum phenotypes of local populations, leading to increased maladaptation and thus increased mortality rates. Subsequent reductions in density-dependent mortality rates because of these reduced population densities, paired with adaptation by natural selection, then reduce overall mortality rates, leading to rebounds in population size, with stochastic movement into and out of local populations, along with other sources of model stochasticity, imposing noise on this oscillatory behavior. Overall, these results show how Geonomics can effectively simulate organismal responses to highly complex environmental scenarios and reveal that these simulations can uncover system behavior that could provide avenues for future investigation.

Discussion

Our validations tests demonstrate that Geonomics simulates molecular evolution in concordance with predictions from theoretical population genetics (Fisher 1923; Wright 1943; Kimura; Szymura and Barton 1986), including dynamics of genetic drift, migration, and selection along clines, and our example applications show that Geonomics is capable of generating accurate and realistic population and landscape genomic data sets under scenarios of varying complexity. Geonomics is embedded in Python (van Rossum 1995; Python Software Foundation 2019), one of the most popular programming languages and one already familiar to many researchers who use bioinformatics. It makes the creation of arbitrarily complex models quick and easy, without even requiring prior Python experience, yet provides advanced users with access to core data structures, enabling broad customization and extension.

Many theoretical questions in population genomics necessitate explicitly spatial study methods, often with full tracking of a population's spatial pedigree (Bradburd and Ralph 2019). Geonomics makes this work more tractable than ever before. Landscape genomics studies draw conclusions about complex, real-world systems, sometimes with direct implications for conservation and management (Epperson et al. 2010; Landguth et al. 2012). Geonomics not only enables the generation of simulated data sets specific to such study systems

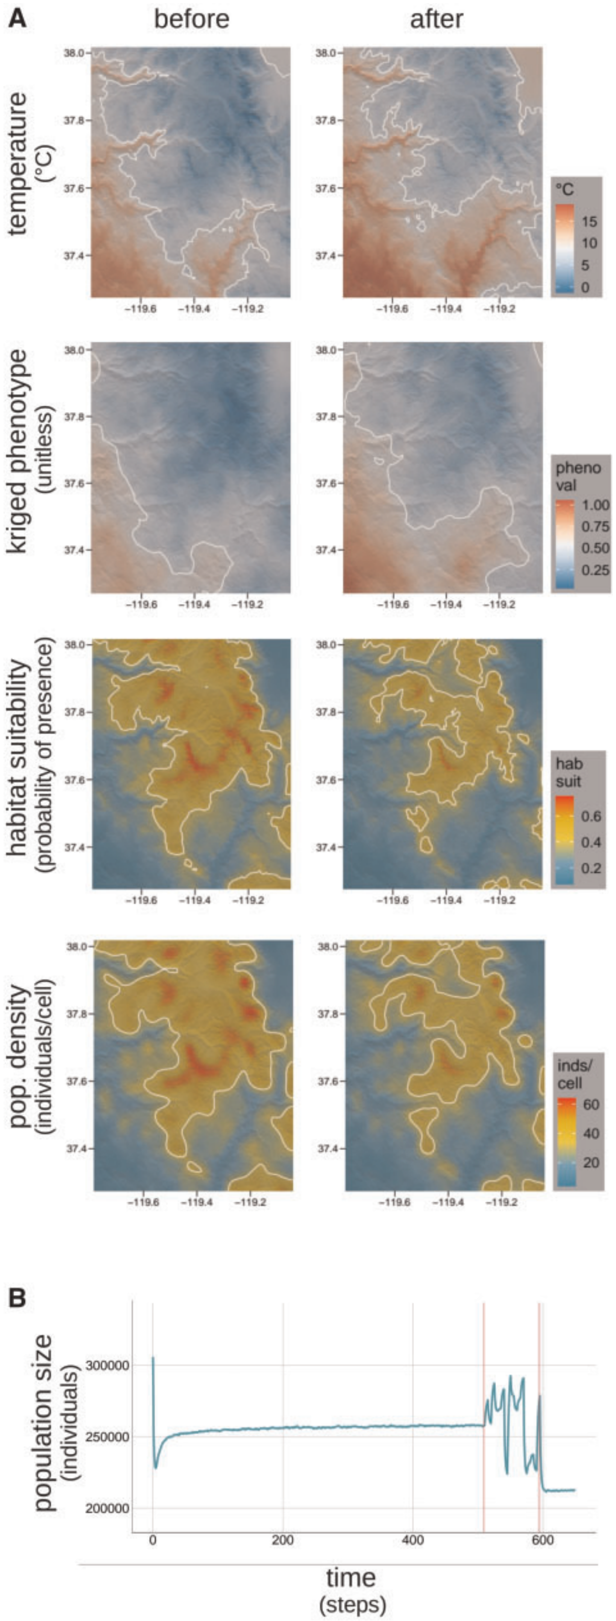

Fig. 5. Polygenic adaptation to climate change in the Yosemite region. (*A*) Hillshade plot comparisons of key variables (mean temperature, phenotype, habitat suitability, and population density) before and after the simulated climate change event. The mean of each variable through time is used to draw midvalue contours on each map (white lines) to help visualize spatial change. As expected, the spatiotemporal shift in temperature (first row) drives a spatially corresponding shift in phenotypes, visualized as a surface kriged from all phenotypic values (second row), and the shift in habitat suitability (third row) likewise drives a corresponding shift in population density, visualized using a 2D kernel density estimator (fourth row). (*B*) The time course of population size. The early drop in population size results from the onset of natural selection after completion of the unplotted burn-in portion of the model. The oscillations and ultimate reduction at the end of the simulation are a result of the climate change event, which occurs during the period bracketed by vertical red lines.

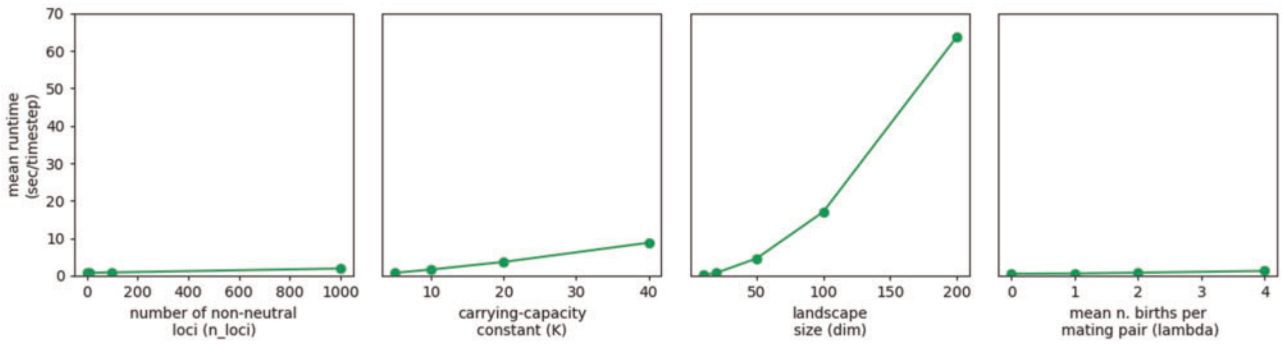

Fig. 6. Average runtime as a function of four major parameters: the number of nonneutral loci (*n_loci*), the carrying-capacity constant (*K*) that determines mean population size, landscape size (*dim*), and the λ parameter of the Poisson distribution from which the number of offspring in each mating event is drawn (*lambda*). Runtime increases with landscape dimension both because of functions whose runtimes scale with the landscape size directly and because of functions whose runtime scales with total population size. The difference between the lines for *K* and *dim* can be taken as an indication of the runtime cost of landscape dimension above and beyond population size effects, which predominates because it is superlinear.

but will also aid the development and testing of analytical methods in landscape genomics, strengthening our ability to draw accurate and reliable inferences from real-world data.

Runtime and Memory

Geonomics models run more slowly and have steeper memory limitations than models written and optimized in compiled languages, such as SLiM (Messer 2013; Haller and Messer 2017; Haller and Messer 2019). However, for users whose scenarios are well served by the design and affordances of Geonomics, what is sacrificed in runtime will be made up for in flexibility, customizability, and ease of use. With a reasonably powerful computer and for moderately sized models, most users should not find runtime or memory a major limitation. Indeed, our first two example applications were run on a laptop computer with 8 GB of RAM and an Intel® Core™ i5-8250U 3.4 GHz quad-core processor. Each run took an average of 271 s (about 0.27 s per time step) for the IBD-IBE model and 144 s (roughly 0.14 s per time step) for the simultaneous selection model. Because the polygenic adaptation example has much higher complexity, approximating the high population density of a small vertebrate, we ran it on a regular-memory node of a computing cluster (the *savio3* partition of UC Berkeley's Savio system) with 96 GB RAM and 2.1 GHz Skylake processors in order to handle the larger memory requirement. This model took considerably longer to run (approximately 7.35 h, running at about 32 s per time step after time for upfront computation of a series of changing movement surfaces) and had a peak memory usage of 25.433 GB, but even highly complex scenarios like this remain tractable on reasonable research timelines.

Given the complexity of Geonomics and the number of parameters a user can modify, numerous parameters and parameter combinations can influence a model's average runtime. We provide a basic runtime analysis (fig. 6), run on the same 8 Gb, quad-core laptop as the examples above. This analysis highlights some basic parameters that are likely to influence a model's average runtime per time step, including the mean population size (as determined by an array of local carrying capacities), the number of offspring per mating event, the size of the landscape, and the number of nonneutral loci in the genome. The effect of landscape size predominates, as runtime scales superlinearly with this parameter. The number of nonneutral loci actually has only slight effects on total runtime, and because neutral genetic data are stored in a set of tables rather than redundantly for each individual (Kelleher et al. 2018) runtime is even less sensitive to the number of neutral loci in the genome (although recombination rate does impact runtime and memory usage), meaning that Geonomics can efficiently simulate genome-scale data sets if provided adequate memory to store the set of recombination pathways that is calculated at the outset. Finally, after the upfront cost of computing recombination pathways and movement surfaces, runtime scales roughly linearly with the number of time steps, barring large demographic changes. That means that the moderately complex scenario in our simultaneous selection example could complete 1,000,000 time steps in approximately 40 h, and even the highly complex scenario in our polygenic adaptation example with more than 200,000 individuals could run through 10,000 generations in about 3.7 days. Hence, Geonomics could even prove useful for research at deeper timescales, for example in phylogeography or geogenomics (Baker et al. 2014).

Caveats

Geonomics uses two unconventional approximations to make complex models tractable within reasonable compute time in an interpreted language. The first is the approximation used to model heterogeneous recombination. Enacting recombination between all neighboring loci each time a gamete is produced would require an extremely large and time-consuming number of random draws. To avoid this, when a model is first created, Geonomics generates and saves, as binary arrays, a large collection of recombination "paths." The number of paths used is set by the user, and directly determines the minimum recombination rate difference that can be modeled. Each path is just a genome-length array that switches between 0 and 1 at each interlocus position where a recombination event should occur. The path can then quickly be used to subset an individual's genome, producing a gamete. As a model runs and gametes are continually produced, these paths are repeatedly shuffled and drawn through, like a deck of cards during multiple rounds of a game. This approach, which we have validated using msprime (see Recombination Test, Supplementary Material online), can lead to memory limitations for models with a large number of paths and a long genome, because the data structure containing the paths is essentially a 2D binary array whose size is the product of these values and because more recombination events require that more trees are recorded in the spatial pedigree. To avoid these problems, for genomic architectures with homogeneous recombination rates, Geonomics provides the option to use an alternative recombination mechanism that simulates recombination on the fly for each new gamete but does so at a cost of increased average runtime per time step.

The second is the approximation used to model the circular distributions from which movement directions are drawn. Conceptually, a movement or dispersal surface is an $x \times y$ array of Von Mises distributions. In practice, each distribution on that surface is represented by a column of angular directions (an "approximation column") drawn, at the time the model is built, from the true, continuous distribution. During a model run, to draw a movement direction from a cell, a random value is sampled from that cell's approximation column. This increases computational efficiency by avoiding large numbers of calls to random number generators during runtime. The accuracy of these approximation columns is a function of their length, which is set by the user. This length will usually not be so constrained that it significantly impacts the accuracy of the approximation, but such a constraint could arise if the movement or dispersal surface undergoes environmental change. In this case, the movement surfaces corresponding to each step of the change event will be generated and stored when the model is first created, and the series of arrays produced could exhaust memory if the landscape is very large and has many environmental change steps. A solution to this problem would feature some combination of decreasing the temporal resolution of the environmental change event, decreasing the landscape size, or decreasing the approximation column length. In all cases, users may check the accuracy of modeled movement by using built-in functions that visualize the composition and behavior of movement and dispersal surfaces.

Materials and Methods

We performed all simulations using the Geonomics Python package (see New Approaches) as described in the Results section and Supplementary Material. The simplest way to get started with Geonomics is to install it via pip. Geonomics uses common, well-established Python packages as required dependencies—Numpy (Harris et al. 2020), Matplotlib (Hunter 2007), Pandas (McKinney et al. 2010), Shapely (Gillies et al. 2007), Scipy (Virtanen et al. 2020), Scikit-learn (Pedregosa et al. 2011), Statsmodels (Seabold and Perktold 2010), Rasterio (Gillies et al. 2019), Bitarray (Schnell et al. 2021), msprime (Kelleher et al. 2016), and tskit (Kelleher et al. 2018)—and offers optional integration of neutral landscape models through the NLMpy package (Etherington et al. 2015). The source code is publicly available on GitHub (https://github.com/drewhart/geonomics), where it is actively maintained and developed.

Conclusions

Geonomics is a Python package designed to make building and running complex landscape genomic models quick and simple. At the same time, it provides a flexible scripting framework that allows advanced users to customize and extend its functionality. We believe Geonomics will prove highly useful for theoretical, empirical, methodological, and applied research in population and landscape genomics, molecular ecology, global change biology, and conservation.

Supplementary Material

Supplementary data are available at *Molecular Biology and Evolution* online.

Acknowledgments

We thank G. Bradburd, T. Dawson, N. Fefferman, M. Fitzpatrick, J. Frederick, N. Graham, M. Kelly, M. McElroy, R. Nielsen, T. O'Connor, D. O'Sullivan, C. Small, E. Westeen, N. Whiteman, G. Wogan, and M. Yuan for feedback and guidance on various iterations of the model that became Geonomics. We also thank I. R. Wang, B. Xu, J. Zhao for contributing to the Python code in its early stages and Berkeley Research Computing for access to the Savio computing cluster. Lastly, we thank M. Terasaki, C. Hart, G. Hart, J. Hart, and M. Tylka for supporting and encouraging a lifetime of curiosity about nature, and C. Adderly, K. Cobain, C. Debussy, L. Hill, S. Jeffes, N. Jones, G. R. Sessions, and T. Yorke for their continual company during many long hours of programming. This work was supported by a Berkeley Fellowship (to D.E.T.H.) and National Science Foundation grant DEB-1845682 (to I.J.W.).

Author Contributions

D.E.T.H. conceived and wrote the software, validations, and example applications, with contributions from I.J.W. D.E.T.H. co-wrote the manuscript and performed analyses. A.P.B. co-wrote the manuscript and contributed analyses. I.J.W. co-wrote the manuscript.

Data Availability

Geonomics is installable from the Python Package Index (https://pypi.org/project/geonomics/) using pip, and the Geonomics source code is available in the Geonomics GitHub repository (https://github.com/drewhart/geonomics). Parameterizations for all simulations run for this study are stored in their respective directories in the source code repository, and other associated data (e.g., raster layers) are available in the ./data directory of a local Geonomics installation. Ancillary code used for the preparation of simulations and execution of analyses for this paper live in a separate GitHub repository (https://github.com/drewhart/geonomics_methods_paper_ancillary_code). Validations tests can be run using the code provided in the respective directories in the source code repository. Demo simulations can be run by loading the Geonomics package into Python and then calling the run_demo function (e.g., "import geonomics as gnx; gnx.run_demo('IBD IBE')"). Using those materials, readers can reproduce stochastically varying but qualitatively equivalent results as those presented in this article.

References

Baker PA, Fritz SC, Dick CW, Eckert AJ, Horton BK, Manzoni S, Ribas CC, Garzione CN, Battisti DS. 2014. The emerging field of geogenomics: Constraining geological problems with genetic data. *Earth-Sci Rev.* 135:38–47.

Barbet-Massin M, Jiguet F, Albert CH, Thuiller W. 2012. Selecting pseudo-absences for species distribution models: how, where and how many? *Methods Ecol Evol.* 3(2):327–338.

Barrett RDH, Laurent S, Mallarino R, Pfeifer SP, Xu CCY, Foll M, Wakamatsu K, Duke-Cohan JS, Jensen JD, Hoekstra HE. 2019. Linking a mutation to survival in wild mice. *Science* 363(6426):499–504.

Bay RA, Harrigan RJ, Underwood VL, Gibbs HL, Smith TB, Ruegg K. 2018. Genomic signals of selection predict climate-driven population declines in a migratory bird. *Science* 359(6371):83–86.

Bradburd GS, Ralph PL. 2019. Spatial population genetics: it's about time. *Annu Rev Ecol Evol Syst.* 50(1):427–449.

Capblancq T, Fitzpatrick MC, Bay RA, Exposito-Alonso M, Keller SR. 2020. Genomic prediction of (mal)adaptation across current and future climatic landscapes. *Annu Rev Ecol Evol Syst.* 51(1):245–269.

Crossley MS, Chen YH, Groves RL, Schoville SD. 2017. Landscape genomics of Colorado potato beetle provides evidence of polygenic adaptation to insecticides. *Mol Ecol.* 26(22):6284–6300.

Daly C, Halbleib M, Smith JI, Gibson WP, Doggett MK, Taylor GH, Curtis J, Pasteris PP. 2008. Physiographically sensitive mapping of climatological temperature and precipitation across the conterminous United States. *Int J Climatol.* 28(15):2031–2064.

Epperson BK, McRae BH, Scribner K, Cushman SA, Rosenberg MS, Fortin MJ, James PMA, Murphy M, Manel S, Legendre P, et al. 2010. Utility of computer simulations in landscape genetics. *Mol Ecol.* 19(17):3549–3564.

Etherington TR, Holland EP, O'Sullivan D. 2015. NLMpy: a python software package for the creation of neutral landscape models within a general numerical framework. *Methods Ecol Evol.* 6(2):164–168.

Fisher RA. 1923. XXI.—on the dominance ratio. *Proc R Soc Edinb.* 42:321–341.

Fournier-Level A, Korte A, Cooper MD, Nordborg M, Schmitt J, Wilczek AM. 2011. A map of local adaptation in *Arabidopsis thaliana*. *Science* 334(6052):86–89.

Franks SJ, Hoffmann AA. 2012. Genetics of climate change adaptation. *Annu Rev Genet.* 46:185–208.

Gillies S. 2007. Shapely contributors. Shapely: manipulation and analysis of geometric objects. Available from: https://github.com/Toblerity/Shapely.

Gillies S, Rasterio contributors. 2019. Rasterio: geospatial raster I/O for Python programmers. Available from: https://github.com/mapbox/rasterio.

Guillaume F, Rougemont J. 2006. Nemo: an evolutionary and population genetics programming framework. *Bioinformatics* 22(20):2556–2557.

Haller BC, Messer PW. 2017. SLiM 2: flexible, interactive forward genetic simulations. *Mol Biol Evol.* 34(1):230–240.

Haller BC, Messer PW. 2019. SLiM 3: forward genetic simulations beyond the wright–fisher model. *Mol Biol Evol.* 36(3):632–637.

Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, Taylor J, Berg S, Smith NJ, et al. 2020. Array programming with NumPy. *Nature* 585(7825):357–362.

Hijmans RJ, Phillips S, Leathwick J, Elith J. 2017. *dismo*: species distribution modeling. R package version 1.1–4. Available from: https://CRAN.R-project.org/package=dismo.

Hoban S, Bertorelle G, Gaggiotti OE. 2012. Computer simulations: tools for population and evolutionary genetics. *Nat Rev Genet.* 13(2):110–122.

Hoffmann AA, Sgrò CM. 2011. Climate change and evolutionary adaptation. *Nature* 470(7335):479–485.

Hunter JD. 2007. Matplotlib: a 2D graphics environment. *Comput Sci Eng.* 9(3):90–95.

Jombart T, Devillard S, Balloux F. 2010. Discriminant analysis of principal components: a new method for the analysis of genetically structured populations. *BMC Genet.* 11(1):94.

Kelleher J, Etheridge AM, McVean G. 2016. Efficient coalescent simulation and genealogical analysis for large sample sizes. *PLoS Comput Biol.* 12(5):e1004842.

Kelleher J, Thornton KR, Ashander J, Ralph PL. 2018. Efficient pedigree recording for fast population genetics simulation. *PLoS Comput Biol.* 14(11):e1006581.

Kim Y, Nielsen R. 2004. Linkage disequilibrium as a signature of selective sweeps. *Genetics* 167(3):1513–1524.

Landguth EL, Bearlin A, Day CC, Dunham J. 2017. CDMetaPOP: an individual-based, eco-evolutionary model for spatially explicit simulation of landscape demogenetics. *Methods Ecol Evol.* 8(1):4–11.

Landguth EL, Cushman SA. 2010. cdpop: A spatially explicit cost distance population genetics program. *Mol Ecol Resour.* 10(1):156–161.

Landguth EL, Cushman SA, Johnson NA. 2012. Simulating natural selection in landscape genetics. *Mol Ecol Resour.* 12(2):363–368.

Lasky JR, Des Marais DL, McKAY JK, Richards JH, Juenger TE, Keitt TH. 2012. Characterizing genomic variation of *Arabidopsis thaliana*: the roles of geography and climate. *Mol Ecol.* 21(22):5512–5529.

Lind BM, Friedline CJ, Wegrzyn JL, Maloney PE, Vogler DR, Neale DB, Eckert AJ. 2017. Water availability drives signatures of local adaptation in whitebark pine (*Pinus albicaulis* Engelm.) across fine spatial scales of the Lake Tahoe Basin, USA. *Mol Ecol.* 26(12):3168–3185.

Manel S, Schwartz MK, Luikart G, Taberlet P. 2003. Landscape genetics: combining landscape ecology and population genetics. *Trends Ecol Evol.*18(4):189–197.

Manel S, Gugerli F, Thuiller W, Alvarez N, Legendre P, Holderegger R, Gielly L, Taberlet P. IntraBioDiv Consortium. 2012. Broad-scale adaptive genetic variation in alpine plants is driven by temperature and precipitation. *Mol Ecol.* 21(15):3729–3738.

McKinney W, Numpy contributors. 2010. Data structures for statistical computing in python. Proceedings of the 9th Python in Science Conference. p. 51–56.

McRae BH. 2006. Isolation by resistance. *Evolution* 60(8):1551–1561.

McRae BH, Dickson BG, Keitt TH, Shah VB. 2008. Using circuit theory to model connectivity in ecology, evolution, and conservation. *Ecology* 89(10):2712–2724.

Messer PW. 2013. SLiM: simulating evolution with selection and linkage. *Genetics* 194(4):1037–1039.

Mountain Research Initiative EDW Working Group. 2015. Elevation-dependent warming in mountain regions of the world. *Nat Clim Change.* 5:424–430.

Neuenschwander S, Hospital F, Guillaume F, Goudet J. 2008. quantiNemo: an individual-based program to simulate quantitative traits with explicit genetic architecture in a dynamic metapopulation. *Bioinformatics* 24(13):1552–1553.

Oksanen J, Guillaume Blanchet F, Friendly M, Kindt R, Legendre P, McGlinn D, Minchin PR, O'Hara RB, Simpson GL, Solymos P, et al. 2019. vegan: Community Ecology Package. R package version 2.5-6. Available from: https://CRAN.R-project.org/package=vegan.

Oyler JW, Dobrowski SZ, Ballantyne AP, Klene AE, Running SW. 2015. Artificial amplification of warming trends across the mountains of the western United States. *Geophys Res Lett.* 42(1):153–161.

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, et al. 2011. Scikit-learn: machine learning in Python. *J Mach Learn Res.* 12:2825–2830.

Pelletier F. 2019. Testing evolutionary predictions in wild mice. *Science* 363(6426):452–453.

Peng B, Kimmel M. 2005. simuPOP: a forward-time population genetics simulation environment. *Bioinformatics* 21(18):3686–3687.

Peng B, Chen HS, Mechanic LE, Racine B, Clarke J, Clarke L, Gillanders E, Feuer EJ. Genetic Simulation Resources: a website for the registration and discovery of genetic data simulators. *Bioinformatics*. 2013 29(8):1101–1102.

Pierce DW, Cayan DR, Thrasher BL. 2014. Statistical downscaling using localized constructed analogs (LOCA). *J Hydrometeorol*. 15(6):2558–2585.

Porter A. 2013. ClineFit v. 2.0, user's manual. Available from: http://people.umass.edu/aporter/ClineFit/ClineFit%20Manual.pdf. Accessed June 24, 2021.

Przeworski M. 2002. The signature of positive selection at randomly chosen loci. *Genetics* 160(3):1179–1189.

Python Software Foundation. 2019. Python language reference, version 3.7. Available from: http://www.python.org. Accessed June 24, 2021.

R Core Team. 2020. R: a language and environment for statistical computing.Vienna, Austria: R Foundation for Statistical Computing.

Rangwala I, Sinsky E, Miller JR. 2013. Amplified warming projections for high altitude regions of the northern hemisphere mid-latitudes from CMIP5 models. *Environ Res Lett.* 8(2):024040.

Rebaudo F, Le Rouzic A, Dupas S, Silvain J-F, Harry M, Dangles O. 2013. SimAdapt: an individual-based genetic model for simulating landscape management impacts on populations. *Methods Ecol Evol*. 4(6):595–600.

Rose BR. 1976. Habitat and prey selection of *Sceloporus occidentalis* and *Sceloporus graciosus*. *Ecology* 57(3):531–541.

Ruth SB. 1978. A comparison of the demography and female reproduction in sympatric western fence lizards (*Sceloporus occidentalis*) and sagebrush lizards (*Sceloporus graciosus*) on Mount Diablo, California. Ph.D dissertation: University of California Berkeley. Available from: https://www.proquest.com/openview/4fc8add6870959ae6228fe7ce5f7cb47. Accessed November 20, 2020.

Schoville SD, Bonin A, François O, Lobreaux S, Melodelima C, Manel S. 2012. Adaptive genetic variation on the landscape: methods and cases. *Annu Rev Ecol Evol Syst*. 43(1):23–43.

Schnell I, Bitarray contributors. 2021. bitarray: efficient arrays of booleans. Available from: https://github.com/ilanschnell/bitarray. Accessed June 24, 2021.

Seabold S, Perktold J. 2010. statsmodels: econometric and statistical modeling with python. In: van der Walt S, Millman J, editors. Proceedings of the 9th Python in Science Conference (SciPy 2010). Austin (TX): SciPy. p. 92–96.

Sexton JP, Hangartner SB, Hoffmann AA. 2014. Genetic isolation by environment or distance: which pattern of gene flow is most common? *Evolution* 68(1):1–15.

Spear SF, Balkenhol N, Fortin M-J, Mcrae BH, Scribner K. 2010. Use of resistance surfaces for landscape genetic studies: considerations for parameterization and analysis. *Mol Ecol*. 19(17):3576–3591.

Stebbins RC. 1948. Additional observations on home ranges and longevity in the lizard *Sceloporus graciosus*. *Copeia* 1948(1):20–22.

Szymura JM, Barton NH. 1986. Genetic analysis of a hybrid zone between the fire-bellied toads, *Bombina bombina* and *B. variegata*, near Cracow in Southern Poland. *Evolution* 40:1141–1159.

Tinkle DW. 1973. A population analysis of the sagebrush lizard, *Sceloporus graciosus* in Southern Utah. *Copeia* 1973(2):284–296.

Tinkle DW, Dunham AE, Congdon JD. 1993. Life history and demographic variation in the lizard *Sceloporus graciosus*: a long-term study. *Ecology* 74(8):2413–2429.

van Rossum G. 1995. Python tutorial. Technical Report CS-R9526, Centrum voor Wiskunde en Informatica (CWI), Amsterdam, the Netherlands.

Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, et al.; SciPy 1.0 Contributors. 2020. SciPy 1.0: fundamental algorithms for scientific computing in python. *Nat Methods*. 17(3):261–272.

Wang IJ. 2013. Examining the full effects of landscape heterogeneity on spatial genetic variation: a multiple matrix regression approach for quantifying geographic and ecological isolation. *Evolution* 67(12):3403–3411.

Wang IJ, Bradburd GS. 2014. Isolation by environment. *Mol Ecol*. 23 (23):5649–5662.

Wright S. 1931. Evolution in Mendelian populations. *Genetics* 16(2):97–159.

Wright S. 1943. Isolation by distance. *Genetics* 28(2):114–138.

Supplemental Material

Geonomics: forward-time, spatially explicit, and arbitrarily complex landscape genomic simulations

Drew E. Terasaki Hart, Anusha P. Bishop, and Ian J. Wang

Validation Tests

Full details and reproducible code and parameter files for each of our validation tests are available from the '/tests/validation/' subdirectory of the source code. We discuss the key details and results of these tests below.

Wright-Fisher test: genetic drift

The Wright-Fisher model of genetic drift models a fixed-size haploid population that turns over completely at each timestep (i.e. generation). The population can have any number of independent, biallelic genetic loci. For each locus, each generation's allele frequency is chosen as a binomial random variable, with the number of trials equal to the population size and the probability of success (i.e. of drawing the '1' allele) equal to the previous generation's '1'-allele frequency. The mean persistence time for an allele (i.e. the expected number of generations for which a locus remains segregating) is:

$$\bar{t}(p) = -4N[(1-p)ln(1-p) + p\ ln(p)] \qquad \text{(S1)},$$

where $2N$ is the number of alleles in the population (such that Ncan represent the diploid population size) and $p$is the frequency of an allele at the locus (Fisher, 1923; Hartl and Clark, 2007; Wright, Sewall, 1930).

The Wright-Fisher model is much simpler than the sorts of models for which Geonomics is designed (as are all of the following validation tests)—it is aspatial, panmictic, features fixed population sizes, and models only neutral loci. Thus, we parameterized Geonomics so as to approximate the model as closely as possible. To emulate aspatiality and panmixia, we used a population on a homogeneous landscape, using isotropic movement, with movement and dispersal distributions that broadly encompass the diagonal width of the landscape, and with no mating radius imposed (to allow panmixia) instead of local mating (i.e. with the mating radius set to 'None'). To enforce complete generational turnover, we set the maximum-age parameter to 1 (i.e. 1 timestep).

While Geonomics does not maintain constant population size, we maintained the carrying-capacity raster at a constant, uniform value, thus maintaining a stationary mean population size. We simulated 250 independent neutral loci (by setting all inter-locus recombination rates to 0.5), with starting ‘1’-allele frequencies of 0.5 (although the actual starting frequencies vary slightly around this value because of sampling error when all individuals’ genotypes are drawn binomially).

We ran the Wright-Fisher approximation test for three values of the carrying-capacity raster (i.e. three values of ‘K_factor’), hence for three mean population sizes (708, 1564, and 2440 individuals). For each mean population size (calculated as the harmonic mean, to account for stochastic fluctuations around the carrying capacity), we compared mean persistence time to that expected by theory, according to equation S1. Figures S1 and S2 show that the results are a close match to the theoretical expectations.

Bottleneck test: population dynamics

Because drift is a stronger evolutionary force in smaller populations, drift accelerates in shrinking populations. If a population undergoes a bottleneck event, the overall effect of drift on the population during that time is expected to be larger than what a constant-size population of equivalent starting size would experience during that time. Thus, mean fixation time should decrease in a bottlenecked population relative to one of constant size.

As with the Wright-Fisher model, to test the effectiveness of Geonomics for modeling a population bottleneck we used a homogeneous landscape with broad distributions for movement and dispersal and without a mating radius in order to emulate aspatiality and panmixia. To simulate a bottleneck event, we created a custom change event in which the population’s carrying-capacity raster is reduced to 30% of its initial value for 50 timesteps (from the 200th to 250th), then returned to its initial value for the remainder of the simulation (through the 300th timestep). These simulations produced a clear signal of accelerated drift during the bottleneck event, with the mean rate of allele-frequency change, calculated in 15-timestep sliding windows, nearly tripling during the period of the bottleneck (Figure S3).

Stepping stone test: population subdivision and genetic differentiation

The stepping-stone model, or one-dimensional island model, is a spatially implicit model. It models a series of subpopulations arranged along a straight line, with migration between all neighboring pairs. As a combined result of divergence by drift and homogenization by effective migration, subpopulations are expected to reach a stationary level of genetic differentiation: migration-drift equilibrium. Theory provides the expected pairwise genetic differentiation between a pair of subpopulations at equilibrium as:

$$F_{ST} = \frac{1}{1+4Nm} \qquad \text{(S2),}$$

where N is the population size and m is the per-generation migration rate, such that Nm can be interpreted as the per-generation number of migrant individuals (Hartl and Clark, 2007).

To approximate the stepping-stone model, we created a Landscape Layer with a diagonal of six equally spaced, equal-sized islands (1.0-valued cells) embedded in a ‘sea’ of 0.0-valued cells. We used this layer as the carrying-capacity raster (Figure S4, left). We parameterized dispersal to be very near to parents’ midpoints, movement distance to be strongly right-skewed, such that the long-distance movement events leading to migration are uncommon, and the mating radius to a value that makes island populations effectively panmictic but that prohibits mating between individuals on separate islands. Genomes contained 100 neutral loci, and we ran the simulation for 5000 timesteps.

Because Geonomics does not model discrete populations, it does not stipulate migration rates between discrete locations on the landscape. Thus, we manually tracked the number of migration events during each timestep for all possible directional migration events (i.e. for all permutations of island pairs), then used that data to calculate all mean migration rates. With those values, we solved equation S2, then compared the resulting F_{ST} expectations to the observed values (calculated from the simulated data using two common methods; Figure S4, right). We also used Discriminant Analysis of Principal Components (DAPC), performed in the R package *adegenet* (Jombart et al. 2008), to visualize population structure.

The results demonstrate that the model approached migration-drift equilibrium, as expected by theory (Figure S5), with all island populations reaching dynamic equilibria around the same mean size. Estimated migration rates and F_{ST} values qualitatively match theoretical expectations: mean migration rate drops off precipitously at greater than one step-distance apart, and genetic differentiation increases to approximate saturation. Values of F_{ST} consistently undershoot the values expected based on estimated migration rates, however, because subpopulations have yet to approach fixation at most loci (which is the expectation implied by expected F_{ST} values close to 1). DAPC demonstrated that the simulation generated the expected population structure of six distinct clusters, one per island (Figure S6).

Contrasting-habitat test: divergent selection

In a population divided between two, divergent selective environments, if there is standing genetic variation for a biallelic locus controlling the trait adapting to those environments then theory predicts that the two subpopulations will diverge at that locus as each moves toward its respective adaptive peak. The rate at which divergence should occur depends on the relative strengths of two opposing evolutionary forces: natural selection, which causes divergence, and gene flow, which causes homogenization. The rate of allele frequency change in either subpopulation at timestep t is expressed as:

$$\delta q = \frac{-spq[q+h(p-q)]}{1-sq(2hp+q)} + m_i q^* - m_o q \qquad \text{(S3),}$$

where $p$and $q$are the frequencies of the beneficial and deleterious alleles in the local subpopulation, $s$is the selection coefficient against the homozygous recessive phenotype, $h$is the degree of dominance of the recessive allele, m_i and m_o are the migration rates into and out of the subpopulation being analyzed, and $q^*$is the frequency of the locally deleterious allele in the alternative subpopulation where it is beneficial (Hartl and Clark, 2007).

This model, like the stepping-stone model, is spatially implicit. To approximate this, we created a landscape with two layers. The first was divided into two equal-sized halves, one valued at 0.0, the other at 1.0; this layer was used as the layer driving natural selection. The second was valued uniformly at 1.0; this was used as the carrying-capacity raster (thus setting uniform population density across the landscape and determining, in sum, the overall carrying capacity of the landscape). We created one monogenic trait whose position was randomly chosen within a genomic architecture of 100 unlinked loci. We ran the model for 1000 timesteps for each of three values of the parameter ϕ (identical to s in equation S3): 0.1, 0.05, and 0.01. Given that Geonomics does not directly define a migration rate parameter, we tracked the number of migration events (i.e. individuals crossing the landscape's horizontal midline) during each timestep, then used that data to solve equation S3.

Results depict clear local adaptation to each of the two halves of the landscape, with spillover of opposite phenotypes and resulting heterozygote births occurring along the border between the two habitats (Figure S7). Allele trajectories in each half of the environment follow qualitatively the increasing and saturating trajectories expected by theory, but reach consistently more divergent allele frequencies than expected based on the theoretical calculation (Figure S8). However, these results are an easily understandable artefact of estimating a spatially implicit, population-based model using a spatially explicit, individual-based one—our method of calculating migration rates includes all individuals who cross the habitat boundary, including the large number who only barely cross and who may even quickly migrate back, such that expected allele-frequency trajectories are based on an overestimation of true gene flow and thus serve as lower bounds on the real trajectories. As further validation, a plot of the mean difference between each individual's phenotypic and environmental values shows a strong decline over model time, with the rate and level of decline increasing as a function of increasing strength of selection (Figure S9). Moreover, logistic regressions show no significant relationships between phenotypic and environmental values at the outset (pseudo-R^2s ≈ 0.0, p-values > 0.1) but show highly significant relationships at the ends of the simulations ($p < 0.0001$ for all values of ϕ), with the amounts of variation explained increasing as a function of selection strength (pseudo-$R^2 = 0.327$ for $\phi = 0.01$, 0.376 for $\phi = 0.05$, and 0.406 for $\phi = 0.1$).

Cline test: local adaptation

In a clinal model, a population adapts locally across an environmental gradient, which is characterized by the extremes of its environmental values and its steepness (i.e. the instantaneous rate of environmental change along it). Local adaptation across this gradient will generate a geographic cline in allele frequencies. The clinal pattern is only expected for loci under selection along the cline (and other loci in linkage). Unlinked loci have no long-term clinal expectation (though they could initially be swept along with the selective locus, and any number could continue to show spurious concordant clinal variation). To detect clinal adaptation, we can fit cline curves to the allele-frequency variation across the environmental gradient for all loci, with the expectation that the clines fit to adaptive loci will mirror the gradient. Numerous equations have been used to fit clines, but one of the most common is the sigmoidal *tanh* function:

$$p_x = \frac{1}{2}(1 + t[\frac{2(x-c)}{w}]) \qquad \text{(S4),}$$

where p is the frequency of the reference allele at position x along the cline, c is the centerpoint of the cline (such that $p_{x=c} = 0.5$), and w is the 'width', which is defined as $w = \frac{1}{slope}$ at point c (Porter, 2013).

To implement the cline model in Geonomics, we created a landscape with two layers. The first layer was an environmental layer—a symmetrical, non-linear gradient between 0-valued and 1-valued halves (Figure S10). The second was a uniformly valued habitat-quality layer, used to set a uniform population density and thus determine the global carrying capacity. We created a monogenic trait whose locus was randomly placed within a genomic architecture of 100 independent loci. The trait had a ϕ of 0.01, with the gradient layer serving as its selective force. We ran the cline model for 1500 timesteps, then used a numerical optimization function (in Python's *scipy* package; Jones *et al.*, 2001) to fit equation S4 for all loci. We plotted all fitted clines on top of the first landscape layer, with the cline for the single selective locus highlighted. The selective locus consistently and clearly stands out as the only locus with a cline matching the expectation of a monotonic pattern mirroring the environmental gradient and spanning nearly the full range of phenotypic values (Figure S11).

Results clearly show a pattern of clinal adaptation across the landscape—despite isolated patches of maladaptive genotypes potentially resulting from occasional long-distance migration events—with a zone of admixture and phenotypic spillover surrounding the cline's center (Figure S10). In a Bonferroni-corrected family of locus-wise logistic regression models of environmental value on genotype, the selective locus consistently stands out as the most significant (p-values of roughly $3x10^{-100}$). Furthermore, a plot of the mean difference between phenotypic and environmental values shows a strong decline over model time (Figure S12), and logistic regressions show no significant relationship between phenotypic and environmental values at the outset (pseudo-$R^2 = 0$, p-value = 0.370) but a significant relationship at the end of the simulation (pseudo-$R^2 = 0.345$, p-value < 0.0001).

Selective sweep test: genetic hitchhiking

Genomic architecture and linkage add important complexity to models of molecular evolution. The most basic model of selection with linkage is that of a selective sweep: a beneficial mutation occurs in a population, falling on a random genomic background, then rises in frequency because of its selective advantage until it becomes fixed, pulling up the frequency of the surrounding haplotype block in the process. The haplotype block is, nevertheless, subject to recombination, which gradually erodes it symmetrically around the beneficial mutation. Thus, the selective-sweep model predicts that once a beneficial mutation occurs —as long as it is not lost early on by chance— it and the haplotype block around it will rise in frequency, the mutation will eventually fix, potentially with some core block around it, and the rest of the block will erode over time. The haplotype block should be clearly visible in genomic data, where it will manifest as a genomic region of reduced diversity and heterozygosity centered on the mutation.

To implement the selective sweep model in Geonomics, we again created a model approximating an aspatial, panmictic population (see Wright-Fisher test for details). We created a single, monogenic trait with a ϕ of 0.1. The trait's locus was manually set to position 500, such that it was at the center of the 1001-locus genome. The genome had a homogeneous recombination rate of 0.001 between all neighboring loci. We manually set the starting '1'-allele frequency at this locus to 0.0 but set the trait to be selected upon by a uniform layer of 1 values, such that all individuals began the model equally unfit (i.e. with a fitness value of $1 - \phi = 0.9$). After burn-in, we iteratively chose a random individual, introduced a '1'-mutation in its genome at locus 50, ran the model for 50 timesteps, and checked whether the '1' allele had reached a frequency greater than 0.05 by that time. We iterated until that check was passed, at which point we declared the mutant allele 'established' and continued to run the model until 2500 timesteps after the novel mutation reached fixation. At three timepoints during that model we calculated and recorded genome-wide nucleotide diversity using a sliding-window approach.

We found that Geonomics successfully and realistically simulated the behavior of a selective sweep. The first adaptive mutant that was not immediately lost by drift rose rapidly in frequency, then fixed. The population's mean fitness increased quickly from 0.9 (the universal fitness value before the mutation was introduced) to 1.00 (the universal fitness value after the sweep was complete; Figure S14). The linkage block around the selected locus became a region of depressed nucleotide diversity (Figure S13, top row) and heightened linkage (Figure S13, bottom row) — the classic signature of a selective sweep.

Recombination test

To provide additional validation of Geonomics' recombination model, we compared the effective recombination rates observed in a Geonomics model to those produced by an msprime simulation using the same recombination map. We produced a recombination map by assigning 999 random, interlocus recombination rates to a 1000-length simulated genome. We drew the rates by taking the first 999 values ≤ 0.5 from a random vector drawn from the distribution ~Beta(0.4, 1.3), producing a left-skewed distribution that nonetheless sampled the full range of physical linkage values. We ran an msprime model using the msprime.RecombinationMap object created from those values and also ran a Geonomics model using those values as the recombination-rate column in a Geonomics custom genomic architecture file. We then plotted the true recombination rates and the observed breakpoint densities from both models, binned within even-width genomic windows. The results show that Geonomics' observed breakpoint densities recapitulate the true recombination rates just as closely as do those of msprime (Figure S15).

Applications

Example 3: Polygenic adaptation to climate change in the Yosemite region

To build this simulation, we first generated a template Geonomics parameters file (using the Geonomics function ‘gnx.make_parameters_file(…)’), then edited it to best emulate our empirical study system (see Code Sample S1). This created a parameters file for a simulation with: (1) three empirical layers mean temperature, precipitation, and habitat suitability), two of which (temperature and habitat suitability) have environmental change events; (2) one species, with one trait adapted to mean temperature; and (3) a data collection design. We set life-history parameters to reasonable approximations of *S. graciosus* biology, based on available literature. We set spatial parameters based on the relationship between the resolution of the environmental rasters and the characteristic scales of *S. graciosus* key life-history traits. The resolution of the rasters is 0.00833°, which at latitude 38° is equal to about 730.984 m in the east-west direction and 927.296 m in the north-south direction, giving each cell a total area of roughly 6.78×10^5 m^2 (67.8 hectares). Using a population density of 208 individuals per hectare (Tinkle, 1973) and the rough estimate that about 10% of the land area covered by our study contains the open habitat favored by *S. graciosus* rather than the more closed habitat favored by the congener *S. occidentalis*, with whom it experiences a large degree of competitive exclusion (Rose 1976), we chose a per-cell carrying capacity (parameter ‘K_factor’) of 67.8 hectares/cell × 208 individuals/hectare × 0.1 proportion of habitat suitable $\approx$ 1410 individuals (which we then further multiplied by 0.1 for computational tractability). We set the reproductive age to 2 years (Tinkle, 1973; Tinkle *et al.*, 1993). We left the sex ratio at unity, given controversy in the literature about whether or not it skewed toward females because of lower male survival rates (Tinkle, 1973; but see Tinkle *et al.*, 1993). We set the number of births per individual to be a Poisson random variable with lambda = 4.464 individuals/clutch × 2 clutches/year × 0.16 survival rate = 1.428, based on an average clutch size of 4.464 across surveyed California populations (Tinkle *et al.* 1993), an average of 2 clutches per year (Tinkle, 1973; Tinkle *et al.*, 1993), and an average rate of survival to the first year of 0.16 (Ruth, 1978). We estimated the mean interannual movement distance as 12.457 m (expressed as 0.01704 cell widths), based on an average of all recorded interannual movement events in Stebbins’ (1948) study of *S. graciosus* home ranges. We used this as an order-of-magnitude estimate for movement, but increased parameters slightly above this value in order to pair the reduced population density we chose for purposes of computational tractability. Thus, we set a mating radius value of 0.5 cell widths, and parameterized movement and dispersal as '~Lognormal($7x10^{-5}$, 0.3)' and '~Lognormal($7.5x10^{-4}$, 1),' respectively. In the absence of any known published estimates, we set the population intrinsic growth rate to 0.5. The full code to perform this analysis is available as Code Sample S2 and in the Yosemite demo script, included in the Geonomics package.

Accessibility

For both of the two most common types of color blindness (protanopia and deuteranopia) we tested the colorblind-friendliness of all of the Matplotlib color palettes used by Geonomics as defaults. For perceptive simulation we used the script provided by Sarjak Thakkar (2018). All color palettes retain interpretability.

Figure Legends

Figure S1: Trajectories for the frequencies of the ‘1’-alleles for 25 of the 250 simulated loci (one line per locus) in a Wright-Fisher model without mutation. We ran simulations for three mean population sizes, as determined by three fixed values of the carrying capacity (‘K_factor’) parameter, until all loci fixed.

Figure S2: Violin plots of mean persistence time distributions across all loci from our Wright-Fisher validation test, shown as a function of harmonic mean population size. Resulting mean persistence times (red dots) are an extremely close match to predictions calculated using Equation S1 (black, horizontal lines).

Figure S3: Ten randomly chosen allele frequency trajectories (top), population size (middle), and mean rate of allele frequency change (bottom; calculated for 15-timestep sliding windows) from the bottleneck validation test. We ran the simulation for 300 timesteps with a 50-timestep bottleneck.

Figure S4: Map of six island populations at the end of the simulation for the stepping-stone validation test (left), produced using ‘model.plot’ in Geonomics, and plot of pairwise F_{ST} values and inter-island migration rates as functions of inter-island distance (right). R^2 values and p-values result from quadratic regressions of F_{ST} values on inter-island distances and log-log regression of mean migration rates on inter-island distances.

Figure S5: Plot of F_{ST} over model time for the stepping-stone validation test. Each line represents a different island pair, with colors corresponding to increasing inter-island distances (from yellow to green).

Figure S6: Results of discriminant analysis on principal components (DAPC) from the stepping-stone validation test, including plots of (A) the individual loadings on the first three discriminant axes, (B) individuals at the final time step color coded by population membership assignments, and (C) DAPC membership probabilities for each of the individuals. The optimal number of PCs to retain (n = 59) was determined through cross-validation using the ‘xvalDapc’ function in the adegenet R package (Jombart et al. 2008).

Figure S7: Map of the population after spatially divergent selection at □ = 0.10 in simulations for the divergence validation test, produced using the ‘model.plot_fitness’ function in Geonomics. Individuals are plotted on top of the selective landscape layer, which is divided into two halves. Outer circles are colored by phenotype, ranging from dark blue to dark red, representing the optimal phenotypes for each environmental background. Inner circles are color and sized by fitness, such that darker-gray, larger inner circles represent less fit individuals. Stochasticity leads to asymmetry in the structure of the hybrid zone, the nature of which varies from one iteration to the next; at the moment when this figure was produced, more blue alleles were present in the red environment than vice versa.

Figure S8: Observed (solid lines) versus expected (dashed lines) allele-frequency trajectories for two contrasting habitats (blue = 0.0-valued; red = 1.0-valued) resulting from divergence test simulations with three selection coefficients:□ = 0.01 (dark), □ =0.05 (medium), and□ = 0.10 (light).

Figure S9: Plot of the mean difference between each individual’s phenotype and environmental value plotted against time, for divergence test simulations with three different selection coefficients (‘phi’). A pattern of background matching, which is indicative of local adaptation, builds up over time. The pattern develops more quickly, and becomes more pronounced, under stronger selection regimes.

Figure S10: Map of the final generation from the cline test simulation on top of the selective landscape layer, with individuals colored by phenotype (outer circles) and fitness (inner circles), as in Figure S7.

Figure S11: Plot of allele-frequency clines (neutral loci in black, selective locus in bold yellow) against the selective landscape layer (horizontal gradient from red to blue) from the cline test.

Figure S12: Plot of the mean difference between each individual’s phenotype and environmental value plotted against time during the cline test simulation. A pattern of background matching builds up over time.

Figure S13: Results of the selective sweep validation test, including nucleotide diversity calculated in 11-locus windows across the genome (top row) and pairwise linkage (R^2) for locus pairs plotted against genetic distance (bottom row). Genetic distance was calculated as the ‘recombination distance’ (the sum of intervening interlocus recombination rates between paired loci).

Figure S14: Mean fitness of the entire population, over the full run of the selective sweep test simulation.

Figure S15: Breakpoint densities, within 50, even-width genomic windows, as calculated from the tskit.TreeSequence results of a Geonomics simulation (red) and an equivalent msprime simulation (blue). The observed densities clearly recapitulate the true recombination rates expressed in the input recombination map (black dashed line).

Figure S16: Results from discriminant analysis of principal components (DAPC) of neutral-locus genotypes at the final time step from the simulations for the isolation by distance (IBD) and by environment (IBE) example, conducted using the R package *adegenet* (Jombart et al. 2008). The optimal number of PCs to retain (n = 7) was determined through cross-validation using the ‘xvalDapc’ function (Jombart et al. 2008). A plot of the individuals’ loadings on the first three discriminant axes (A), with each individual colored according to its DAPC-derived population membership assignment (C), recapitulates their spatial arrangement (B). Beyond showing general patterns of IBD, our simulated neutral genetic data clearly match the expected hierarchical population structure: distinct clusters separated by the central barrier, with the subclusters further differentiated along the environmental gradients running in opposite directions on either side.

Figure S17: The mean difference between individuals’ phenotypes and environmental values plotted against time, resulting from the simulations for the simultaneous selection example application. Values decrease over time for both traits, reflecting the buildup of a pattern of background matching.

Figure S18: Results of the simultaneous selection simulation when selection is excluded ($\phi = 0$). Individuals are colored by phenotype for the trait corresponding to each layer and show no signal of background matching, as expected.

Figure S19: The mean difference between individuals' phenotypes and environmental values plotted against time, during the simultaneous selection simulation, when selection is excluded ($\phi = 0$). These values show no decreasing trend over time, unlike in the model including selection (Figure S16).

Code Samples

Code Sample S1:

```
>>> gnx.make_parameters_file(filepath=‘yosemite_params.py’,
                             layers=[{‘type’: ‘file’, ‘change’: True},
                                     {‘type’: ‘file’, ‘change’: True},
                                               {‘type’: ‘file’, ‘change’: False}],
                           species=[{‘movement’: True,
                                                 ‘movement_surface’: True,
                                        ‘genomes’: True, ‘n_traits’: 1}],
                            data=True)
```

Code Sample 2:

```
>>> model = gnx.make_model(filepath=‘yosemite_params.py’)
>>> model.walk(100000, mode = ‘burn’)
>>> model.walk(500, mode = ‘main’)
```

References

Fisher, R.A. (1923). XXI.—On the Dominance Ratio. Proc. R. Soc. Edinb. *42*, 321–341.

Hartl, D.L., and Clark, A.G. (2007). Principles of Population Genetics: Fourth Edition (Sunderland, Massachusetts: Sinauer Associates, Inc. Publishers).

Jones E, Oliphant E, Peterson P, *et al.* (2001). SciPy: Open Source Scientific Tools for Python, http://www.scipy.org/.

Porter, A. (2013). ClineFit v. 2.0, User's Manual.

Rose, Barbara R. 1976. Habitat and prey selection of *Sceloporus occidentalis* and *Sceloporus graciosus*. Ecology. 57, 3:531-541.

Ruth, SB. 1978. A Comparison of the demography and female reproduction in sympatric western fence lizards (*Sceloporus occidentalis*) and sagebrush lizards (*Sceloporus graciosus*) on Mount Diablo, California. Ph.D dissertation: University of California Berkeley. URL: https://elibrary.ru/item.asp?id=7216038.

Stebbins, Robert C. 1948. Additional observations on home ranges and longevity in the lizard *Sceloporus graciosus*. Copeia. 1: 20-22.

Tinkle, Donald W. 1973. A population analysis of the sagebrush lizard, *Sceloporus graciosus* in Southern Utah. Copeia. 2:284-296.

Tinkle DW, Dunham AE, Congdon JD. 1993. Life history and demographic variation in the lizard *Sceloporus graciosus*: a long-term study. Ecology. 74, 8:2413-2429.

Thakkar S. 2018. Simulate-Correct-ColorBlindness. GitHub Repository: https://github.com/tsarjak/Simulate-Correct-ColorBlindness. Commit: 95edbbbecaa75e4869e0b7c12126a6965445c93c.

Wright, Sewall (1930). Evolution in Mendelian populations. Genetics *16*, 97–159.

Figures

Figure S1:

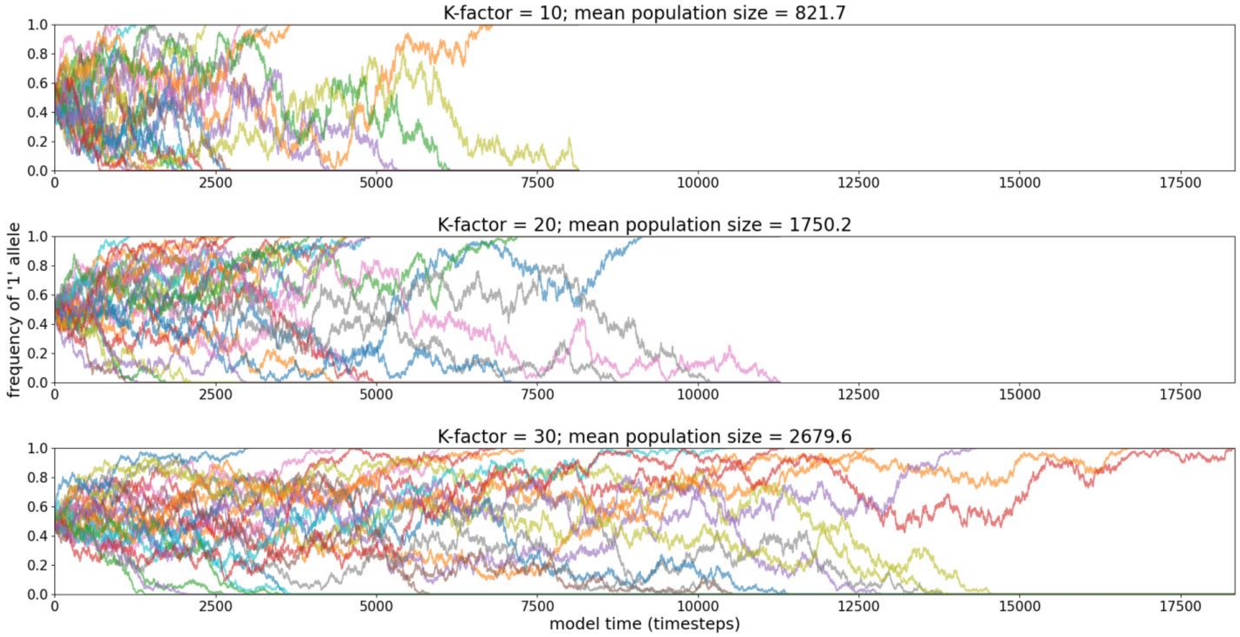

Figure S2:

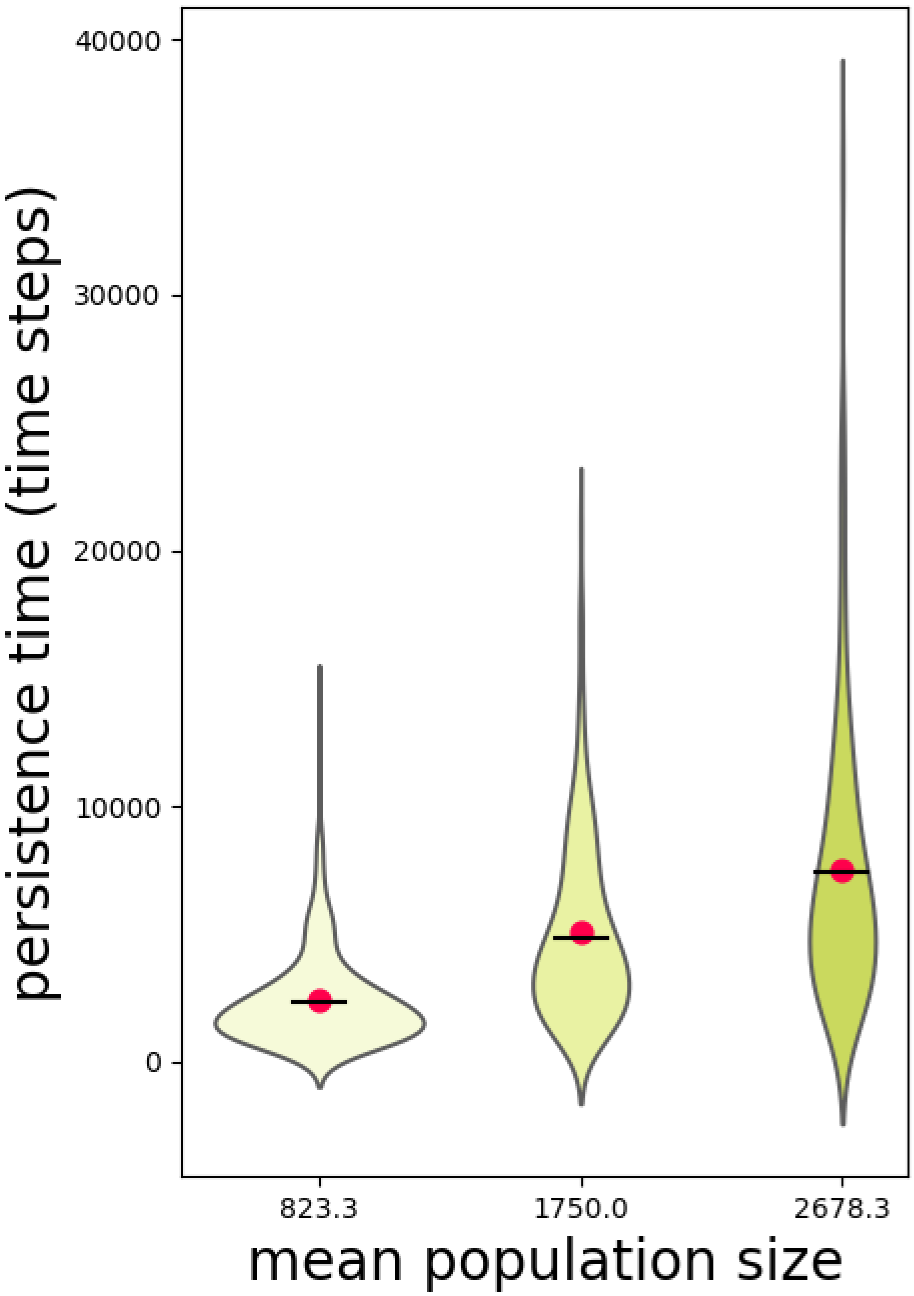

persistence time (time steps)
40000
30000
20000
10000
0
823.3
1750.0
2678.3
mean population size

Figure S3:

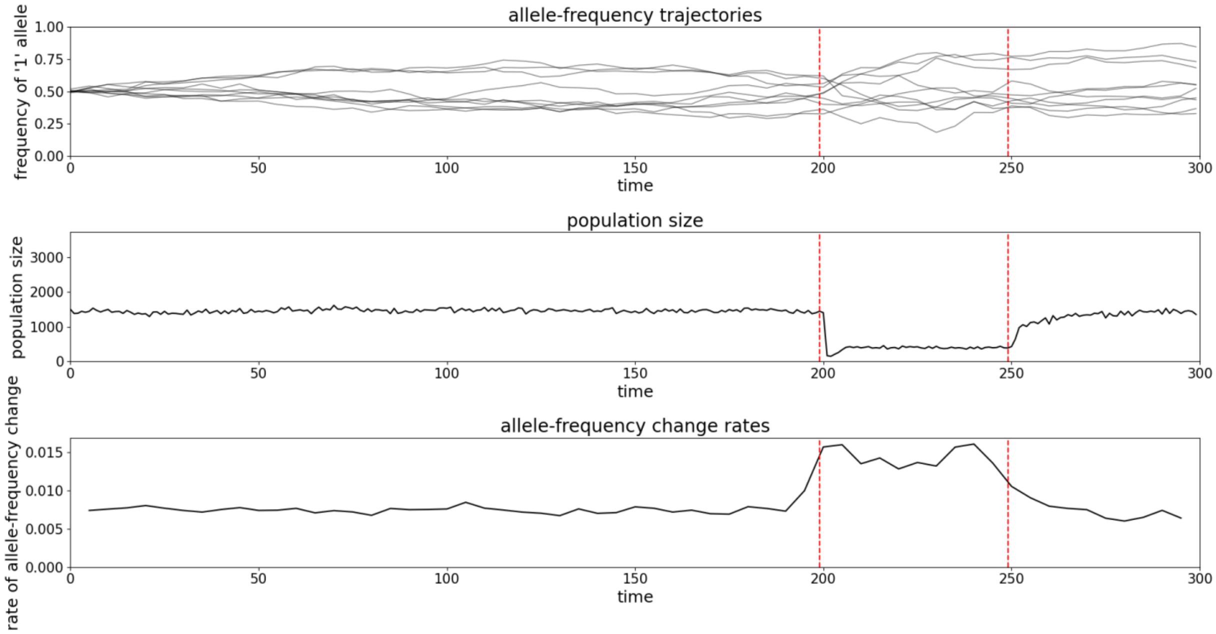
allele-frequency trajectories
frequency of '1' allele
1.00
0.75
0.50
0.25
0.00
0
50
100
150
200
250
300
time
population size
population size
3000
2000
1000
0
0
50
100
150
200
250
300
time
allele-frequency change rates
rate of allele-frequency change
0.015
0.010
0.005
0.000
0
50
100
150
200
250
300
time

Figure S4:

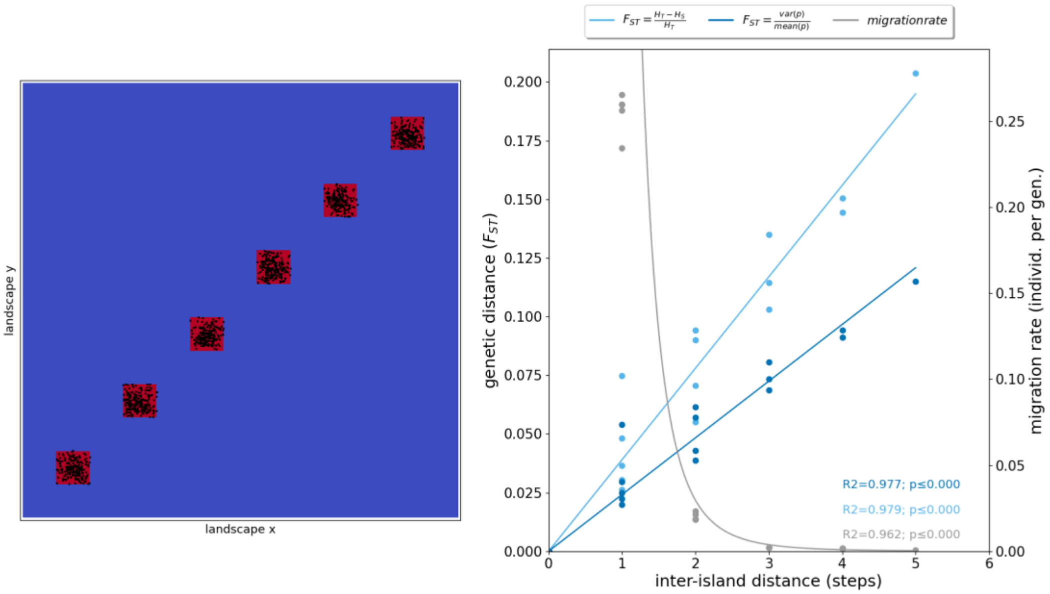
landscape y
landscape x
$F_{ST} = \frac{H_T - H_S}{H_T}$
$F_{ST} = \frac{var(p)}{mean(p)}$
migrationrate
genetic distance (F_{ST})
0.000
0.025
0.050
0.075
0.100
0.125
0.150
0.175
0.200
migration rate (individ. per gen.)
0.00
0.05
0.10
0.15
0.20
0.25
inter-island distance (steps)
0
1
2
3
4
5
6
R2=0.977; p≤0.000
R2=0.979; p≤0.000
R2=0.962; p≤0.000

Figure S5:

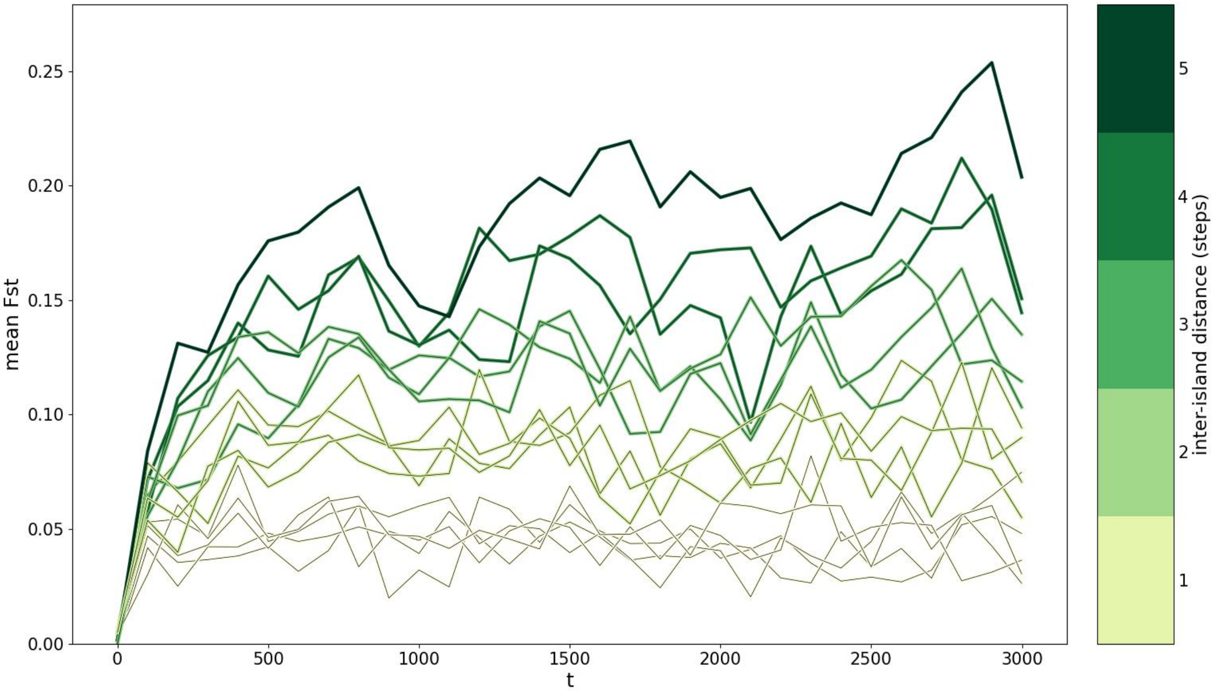
mean Fst
0.00
0.05
0.10
0.15
0.20
0.25
t
0
500
1000
1500
2000
2500
3000
inter-island distance (steps)
1
2
3
4
5

Figure S6:

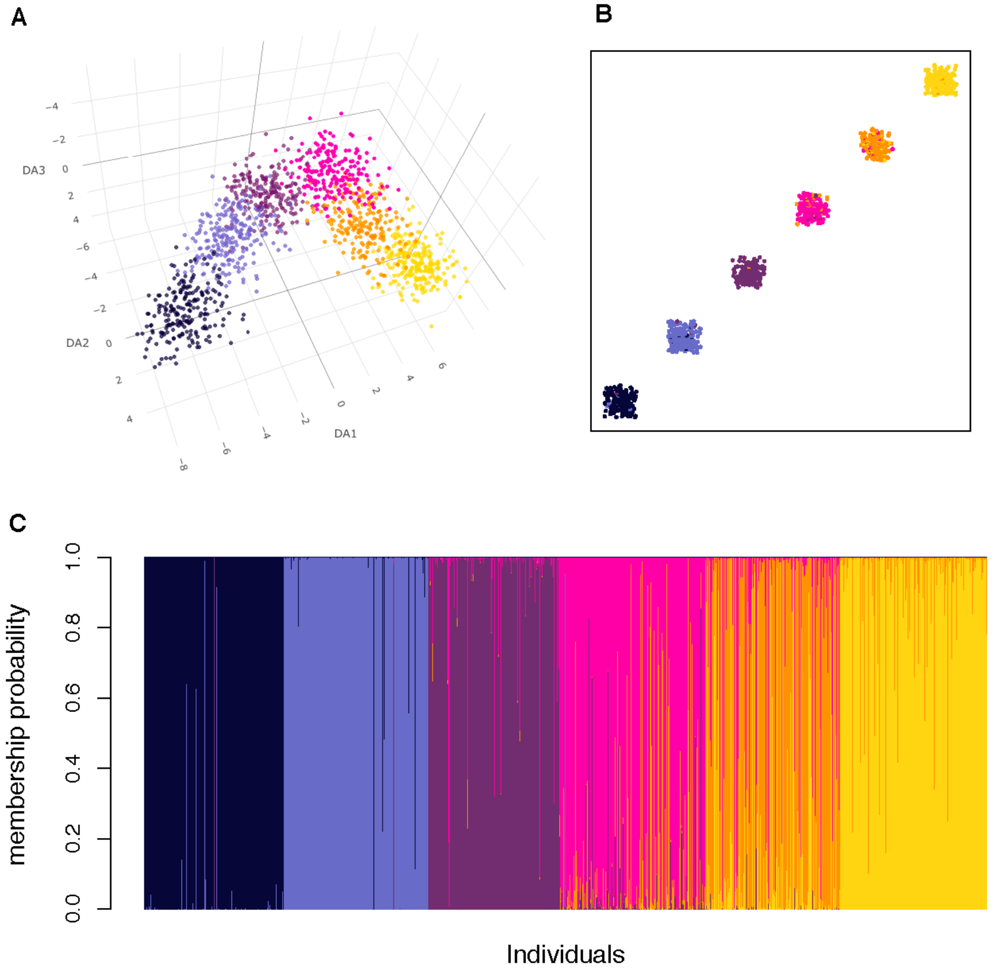
A
B
C
DA3
−4
−2
0
2
4
DA2
−6
−4
−2
0
2
4
DA1
−8
−6
−4
−2
0
2
4
6
membership probability
0.0
0.2
0.4
0.6
0.8
1.0
Individuals

Figure S7:

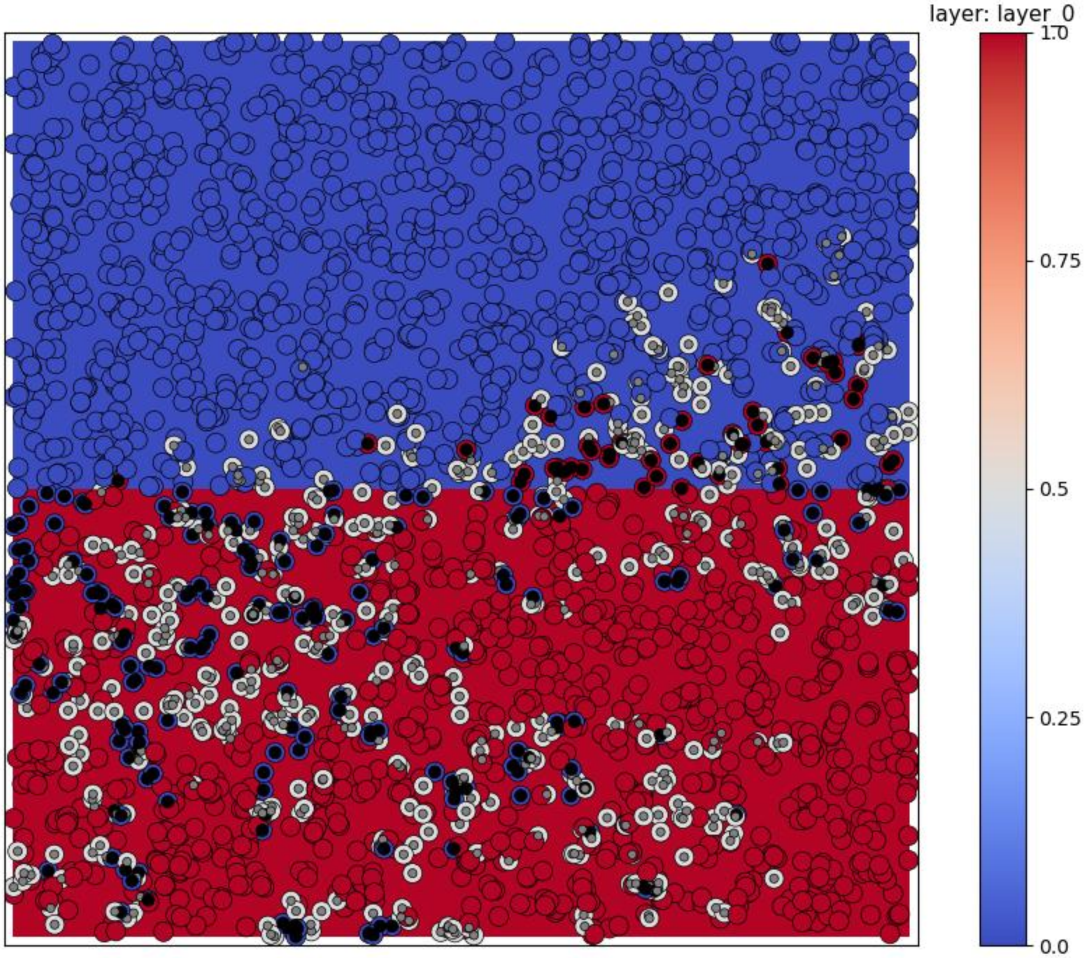

Figure S8:

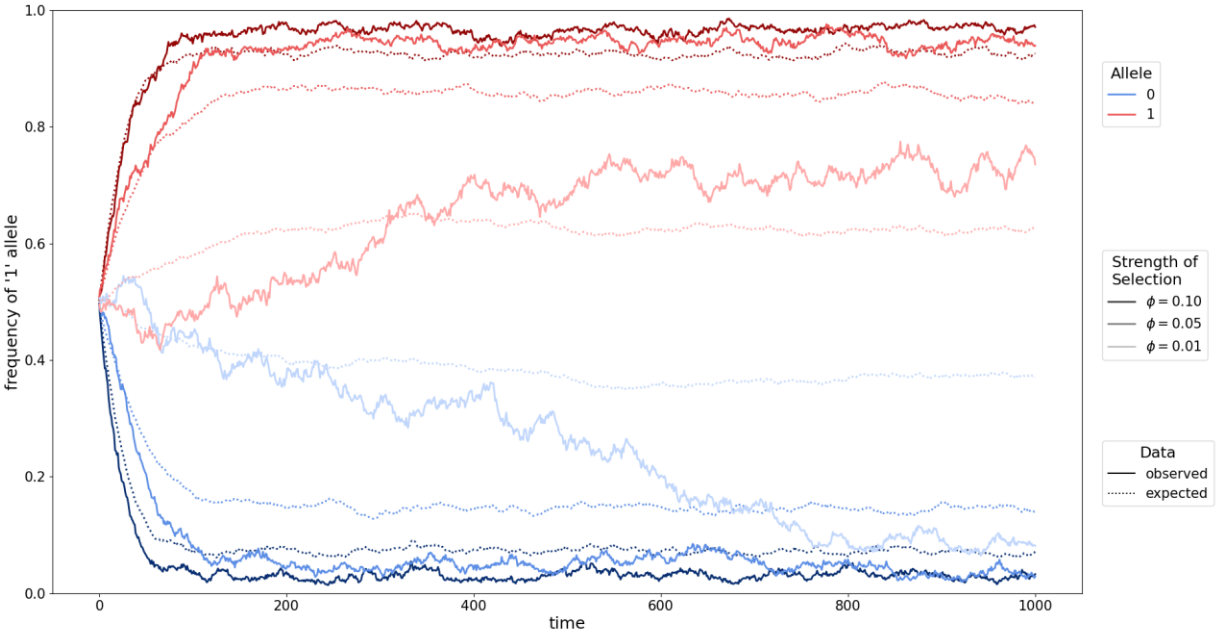

1.0
0.8
0.6
0.4
0.2
0.0
frequency of '1' allele
0
200
400
600
800
1000
time
Allele
0
1
Strength of Selection
$\phi = 0.10$
$\phi = 0.05$
$\phi = 0.01$
Data
observed
expected

Figure S9:

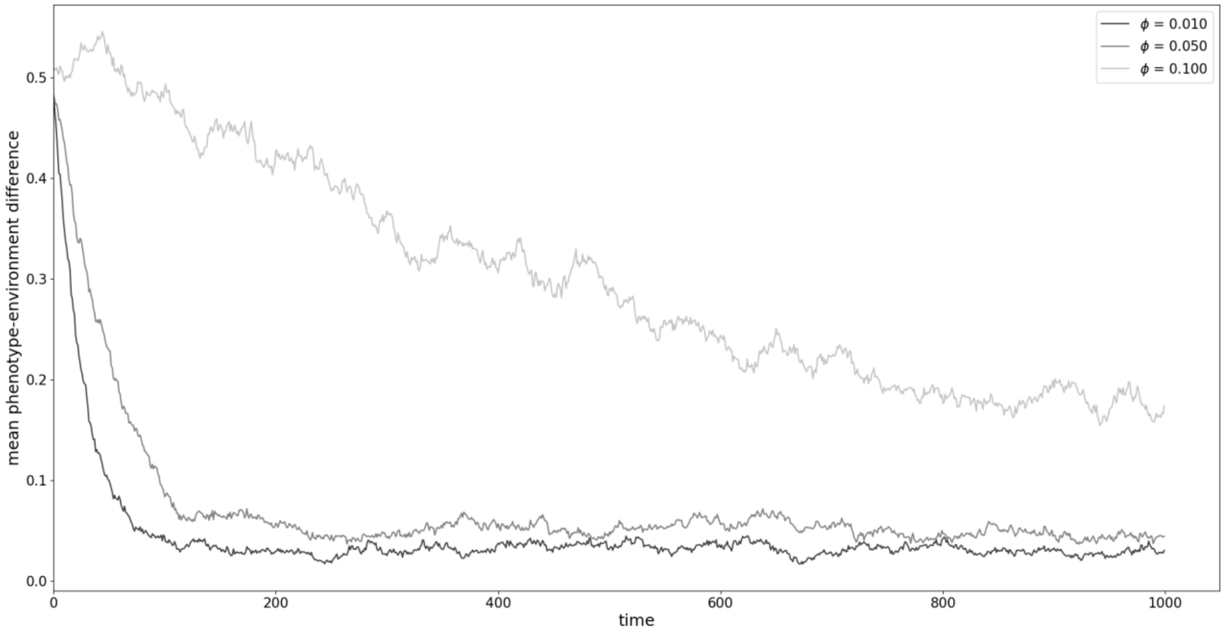
ϕ = 0.010
ϕ = 0.050
ϕ = 0.100
mean phenotype-environment difference
0.0
0.1
0.2
0.3
0.4
0.5
time
0
200
400
600
800
1000

Figure S10:

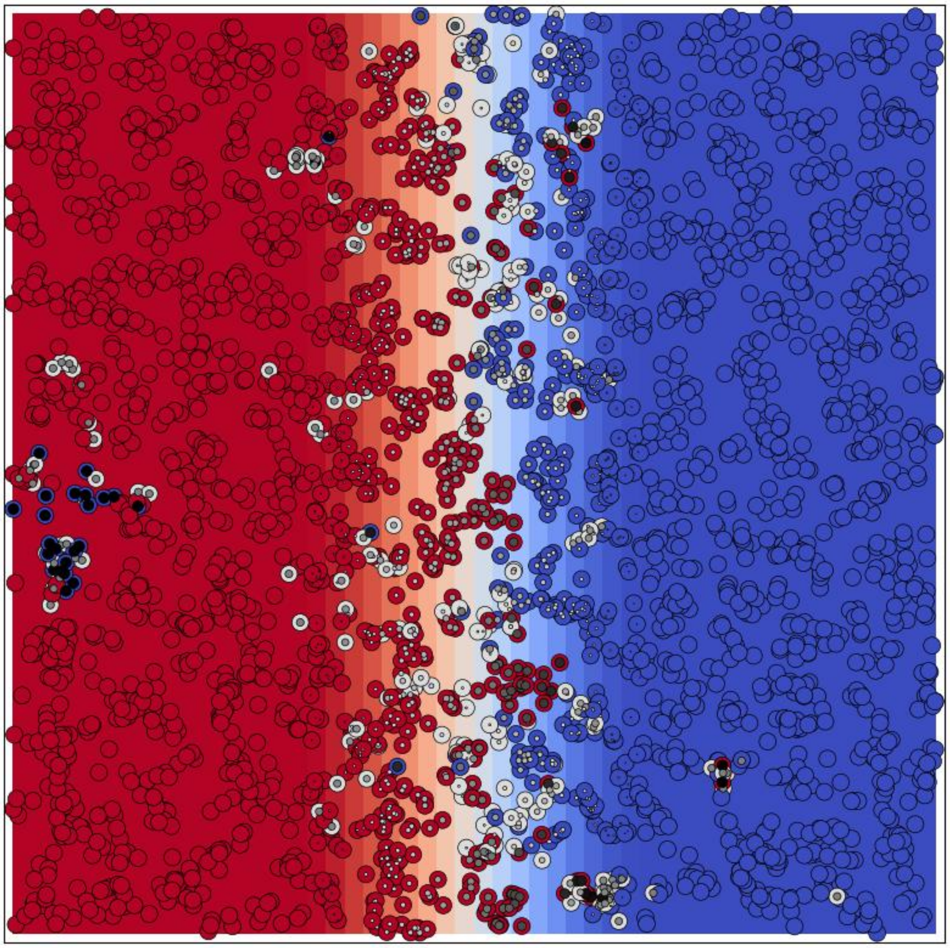

Figure S11:

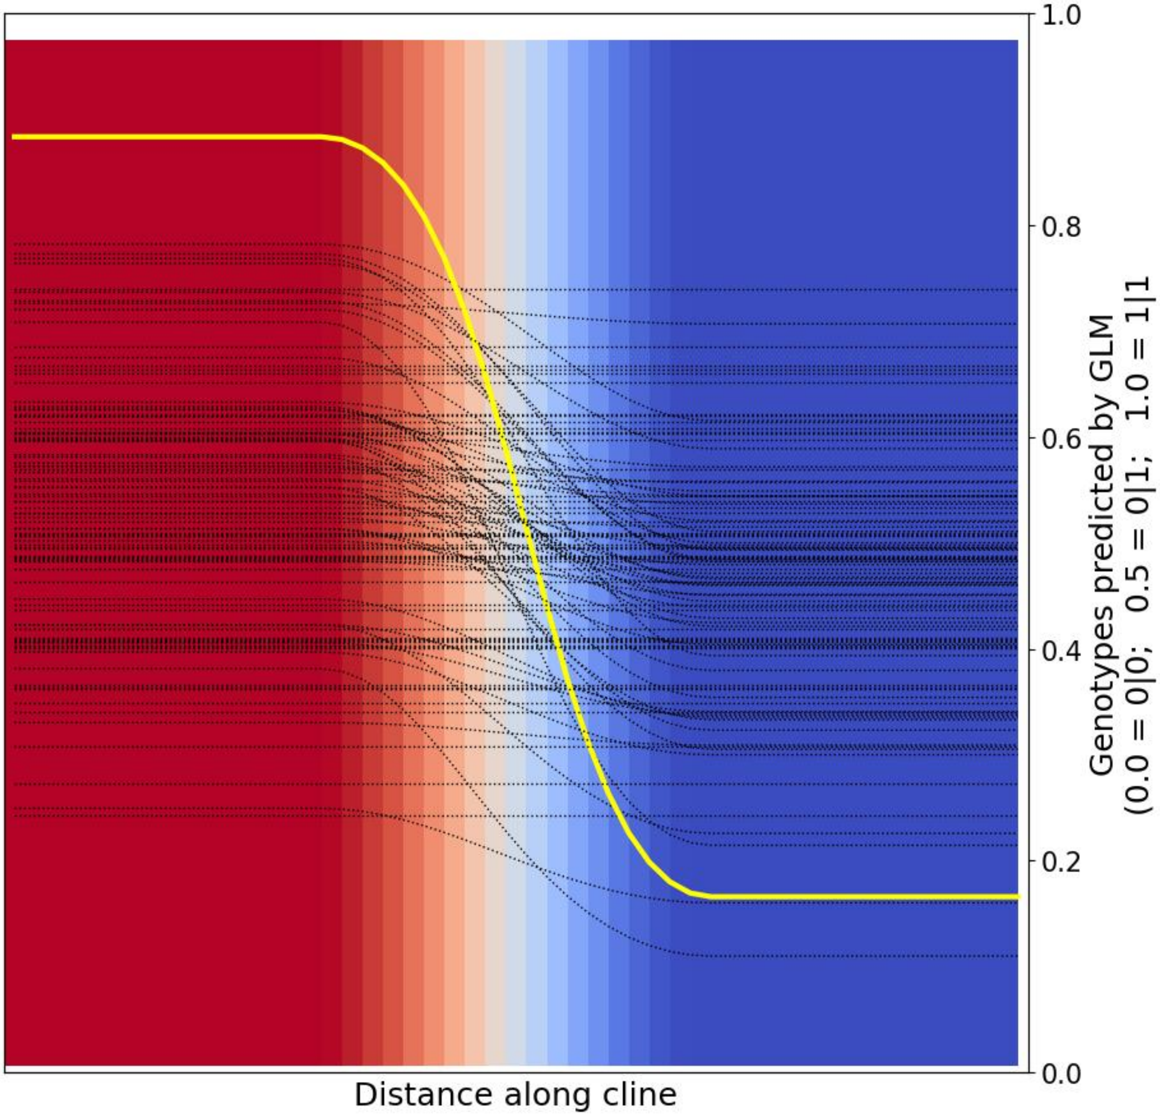
1.0
0.8
0.6
0.4
0.2
0.0
Genotypes predicted by GLM
(0.0 = 0|0; 0.5 = 0|1; 1.0 = 1|1
Distance along cline

Figure S12:

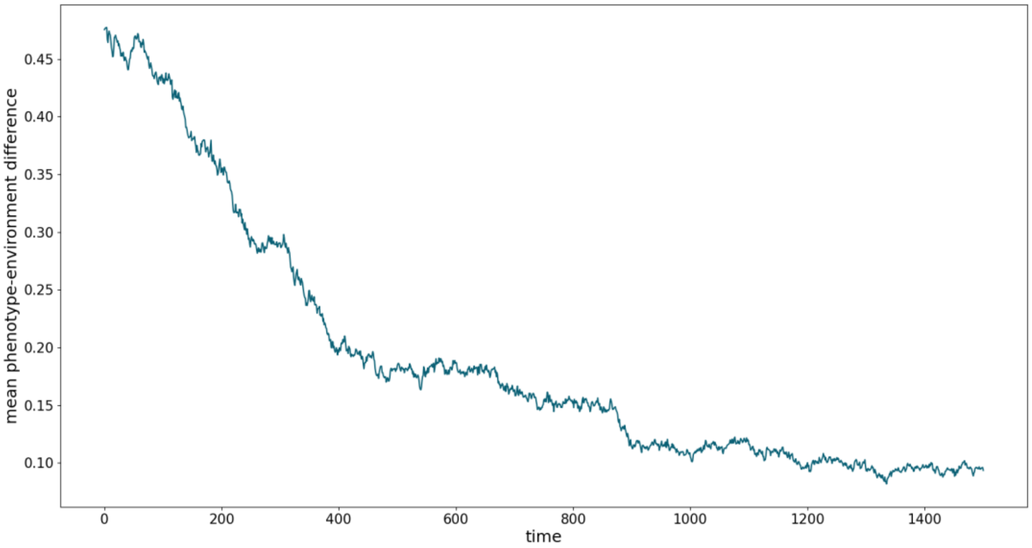
mean phenotype-environment difference
0.45
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0
200
400
600
800
1000
1200
1400
time

Figure S13:

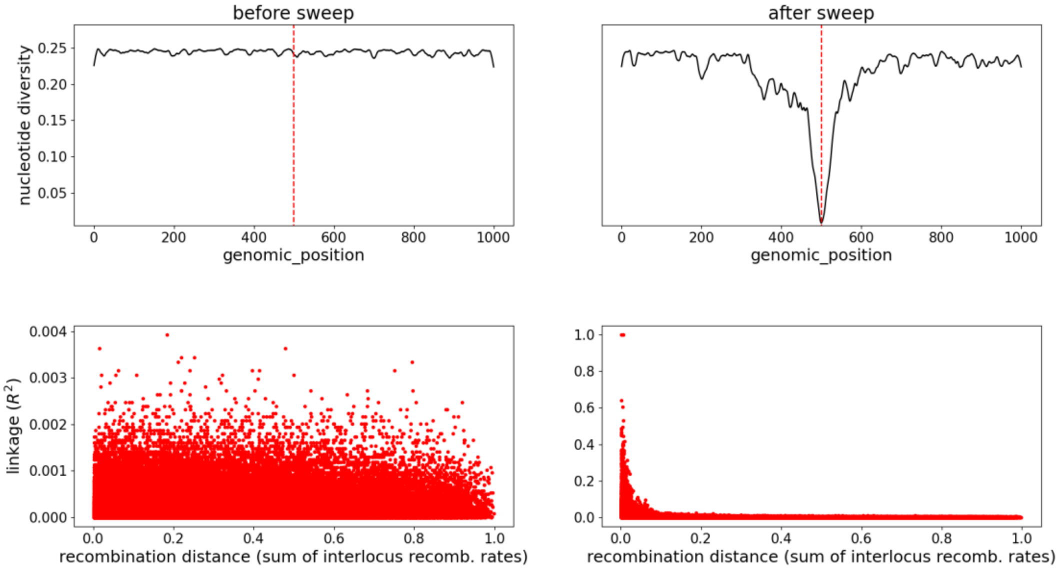
before sweep
after sweep
nucleotide diversity
0.25
0.20
0.15
0.10
0.05
0
200
400
600
800
1000
genomic_position
genomic_position
linkage (R^2)
0.004
0.003
0.002
0.001
0.000
1.0
0.8
0.6
0.4
0.2
0.0
recombination distance (sum of interlocus recomb. rates)
recombination distance (sum of interlocus recomb. rates)

Figure S14:

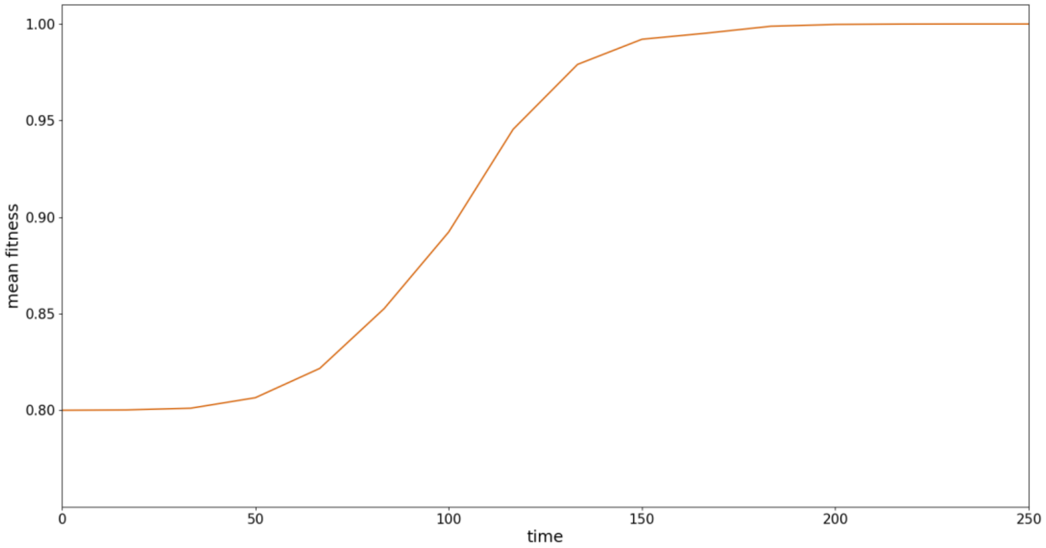

Figure S15:

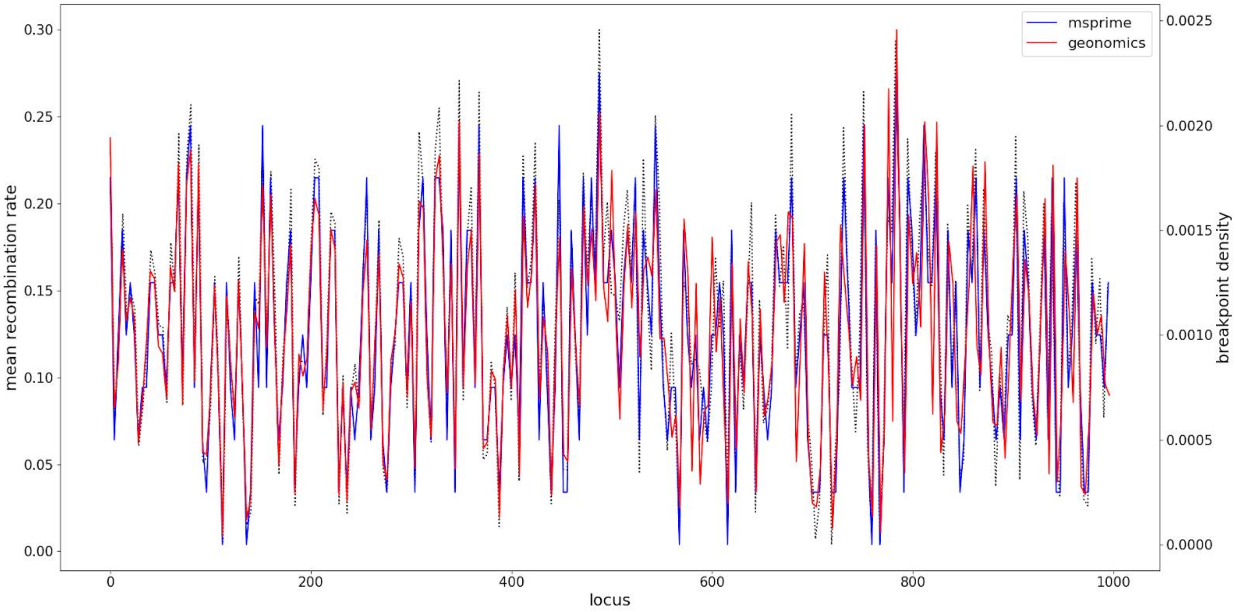
msprime
geonomics
mean recombination rate
0.00
0.05
0.10
0.15
0.20
0.25
0.30
breakpoint density
0.0000
0.0005
0.0010
0.0015
0.0020
0.0025
locus
0
200
400
600
800
1000

Figure S16:

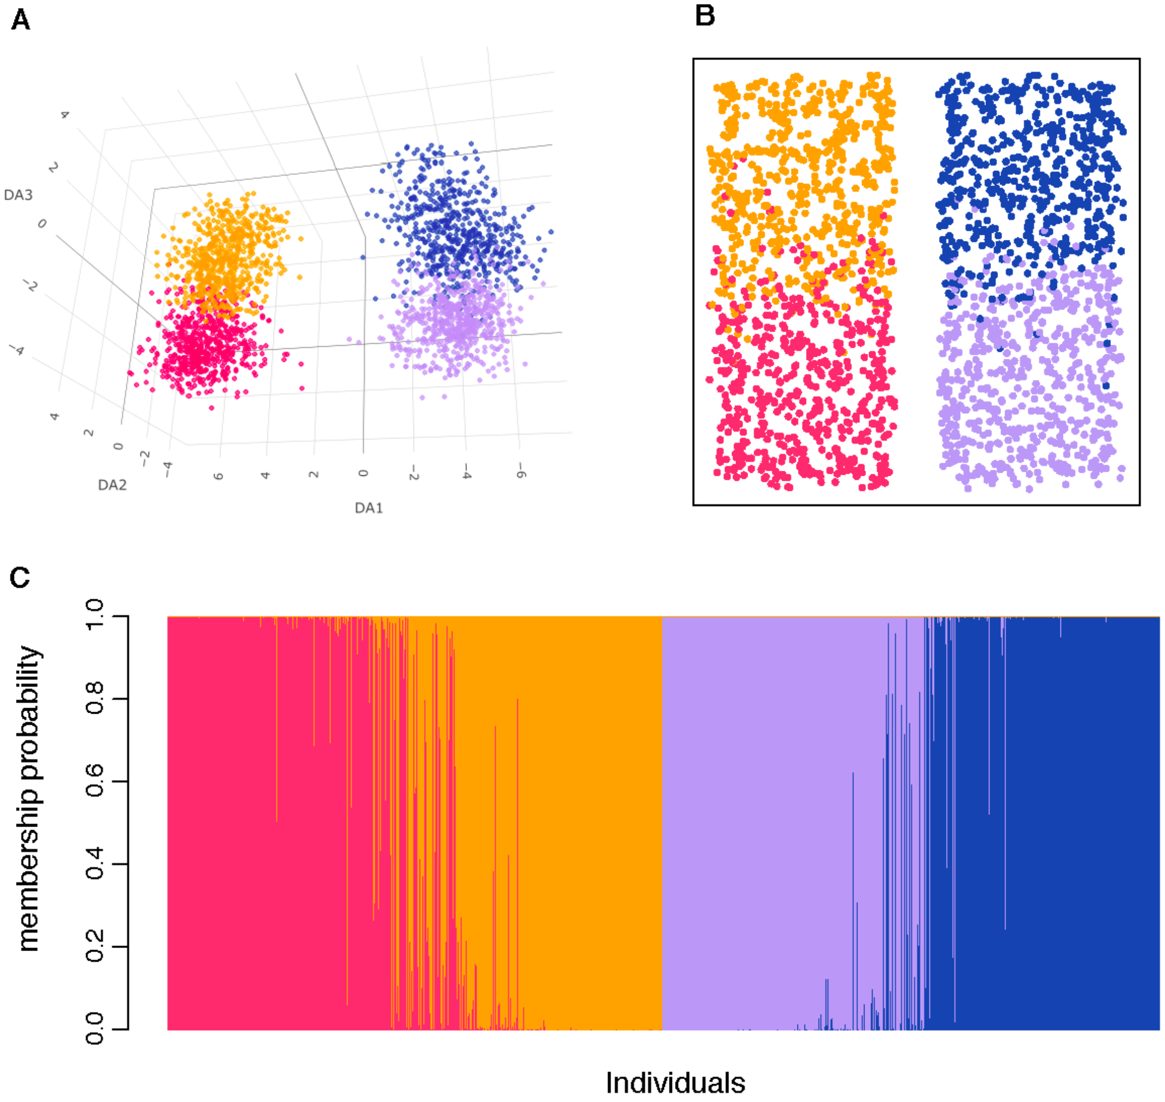
A
DA3
4
2
0
−2
−4
DA2
4
2
0
−2
−4
DA1
6
4
2
0
−2
−4
−6
B
C
membership probability
0.0
0.2
0.4
0.6
0.8
1.0
Individuals

Figure S17:

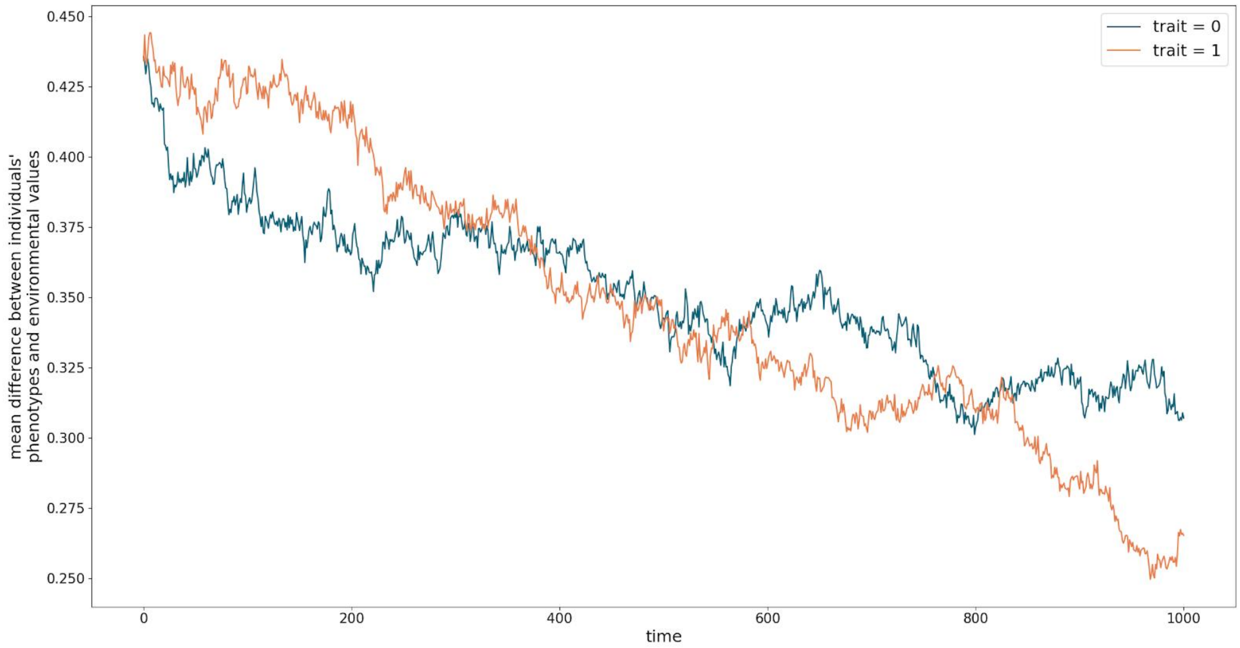

trait = 0
trait = 1
mean difference between individuals'
phenotypes and environmental values
0.450
0.425
0.400
0.375
0.350
0.325
0.300
0.275
0.250
0
200
400
600
800
1000
time

Figure S18:

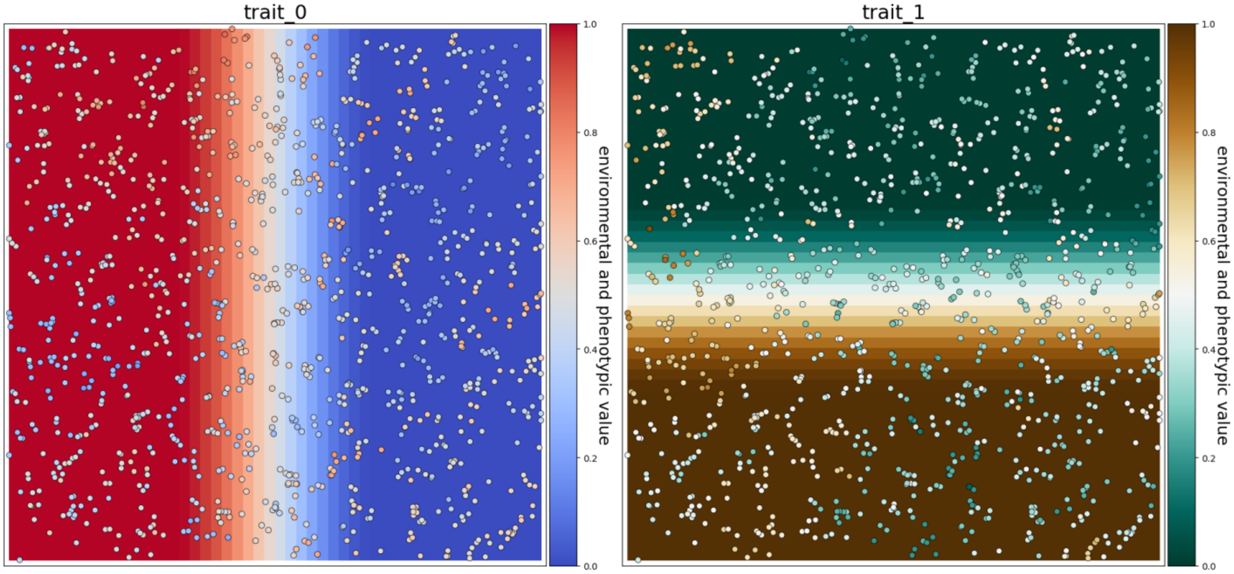
trait_0
trait_1
1.0
0.8
0.6
0.4
0.2
0.0
environmental and phenotypic value

Figure S19:

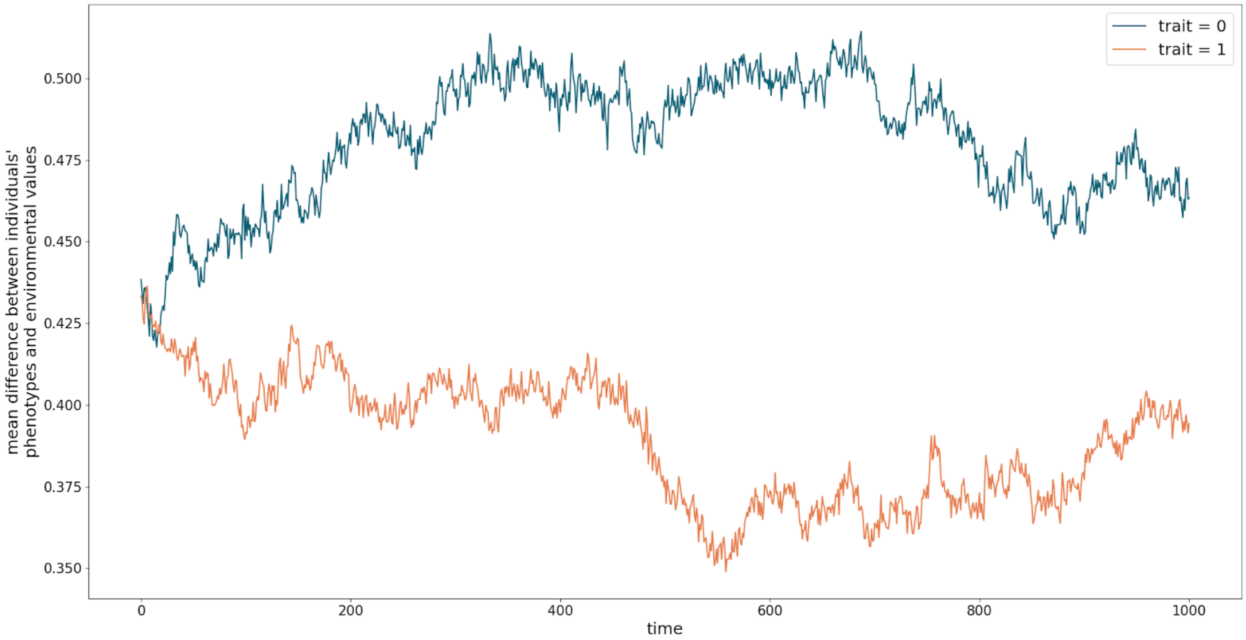

396