





# Accessible, Realistic Genome Simulation with Selection Using `stdpopsim`

Graham Gower <sup>1,†</sup> Nathaniel S. Pope <sup>2,†</sup> Murillo F. Rodrigues <sup>2,3,†</sup> Silas Tittes <sup>2,†</sup>  
Linh N. Tran <sup>4,†</sup> Ornob Alam <sup>5</sup> Maria Izabel A. Cavassim <sup>6</sup> Peter D. Fields <sup>7</sup>  
Benjamin C. Haller <sup>8</sup> Xin Huang <sup>9,10</sup> Ben Jeffrey <sup>11</sup> Kevin Korfmann <sup>2</sup>  
Christopher C. Kyriazis <sup>6</sup> Jiseon Min <sup>2</sup> Inés Rebollo <sup>12</sup> Clara T. Rehmann <sup>2</sup>  
Scott T. Small <sup>2</sup> Chris C. R. Smith <sup>2</sup> Georgia Tsambos <sup>13</sup> Yan Wong <sup>11</sup> Yu Zhang <sup>14</sup>  
Christian D. Huber <sup>15</sup> Gregor Gorjanc <sup>14</sup> Aaron P. Ragsdale <sup>16</sup> Ilan Gronau <sup>17</sup>  
Ryan N. Gutenkunst <sup>4</sup> Jerome Kelleher <sup>11</sup> Kirk E. Lohmueller <sup>6</sup> Daniel R. Schrider <sup>18</sup>  
Peter L. Ralph <sup>2,19</sup> Andrew D. Kern <sup>2,\*</sup>

<sup>1</sup>Section for Molecular Ecology and Evolution, Globe Institute, University of Copenhagen, Copenhagen, Denmark

<sup>2</sup>Institute of Ecology and Evolution, University of Oregon, Eugene, OR 97402, USA

<sup>3</sup>Division of Genetics, Oregon National Primate Center, Oregon Health and Science University, Beaverton, OR 97006, USA

<sup>4</sup>Department of Molecular and Cellular Biology, University of Arizona, Tucson, AZ 85721, USA

<sup>5</sup>Center for Genomics & Systems Biology, New York University, New York, NY 10003, USA

<sup>6</sup>Department of Ecology and Evolutionary Biology, University of California, Los Angeles, Los Angeles, CA, USA

<sup>7</sup>The Jackson Laboratory, Bar Harbor, ME 04609, USA

<sup>8</sup>Department of Computational Biology, Cornell University, Ithaca, NY, USA

<sup>9</sup>Department of Evolutionary Anthropology, University of Vienna, Vienna, Austria

<sup>10</sup>Human Evolution and Archaeological Sciences (HEAS), University of Vienna, Vienna, Austria

<sup>11</sup>Big Data Institute, Li Ka Shing Centre for Health Information and Discovery, University of Oxford, Oxford OX3 7LF, UK

<sup>12</sup>Department of Agronomy and Plant Genetics, University of Minnesota, St. Paul, MN 55108, USA

<sup>13</sup>Department of Genome Sciences, University of Washington, Seattle, WA 98195, USA

<sup>14</sup>The Roslin Institute and Royal (Dick) School of Veterinary Studies, University of Edinburgh, Edinburgh EH25 9RG, UK

<sup>15</sup>Department of Biology, Pennsylvania State University, University Park, PA 16802, USA

<sup>16</sup>Department of Integrative Biology, University of Wisconsin-Madison, Madison, WI, USA

<sup>17</sup>Efi Arazi School of Computer Science, Reichman University, Herzliya 4610101, Israel

<sup>18</sup>Department of Genetics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA

<sup>19</sup>Department of Data Science, University of Oregon, Eugene, OR 97402, USA

<sup>†</sup>These authors contributed equally to the article.

\*Corresponding author: E-mail: [adkern@uoregon.edu](mailto:adkern@uoregon.edu).

Associate editor: Yuseob Kim

## Abstract

Selection is a fundamental evolutionary force that shapes patterns of genetic variation across species. However, simulations incorporating realistic selection along heterogeneous genomes in complex demographic histories are challenging, limiting our ability to benchmark statistical methods aimed at detecting selection and to explore theoretical predictions. `stdpopsim` is a community-maintained simulation library that already provides an extensive catalog of species-specific population genetic models. Here, we present a major extension to the `stdpopsim` framework that enables simulation of various modes of selection, including background selection, selective sweeps, and arbitrary distributions of fitness effects (DFE) acting on annotated subsets of the genome (for instance, exons). This extension maintains `stdpopsim`'s core principles of reproducibility and accessibility while adding support for species-specific genomic annotations and published DFE estimates. We demonstrate the utility of this framework by comparing methods for demographic inference, DFE estimation, and selective sweep detection across several species and scenarios. Our results demonstrate the robustness of demographic inference methods to selection on linked sites, reveal the sensitivity of DFE-inference methods to model assumptions, and show how genomic features, like recombination rate and functional sequence density, influence power to detect selective sweeps. This extension to `stdpopsim` provides a powerful new resource for the population genetics community to explore the interplay between selection and other evolutionary forces in a reproducible, user-friendly framework.

Received: March 23, 2025. Revised: August 12, 2025. Accepted: August 18, 2025

© The Author(s) 2025. Published by Oxford University Press on behalf of Society for Molecular Biology and Evolution.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact [reprints@oup.com](mailto:reprints@oup.com) for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact [journals.permissions@oup.com](mailto:journals.permissions@oup.com).

## Introduction

Selection is a fundamental force in evolution, shaping the genetic diversity of populations and driving the adaptation of species to their environments. The effects of selection on genetic variation are complex, and can be difficult to disentangle from other evolutionary processes such as mutation, recombination, and genetic drift (e.g. Lewontin and Krakauer 1973; Ohta 1973; Kimura 1985; Tajima 1989; Gillespie 1991; Kern and Hahn 2018). For instance, changes in population size can lead to fluctuations in genetic diversity across a recombining chromosome that can mimic the effects of selection (Simonsen et al. 1995), leading to spurious inferences about the strength and targets of genetic adaptation (Simonsen et al. 1995; Akey et al. 2004; Nielsen et al. 2005). In turn, selection can confound our ability to infer demographic history from allele frequencies (Ewing and Jensen 2016; Schrider et al. 2016) and estimates of inverse coalescent rate (Schrider et al. 2016; Johri et al. 2021; Cousins et al. 2024). Thus, it is imperative to jointly account for the effects of selection and demography when inferring evolutionary history from genetic data (e.g. Sheehan and Song 2016; Johri et al. 2020), yet the number of tools for doing so remain limited.

To meet the growing need for the interpretation, analysis, and exploration of realistic and complex evolutionary models, the field of population genetics has increasingly turned to simulation. Simulation in population genetics has a long history including both backward-in-time coalescent simulations (Kingman 1982; Hudson 1983; Hudson et al. 1990) and forward-in-time simulations of complex demography and selection (e.g. Gillespie 1984; Hernandez 2008; Thornton 2014; Haller and Messer 2019; Gaynor et al. 2021). The development of simulation tools has been driven by several needs: understanding the effects of complex evolutionary processes on genetic variation (e.g. Peischl et al. 2013; Thornton 2019; Galloway et al. 2020; Zhang et al. 2020), obtaining a null model for hypothesis testing (e.g. Hudson et al. 1992, 1994; Sabeti et al. 2002), exploring the power and limitations of statistical methods (e.g. Przeworski 2002; MacLeod et al. 2014), and increasingly, providing a basis for machine learning and other simulation-based inference methods (e.g. Beaumont et al. 2002; Lin et al. 2011; Kern and Schrider 2018; Mughal and DeGiorgio 2019; Sanchez et al. 2021; Wang et al. 2021; Smith and Kern 2023; Zhang et al. 2023). However, joint simulation of complex demography and selection on realistic genomes is challenging, requiring many choices including the strength and rate of selection, and a parameterized demographic model, as well as implementation of genetic maps and other genomic features related to the strength of selection. This is a daunting task and can be a barrier to the adoption of simulation-based methods in population genetics. Furthermore, different simulation and inference tools often use different structures and parameter scalings for their models, making it difficult to directly compare across these tools and the studies that use them.

In light of the complexities associated with estimating and simulating population genetic models, it is no surprise that a lingering challenge with simulation in population genetics has been reproducibility and the ability to share and compare results among researchers (e.g. Ragsdale et al. 2020). This challenge has been addressed in part by the development of community resources for sharing and distributing simulation software via the `stdpopsim` project (Adrion et al. 2020).

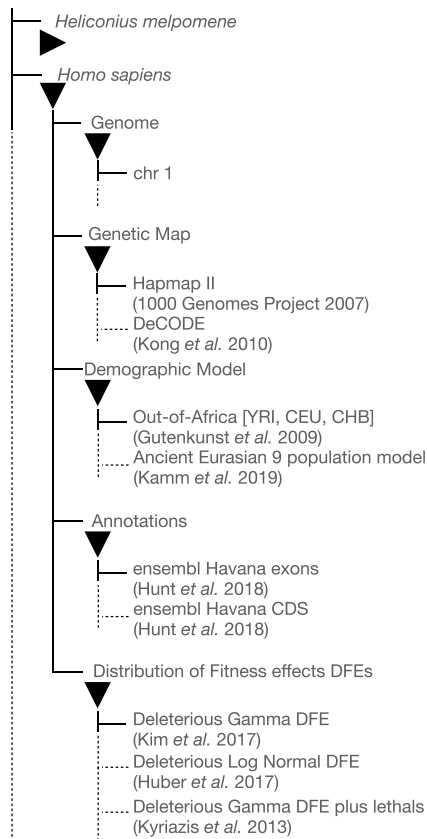
`stdpopsim` provides a standardized interface for accessing a wide range of population genetic models and has been widely adopted by the community (e.g. Speidel et al. 2021; Wang et al. 2021; DeGiorgio and Szpiech 2022; Yang et al. 2022; Browning et al. 2023; Schweiger and Durbin 2023; Temple et al. 2024; Haag et al. 2025). Yet, previous releases of `stdpopsim` did not include models of selection, which is a major limitation for more empirically motivated applications of population genetic simulation. In particular, modeling selection through simulation is critical for understanding processes such as adaptation (e.g. Thornton 2019; Hartfield and Glémin 2024), the effects of selective sweeps (e.g. Braverman et al. 1995; Fay and Wu 2000; Przeworski 2002; Przeworski et al. 2005; Schrider et al. 2015), and the impact of background selection on genetic diversity (e.g. Charlesworth et al. 1993, 1995; Williamson and Orive 2002; Ewing and Jensen 2016; Torres et al. 2020). Ideally, one would like to have a simple, unified framework for simulating both neutral and non-neutral evolutionary processes, and to be able to compare the results of these simulations to empirical data in a manner that is both accessible to a wide range of researchers and highly reproducible. Furthermore, the framework should include the complex realities of genomes, including heterogeneous recombination rates, variation in the size and density of functional elements, and nonequilibrium demographic histories. Moreover, one should be able to model all of these features simultaneously using estimates and annotations obtained from a species of interest (e.g. Schrider 2020; Rodrigues et al. 2024). `stdpopsim` provides a framework for such models, thus lowering the barrier to production of complex, multifaceted simulations, and allowing a wider range of researchers to make use of high-quality simulations in developing or testing their work. A central goal of the `stdpopsim` project is open, collaborative, and inclusive work; we view capacity-building across the community as equal in importance to software development.

Here, we provide an overview of a major new addition to `stdpopsim`—the inclusion of models of selection. We begin by describing the models of selection that we have implemented, as well as the parameter choices that are available to the user. This includes models of background selection, selective sweeps, and models of the distribution of fitness effects (DFE). Further, we describe how these can be applied not only genome-wide but to genomic regions with a particular annotation, specific to a given species available through the `stdpopsim` application programming interface (API), and combined with a parameterized model of demography and recombination to provide a realistic simulation of genetic variation in a population. We then provide a series of examples of how these models can be used to benchmark and compare the performance of different methods for using genetic data to infer demographic history, the distribution of fitness effects, and the targets of recent positive selection. The examples use genomes of both humans (*Homo sapiens*) and the critically endangered Vaquita porpoise (*Phocoena sinus*), which provides an interesting comparison due to its small recent population size.

## Implementing Selection in `stdpopsim`

`stdpopsim` previously provided a wide range of models of neutral demographic history from roughly two dozen species (Lauterbur et al. 2023). This was accomplished through the use of a standardized interface and data structure that allows for the curated addition of new species and models, and can

## (a) Catalog



## (b) Python API

```
example_with_selection.py
1 import stdpopsim
2
3 species = stdpopsim.get_species("HomSap")
4 contig = species.get_contig(
5     "chr22",
6     genetic_map="HapMap_GRCh37"
7 )
8 model = species.get_demographic_model(
9     "OutOfAfrica_3G09"
10 )
11 samples = {"YRI": 5, "CEU": 5, "CHB": 5}
12 # Adding DFE to exons
13 dfe = species.get_dfe("Gamma_K17")
14 exons = species.get_annotations("ensembl_havana_104_exons")
15 exon_intervals = exons.get_chromosome_annotations("chr22")
16 contig.add_dfe(intervals=exon_intervals, DFE=dfe)
17 engine = stdpopsim.get_engine("slim")
18 # Performing the simulation
19 ts = engine.simulate(
20     model,
21     contig,
22     samples,
23     seed=236,
24     slim_scaling_factor=1,
25     slim_burn_in=10,
26 )
```

## (c) Command Line Interface

```
$ stdpopsim -e slim --slim-scaling-factor 1 HomSap \
-c chr22 \
--dfe Gamma_K17 \
--dfe-annotation ensembl_havana_104_exons \
-o example.ts -d OutOfAfrica_3G09 YRI:5 CEU:5 CHB:5
```

**Fig. 1.** a) A schematic of the `stdpopsim` catalog and the different components that can be used in simulations. b) The Python API and c) the command line interface. The user specifies a species, a portion of the genome to simulate, and optionally a genetic map, a model of demography, and a model of selection that itself is composed of a distribution of fitness effects (DFE) and a set of functional annotations.

use either `msprime` (Baumdicker et al. 2021) or `SLiM` (Haller and Messer 2019) as the backend engine for efficient population genetic simulation. To extend `stdpopsim`'s functionality to include models of selection, we introduce two new classes to the `stdpopsim` API: the `DFE` class and the `Annotation` class. The `DFE` class provides an interface for specifying a distribution of the fitness effects (DFE) of new mutations. Each `DFE` object specifies a list of mutation types and their proportions, where each mutation type is associated with a dominance coefficient, the distribution of selection coefficients (e.g. fixed-value, normally distributed, gamma-distributed), and the parameters of this distribution. `DFE` objects can also model a relationship between the selection and dominance coefficients (Huber et al. 2018, see Methods). Each `Annotation` object represents a collection of functional genomic elements (e.g. exons, introns, intergenic regions, etc.) whose coordinates are specified using a GFF3 annotation file. More concretely, consider implementing a model of selection consistent with the DFE inferred by Kim et al. (2017) for exons in the human genome. The associated `DFE` object is defined with two mutation types, neutral and deleterious, corresponding to synonymous and nonsynonymous sites, with their proportions set to their expected prevalence in exons. The neutral mutation type is associated with a fixed selection coefficient (0), and the deleterious mutation type is associated with a dominance coefficient of 0.5 and a gamma-distributed selection coefficient with mean and shape as inferred by Kim et al. (2017).

As with other aspects of `stdpopsim`, we intend the selection models to be as biologically accurate as possible. Thus, the `stdpopsim` catalog provides `Annotation` objects based on species' publicly available functional genomic elements and `DFE` objects based on published DFE estimates (Fig. 1a). By using annotation objects, one can, for example, simulate populations where selected mutations (whose selection coefficients are drawn from a specified DFE) occur within annotation regions (e.g. exons), while all mutations occurring outside of these regions do not affect fitness. There are relatively few species with published DFE estimates in the literature. `stdpopsim` currently has implemented DFEs for five species in the catalog: *Arabidopsis thaliana* (Huber et al. 2018), *Drosophila melanogaster* (Ragsdale et al. 2016; Huber et al. 2017; Zhen et al. 2021), humans (Huber et al. 2017; Kim et al. 2017; Castellano et al. 2019; Kyriazis et al. 2023; Rodrigues et al. 2024), *Mus musculus* (Booker et al. 2021), and the vaquita porpoise *Phocoena sinus* (Robinson et al. 2022). While the DFEs are from a limited number of species, these DFEs can be applied to other species if and when it is suitable, namely, when a species of interest lacks a published DFE in the catalog. For instance, one could simulate a butterfly species with a gamma-distributed DFE originally estimated from humans, or a human population with a DFE estimated from *Drosophila*. Indeed, the DFE labeled `HomSap/Mixed_K23` is recommended by Kyriazis et al. (2023) for "general" use across mammals, when lacking a more specific DFE. Furthermore, the user can specify a custom DFE and

provide their own annotations of functional elements to simulate selection in a species for which we do not yet have a published DFE included in the catalog.

From a user's perspective, a model of selection is specified by pairing a DFE object with a collection of genomic segments defined using an Annotation object. Multiple such pairings can be defined in one simulation to provide a rich model of how selection may vary along a chromosome. For example, coding sequences may be associated with one DFE and non-coding sequences in exons with another DFE. We note that if a genomic interval belongs to more than one Annotation object, then it will be eventually associated with the last Annotation object added to the simulated contig and removed from any previously added annotations. Simulations with selection are then implemented by specifying a species, a portion of the genome to simulate, a genetic map if available, a model of demography, and pairings of DFEs with genomic segments using either the Python API or the command-line interface (Fig. 1b and c). Simulations with selection are generated using the SLiM simulation engine, which is perhaps the most flexible simulation engine for modeling selection available to date (msprime's selection models are relatively limited and thus not currently used by stdpopsim).

While the DFE class facilitates the simulation of non-neutral variation across a sequence, it does not provide a means of specifying the origin and fate of particular selected alleles. Thus, we have also implemented an API in stdpopsim (stdpopsim.selective\_sweep) that enables selective sweeps to be simulated with a high degree of control. This function augments a demographic model by introducing a mutation at a particular time and position along the sequence; sets the selection coefficient of the mutation to a positive value over a given time interval; and conditions the simulation such that the mutation is above a particular frequency at the beginning and end of the sweep period. The conditioning is implemented via rejection sampling; if the advantageous allele is lost or does not have a frequency exceeding a user-specified threshold at the beginning or end of the sweep, the simulation is restarted from the introduction of the advantageous allele. If there are competing advantageous alleles (e.g. simultaneous sweeps at distinct loci) then the simulation is restarted from the introduction of the first of these (see Methods). As this functionality is implemented on top of the existing stdpopsim API, it may be combined with other models of selection and demography to simulate the combined effects of multiple, disparate evolutionary processes on genetic variation in a population.

At this release (version 0.3.0) the stdpopsim catalog comprises 27 species (i.e. genome representations), including vertebrates, invertebrates, plants, and bacteria, for which we have included 35 demographic models, 38 genetic maps, 12 DFEs, and 8 annotation objects. Additions to the catalog are ongoing—we point the reader to our previously published report detailing this effort (Lauterbur et al. 2023). We welcome contributions from the community and encourage those interested in contributing to visit <https://popsim-consortium.github.io/stdpopsim-docs> or contact one of the PopSim Consortium members for guidance.

## Inference of $N_e(t)$ in the Context of Selection

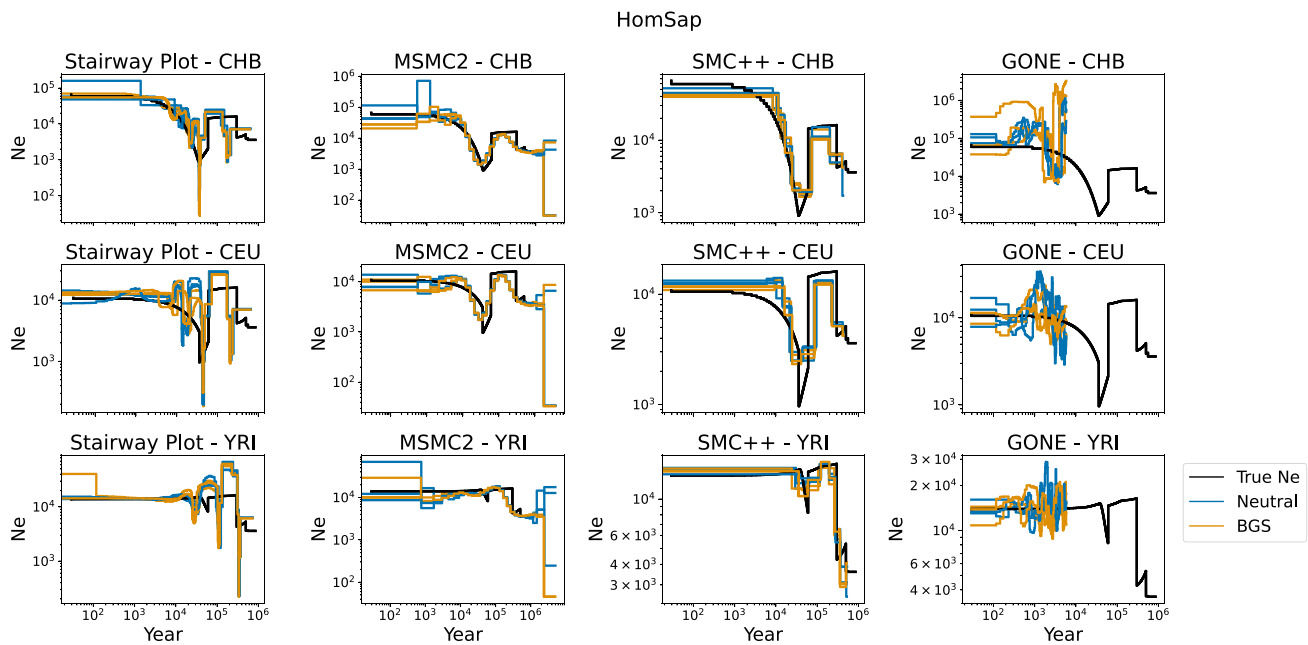
One of the most common applications of population genetic inference is to estimate the effective population size over time,  $N_e(t)$ , from genetic data. This can be done using a variety

of methods, including the sequential Markovian coalescent (Li and Durbin 2011; Terhorst et al. 2017; Schiffels and Wang 2020), through use of the site frequency spectrum (SFS) (Liu and Fu 2020), as well as through linkage-disequilibrium information (Santiago et al. 2020). Since all of these methods assume neutrality, they are typically applied to genomic segments that are not expected to be directly affected by selection (e.g. by masking exons and conserved elements). However, selection acting on linked sites—even those that are masked out of the analysis—can bias estimates of  $N_e(t)$  away from census population sizes by increasing the rate of coalescence in the genome (e.g. Schrider et al. 2016). Using stdpopsim, we can readily examine how different methods for inferring  $N_e(t)$  are influenced by selection on linked sites, which in the case of negative selection is commonly referred to as *background selection* (Charlesworth et al. 1993; Hudson and Kaplan 1995).

To do this, we simulated human genomes with and without selection. In both scenarios, we ran three replicate simulations under the out of Africa (OOA) demographic model of Ragsdale and Gravel (2019) using a genetic map from the HapMap Project (International HapMap Consortium 2007) (see Methods). Simulations with selection were implemented by modeling purifying selection on nonsynonymous mutations using a DFE inferred by Kim et al. (2017). A basic summary of nucleotide diversity from these simulations shows that the presence of selection decreases polymorphism only slightly at the genome-wide level, but more substantially within exonic regions (Fig. S1).

Following the simulation, inference of  $N_e(t)$  was conducted from each of the six simulated datasets using four methods: stairwayplot (Liu and Fu 2020), msmc2 (Schiffels and Wang 2020), SMC++ (Terhorst et al. 2017) and GONE (Santiago et al. 2020). Although the estimates produced by the four methods differ from one another, for each method the estimates produced from data with selection (Fig. 2; orange) and without selection (Fig. 2; blue) are fairly similar. Thus, this particular parameterization of background selection and demography, in combination with selection only at exonic sites, does not appear to notably bias the estimates produced by these methods. This makes sense in light of our observation that genetic diversity is mostly affected within exonic regions and presumably sites very closely linked (Fig. S1), thus masking of exons should be sufficient to remove the major effects of selection. As noted in previous studies that compared demographic inference methods on neutral simulations (Adrion et al. 2020), we see that each method has its strengths and weaknesses. Methods based on the sequential Markovian coalescent (SMC++ and msmc2) appear to produce the most accurate trajectories of  $N_e(t)$  overall in this setting, although they tend to under-estimate the population bottleneck associated with the population split. On the other hand, stairwayplot, which utilizes the site frequency spectrum, produces more noisy estimates, but it does not appear to under-estimate the population bottleneck. The inference of GONE, which utilizes linkage-disequilibrium, is targeted toward demographic changes in more recent time frames (~200 generations), explaining its noisy estimates for all but the most recent time periods.

To expand these observations to additional species and to highlight the ease of comparison enabled by stdpopsim, we performed a similar analysis using the vaquita porpoise *Phocoena sinus*. The vaquita offers an intriguing contrast to the demographic history of humans. While human populations continue to expand, the vaquita is critically endangered,



**FIG. 2.** Inference of  $N_e(t)$  from human genomes simulated under an out-of-Africa demographic model with and without purifying selection on exons. Rows correspond to the three extant populations in the simulation: CHB, CEU, and YRI. Columns correspond to the four inference methods: `stairwayplot`, `msmc2`, `SMC++` and `GONE`. Each plot depicts the inferred  $N_e(t)$  on the three datasets simulated without selection (blue) and the three datasets simulated with the background effects of purifying selection on exons (orange), alongside the true values of  $N_e$  used in simulation (black). The X axis depicts time back from present (in years) and the Y axis depicts  $N_e$  (number of individuals).

with only approximately 10 individuals remaining. Previous demographic inference from genetic data (Robinson et al. 2022) found a long-term effective population size of 4,485 individuals which then decreased to 2,807 individuals 2,162 generations ago. The population then severely decreased even more to its current size in the last few generations. We simulated 100 genomes under a two-epoch demographic model inferred for vaquita by Robinson et al. (2022), and a constant recombination rate across the genome (see Methods). This is particularly interesting because the model was inferred assuming neutrality (with `dadi`) and used to inform conservation efforts for this critically endangered species. In simulations with selection, we applied the DFE inferred by Robinson et al. (2022) to exons annotated from the vaquita genome assembly. This DFE implements a relationship between selection and dominance in which very deleterious mutations (with selection coefficient  $s < -0.1$ ) are fully recessive (dominance coefficient  $h = 0$ ) and weakly deleterious mutations ( $s \geq -0.001$ ) are nearly additive ( $h = 0.4$ ; see Methods). As in the human simulations, we see that the presence of selection decreases polymorphism in genomes simulated with selection, but that the effect is only readily apparent within exonic regions (Fig. S2). For these simulations, the SFS-based method (`stairwayplot`) produces more accurate inference than the two methods based on the sequential Markovian coalescent (`SMC++` and `msmc2`; see Fig. S3). Regardless of these differences, we see that background selection again does not appear to considerably influence the inference of any of the four methods for this species, at least for the parameterization used here, even though vaquita has a very different evolutionary history than humans.

### Estimation of the Distribution of Fitness Effects

Another common application of population genetic inference is to estimate the distribution of fitness effects (DFE) of new

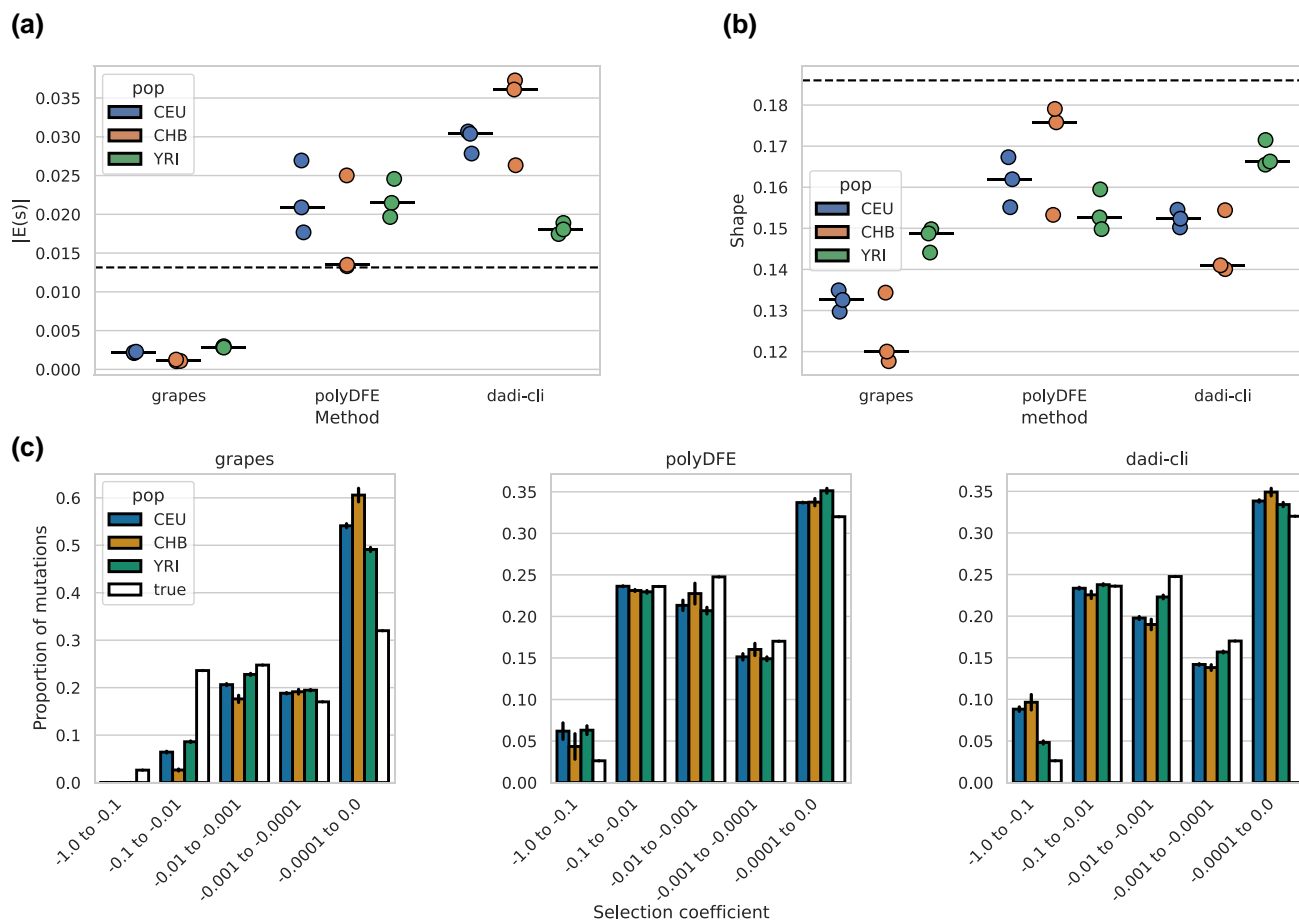
mutations from genetic data. The new framework of modeling selection in `stdpopsim` makes it an ideal tool for easily benchmarking and comparing the performance of different methods for inferring the DFE. Here, we considered three DFE-inference methods: `GRAPES` (Galtier 2016), `polyDFE` (Tataru and Bataillon 2020), and `dadi-cli` (Huang et al. 2023).

A challenge when comparing different inferred and simulated distributions of the selection coefficient ( $s$ ) is differing conventions in the literature regarding its definition. The simulator `SLiM` and the inference tool `GRAPES` define the selection coefficient  $s$  such that a homozygote has fitness  $1 + s$ . On the other hand, `dadi-cli` and `polyDFE` define the selection coefficient such that a homozygote has fitness  $1 + 2s$ . Moreover, inferred distributions of selection coefficients are typically scaled by an inferred ancestral effective population size  $N_e$ , and different methods assume different scaling. `dadi-cli` infers the distribution of  $2N_e s$  while `polyDFE` and `GRAPES` infer the distribution of  $4N_e s$ . To allow direct comparison, here we have normalized all inference results to the convention used by `SLiM` (also used by `stdpopsim`; see Methods).

We start by examining DFE inference from human genomes simulated under an out-of-Africa demographic model with selection acting on exons (see Methods). In this setting, the DFE is inferred separately for each of the three simulated populations using each of the three DFE-inference methods. In each case we note that the DFE is estimated from segregating sites only, without the use of substitution data, as `stdpopsim`'s models currently do not contain outgroup populations or species that can be used to estimate divergence.

Overall, inferred DFEs varied across populations and inference methods (Fig. 3, Table 1). Note that while the estimation accuracy for the DFE parameters themselves is informative (Fig. 3a and b), examining the difference between the actual and inferred distributions of  $s$  can be more helpful for showing how well a method has recovered the true DFE (Fig. 3c). The

## HomSap



**Fig. 3.** Inferred versus simulated DFEs from simulated human genomes. Simulations were performed using a human out-of-Africa demographic model with a gamma-distributed DFE acting on exons (see Methods). DFE is inferred separately for each of the three extant populations (CEU, CHB, and YRI) by three different methods: *GRAPES*, *polyDFE*, and *dadi-cli*. a and b) Mean absolute value of selection coefficient ( $|E(s)|$ ) and shape parameter are shown for each DFE inferred from all three simulation replicates, with median values marked by horizontal bars and simulated values represented by dashed horizontal lines. c) Binned distribution of  $s$  implied by the average DFE inferred from the three simulated datasets (averaging the inferred gamma parameters), with error bars depicting the standard error of the mean for each bin. White bars represent the distribution used in the simulation.

**Table 1.** Performance of DFE-inference methods in four simulation scenarios

Species ID	Demography	MAE $E s $ dadi-cli	MAE $E s $ grapes	MAE $E s $ polyDFE	MAE shape dadi-cli	MAE shape grapes	MAE shape polyDFE
HomSap	Constant	<b>0.0012</b>	0.0101	0.0023	0.030	0.0086	<b>0.0068</b>
HomSap	OOA SR19	0.014	0.0111	<b>0.0072</b>	0.031	0.051	<b>0.024</b>
PhoSIn	Constant	<b>0.024</b>	0.025	<b>0.024</b>	<b>0.23</b>	0.25	<b>0.23</b>
PhoSIn	Vaquita2Epoch 1R22	<b>0.024</b>	0.025	<b>0.024</b>	<b>0.18</b>	0.23	0.21

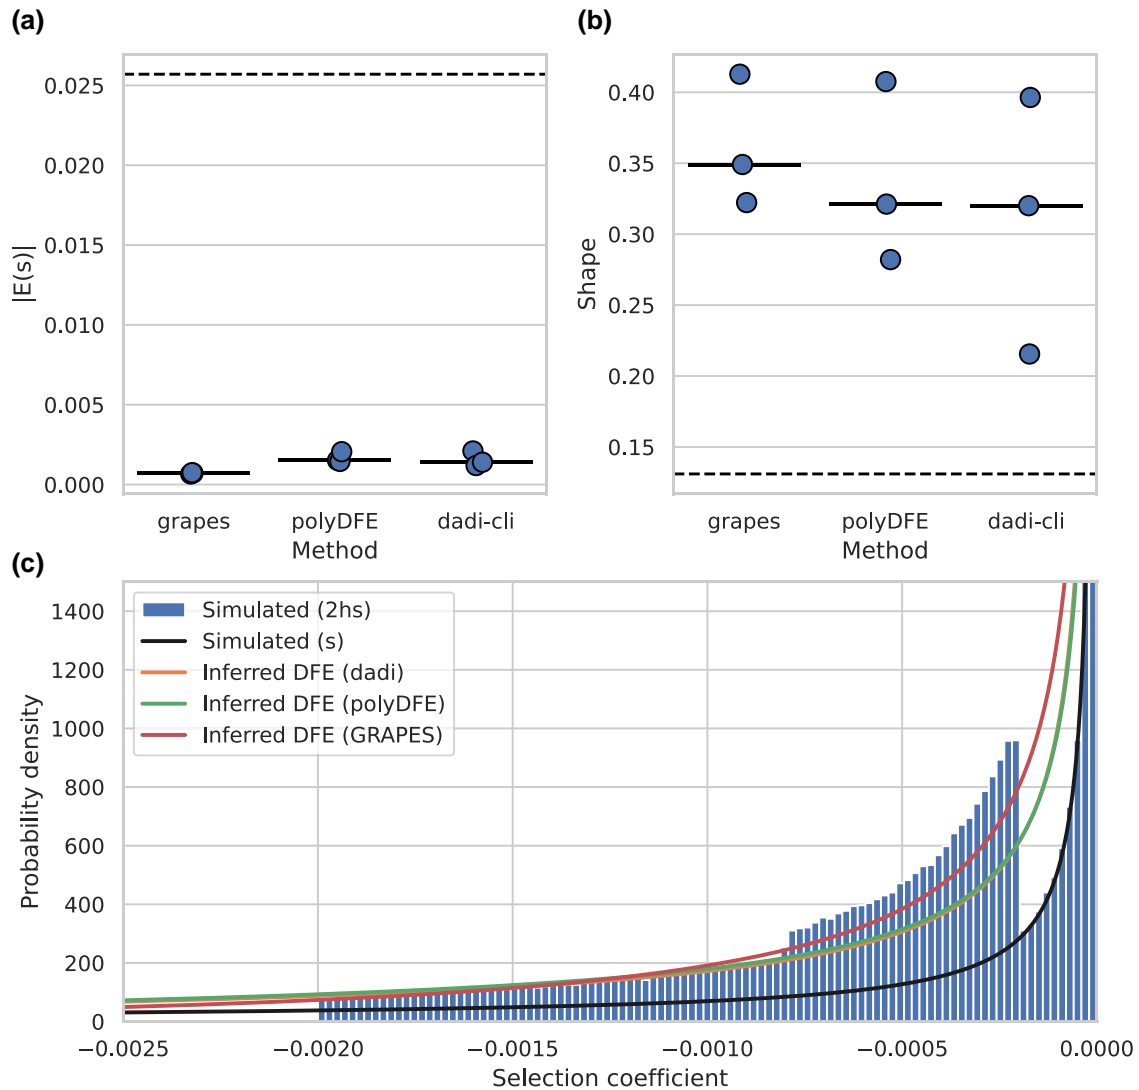
For each of the three methods, we report the mean absolute error (MAE) of the mean selection coefficient and shape of its inferred DFE. In each simulation scenario, the lowest MAEs are highlighted in bold.

distributions of  $s$  inferred by *polyDFE* and *dadi-cli* had higher mean values (in absolute value) but smaller shape parameters, while *GRAPES* inferred distributions with lower mean values and shape parameters. As a result, *polyDFE* and *dadi-cli* over-estimate the fraction of mutations under strong selection ( $s < -0.1$ ), while *GRAPES* considerably over-estimates the fraction of mutations under weak selection ( $s \geq -0.0001$ ).

Population history appears to influence different methods in slightly different ways. For instance, *dadi-cli* and *GRAPES* produce more accurate estimates from the YRI genomes (relative to the other two populations), whereas *polyDFE* shows

very similar accuracy across all populations. To factor out the influence of demographic changes, we also analyzed genomes generated using similar simulations but with no population size changes (Fig. S4, Table 1). In this simpler setting, *polyDFE* overlapped the true mean selection coefficient and shape parameter, whereas the other methods do not show as close a correspondence. *GRAPES* exhibits similar biases as observed in the simulations with more complex demography, but with milder under-estimation of the shape parameter, whereas *dadi-cli* shows an inverse pattern with overestimated shape parameter and slightly underestimated mean. While providing possible insights into the effects of the demographic history on

## PhoSIn



**Fig. 4.** Inferred versus simulated DFEs from simulated vaquita genomes. Simulations were performed using a two-epoch model of vaquita porpoise demography with a gamma-distributed DFE acting on nonsynonymous mutations with a relationship between the selection coefficient ( $s$ ) and dominance coefficient ( $h$ ) (see Methods). The DFE is inferred by analyzing all simulated genomes jointly by one of three different methods: GRAPES, polyDFE, and dadi-cli. a and b) Mean absolute value of selection coefficient ( $|E(s)|$ ) and shape parameter are shown for each DFE inferred from all three simulation replicates, with median values marked by horizontal bars and simulated values represented by dashed horizontal lines. c) Binned distribution of  $2hs$  implied by the simulated DFE compared with the distribution of  $s$  implied by the average DFE inferred for each method from the three simulated replicates (averaging the inferred gamma parameters). The distribution of  $2hs$  is multimodal because of the simulated relationship between  $h$  and  $s$  (see text).

accuracy of different DFE-inference methods, substantiating these insights will require more than three replicate simulations per scenario.

Next, we examined the performance of these three DFE-inference methods when assuming a different DFE. To do this, we again turned to the vaquita. Unlike the human simulations where all mutations were assumed to be additive ( $h = 0.5$ ), the vaquita DFE applied a gamma-distributed selection coefficient to mutations in exons with a relationship between the selection and dominance coefficients (see Methods). Specifically, more deleterious mutations in the vaquita DFE are assumed to be more recessive. In this setting, all methods perform uniformly worse than in the human genome simulations, with consistent underestimation of the mean selection coefficient and overestimation of the shape parameter

(Fig. 4a and b, Table 1). Similar patterns are observed when re-running the simulations with a simpler demographic model without population size changes (Fig. S5, Table 1). We suspect that this is due to the fact that all DFE-inference methods used here assume an additive model ( $h = \frac{1}{2}$ ), while mutations in the simulated genomes were more recessive ( $h < \frac{1}{2}$ ), and their dominance decreased with the strength of selection. To examine the influence of this model misspecification, we compared the inferred distributions of  $s$  to the simulated distribution of  $2hs$ , which would have been equal to  $s$  under an additive model (Fig. 4c). Reassuringly, all three methods infer distributions that appear to fit the main mode of the simulated distribution of  $2hs$ . This is likely because deleterious alleles are typically at low frequency and thus heterozygous, where their selective effect is  $hs$  rather than  $s$ .

## Detection of Selective Sweeps

We next highlight how `stdpopsim` can be used to benchmark and compare the performance of different methods for detecting selective sweeps from genetic data. Because simulation of selective sweeps is time consuming and evaluation of statistical power requires many replicate simulations, we focused on human chromosome 1 and generated replicate simulations on 5 cM segments of the chromosome in 93 different locations. Selective sweeps were simulated by introducing a beneficial mutation in the center of the 5 cM segment with a moderately strong selection coefficient ( $s = 0.03$ ;  $2N_e s \sim 600$ ; see Methods). Each segment was simulated 1,000 times, while keeping only replicates in which the selected allele reached a frequency of 95% or greater. These sweeps were modeled in the context of the three population out-of-Africa model from Gutenkunst et al. (2009) and background selection from deleterious mutations in exons. Sweep-detection methods typically base their inference on checking whether a test statistic exceeds a certain threshold, which is determined according to some empirical null distribution. In our simulations, we considered two different null distributions: one based on neutral simulations, and another based on simulations with background selection from deleterious mutations in exons (see Methods).

We evaluated the statistical power to detect selective sweeps for three methods: `sweepfinder2` (DeGiorgio et al. 2016), `diploshic` (Kern and Schrider 2018), and reduced diversity ( $\pi$ ). For each method, we determined a threshold for its test statistic using each of the two null models separately—neutral and with background selection (BGS)—and controlling for a 5% false positive rate. We then applied each method to data simulated with sweeps in one of two populations (YRI and CEU) in each of the simulated locations along human chromosome 1 (Fig. 5). First and most strikingly, the power to detect sweeps varies considerably along the chromosome, and this variation is much greater than the variation in power between CEU and YRI or between null models. This is likely due to the effects of recombination rate heterogeneity (Fig. 5). Next, we notice that power to detect sweeps is slightly lower when the test statistic threshold is determined using simulated data with background selection (BGS null model). This is likely because the strength of selection acting on exons is relatively weak, and the sweep detection methods we have used are relatively robust to the effects of background selection. Finally, the power to detect sweeps is generally higher in the YRI sample than in the CEU sample, which is consistent with the fact that YRI has a larger ancestral effective population size and thus both stronger selection and more variation to detect a sweep (e.g. Simonsen et al. 1995). Conversely, the CEU population experienced the out-of-Africa bottleneck that may mute the signal of a sweep.

We compared the overall performance of the three sweep-detection methods using receiver operating characteristic (ROC) curves (Fig. S6). Performance varies for each of the sampled populations. In the CEU sample, `sweepfinder2` and `diploshic` perform similarly with a high area under the curve (AUC) score, (AUC = 0.906 and 0.903, respectively), but both underperform the simple method that uses reduced diversity ( $\pi$ ) as an indicator (AUC = 0.932). In the YRI sample, `diploshic` outperforms the other two methods, with an AUC of 0.979, followed by using outliers of  $\pi$  with an AUC of 0.957, and then `sweepfinder2` with an AUC of 0.939. We note that we trained `diploshic` on simulations

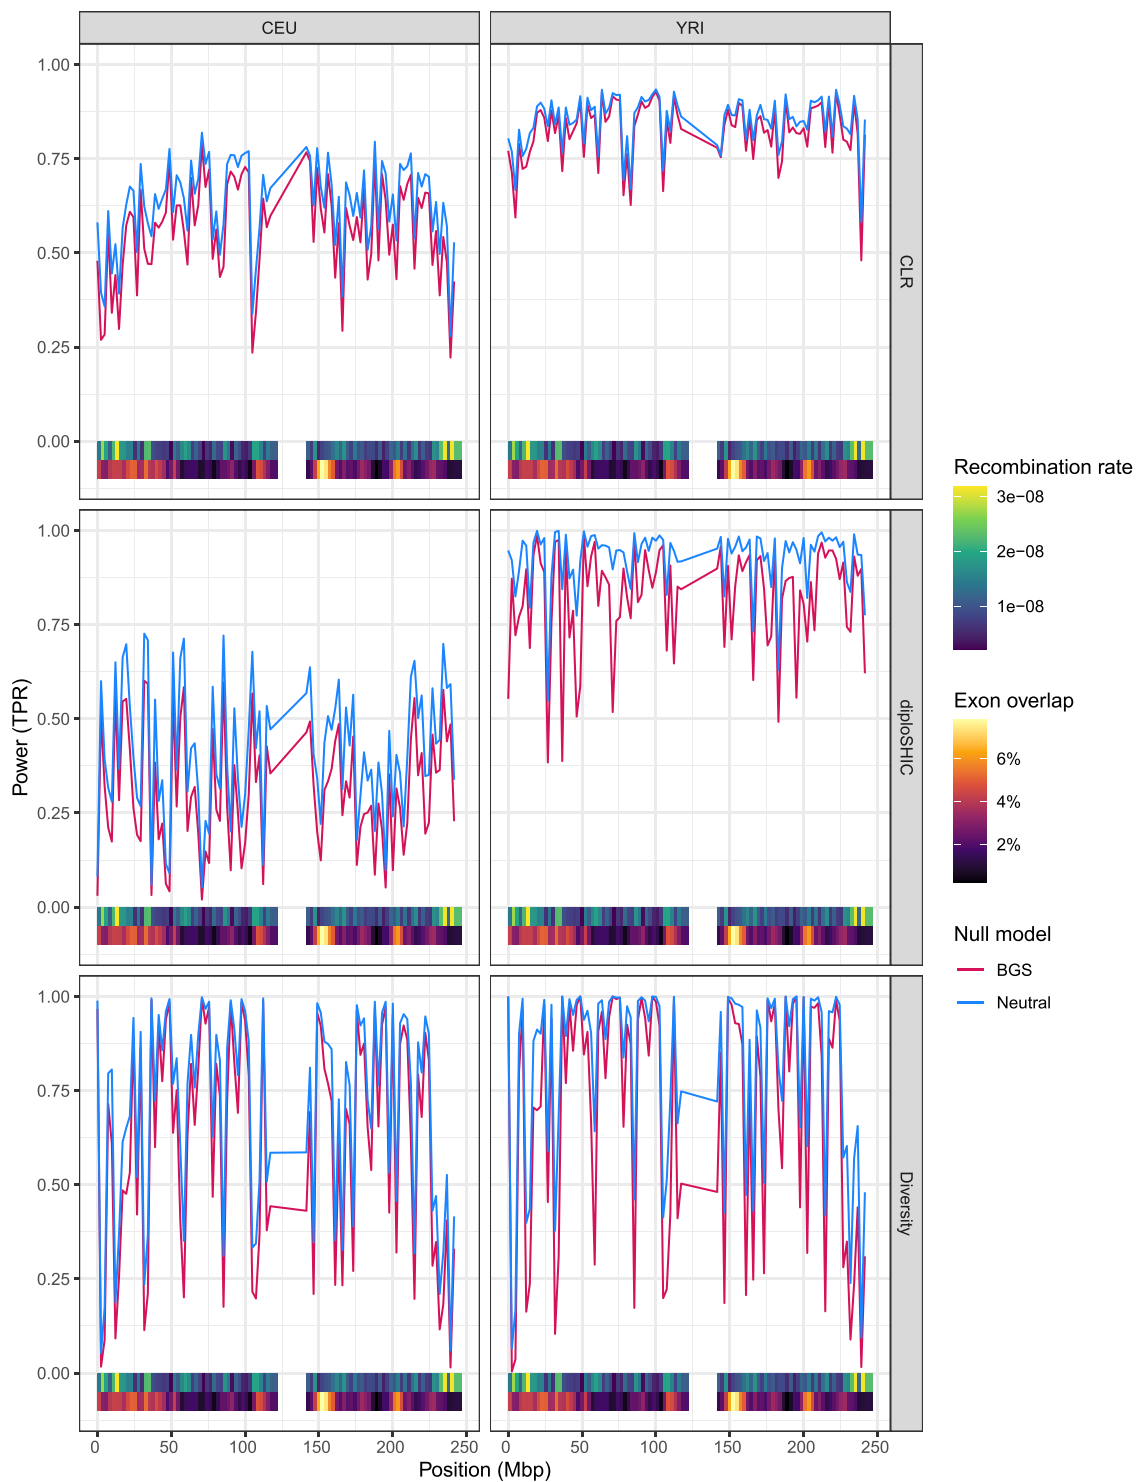
under a constant size population, which more closely matches the demographic history of YRI, than that of CEU, which is characterized by a strong recent bottleneck followed by population expansion. Interestingly, the simple method that uses reduced diversity ( $\pi$ ) as an indicator for a selective sweep performs very well overall (especially on the CEU sample). However, its accuracy is much more variable along the chromosome compared to the other two methods (see Fig. 5), and for quite a few locations along the chromosome it has power below 50% when applied to both sample sets. `sweepfinder2` appears to have the least variation in power along the chromosome or across populations.

To further dissect the reasons for high variance along the chromosome in power to detect sweeps, we examined the relationship between power and local recombination rate (Fig. 6). We see that power to detect sweeps is a decreasing function of local recombination rate for `sweepfinder2` and reduced diversity, which is consistent with sweeps having a larger genomic footprint in regions of low recombination. `diploshic` on the other hand shows a slight increase in power with recombination rates, presumably because it is designed to distinguish regions where a sweep happened from regions linked to a sweep, and if the recombination rate is too low (or selection too strong) then the genomic window examined by `diploshic` may not be large enough to exhibit enough of a recovery in patterns of diversity flanking the sweep. These results confirm that reduced genetic diversity performs well as an indicator for selective sweeps only in regions with low recombination rate. The other two methods, which were specifically designed for sweep detection, are considerably more robust to the effects of recombination rate than the simple indicator of reduced genetic diversity. These methods are particularly robust to recombination rate variation when applied to the YRI sample, especially when using a neutral null model to calibrate the test statistics. The other null model, which assumes background selection from exons, appears to consistently lead to reduced power for detecting sweeps regardless of the recombination rate. Interestingly, we do not see a clear correlation between power to detect sweeps and local exon density (Fig. S7). This is somewhat surprising given that background selection may be expected to interfere with the sweep signature (reviewed by Stephan 2010). This might suggest that the effects of realistic background selection on genome variation may not be as similar to those of sweeps as previously thought (Schrider 2020).

## Discussion

In this article, we have presented an important new addition to the `stdpopsim` library that enables the simulation of genetic variation in the presence of selection. We have demonstrated the utility of this new API by showing several illustrative examples of how it can be used to benchmark the performance of different methods for demographic history inference, the inference of the distribution of fitness effects of new mutations, and the detection of selective sweeps. This allows for a wide range of models of selection to be easily simulated and compared to empirical data, with reproducibility, computational efficiency, and rigor.

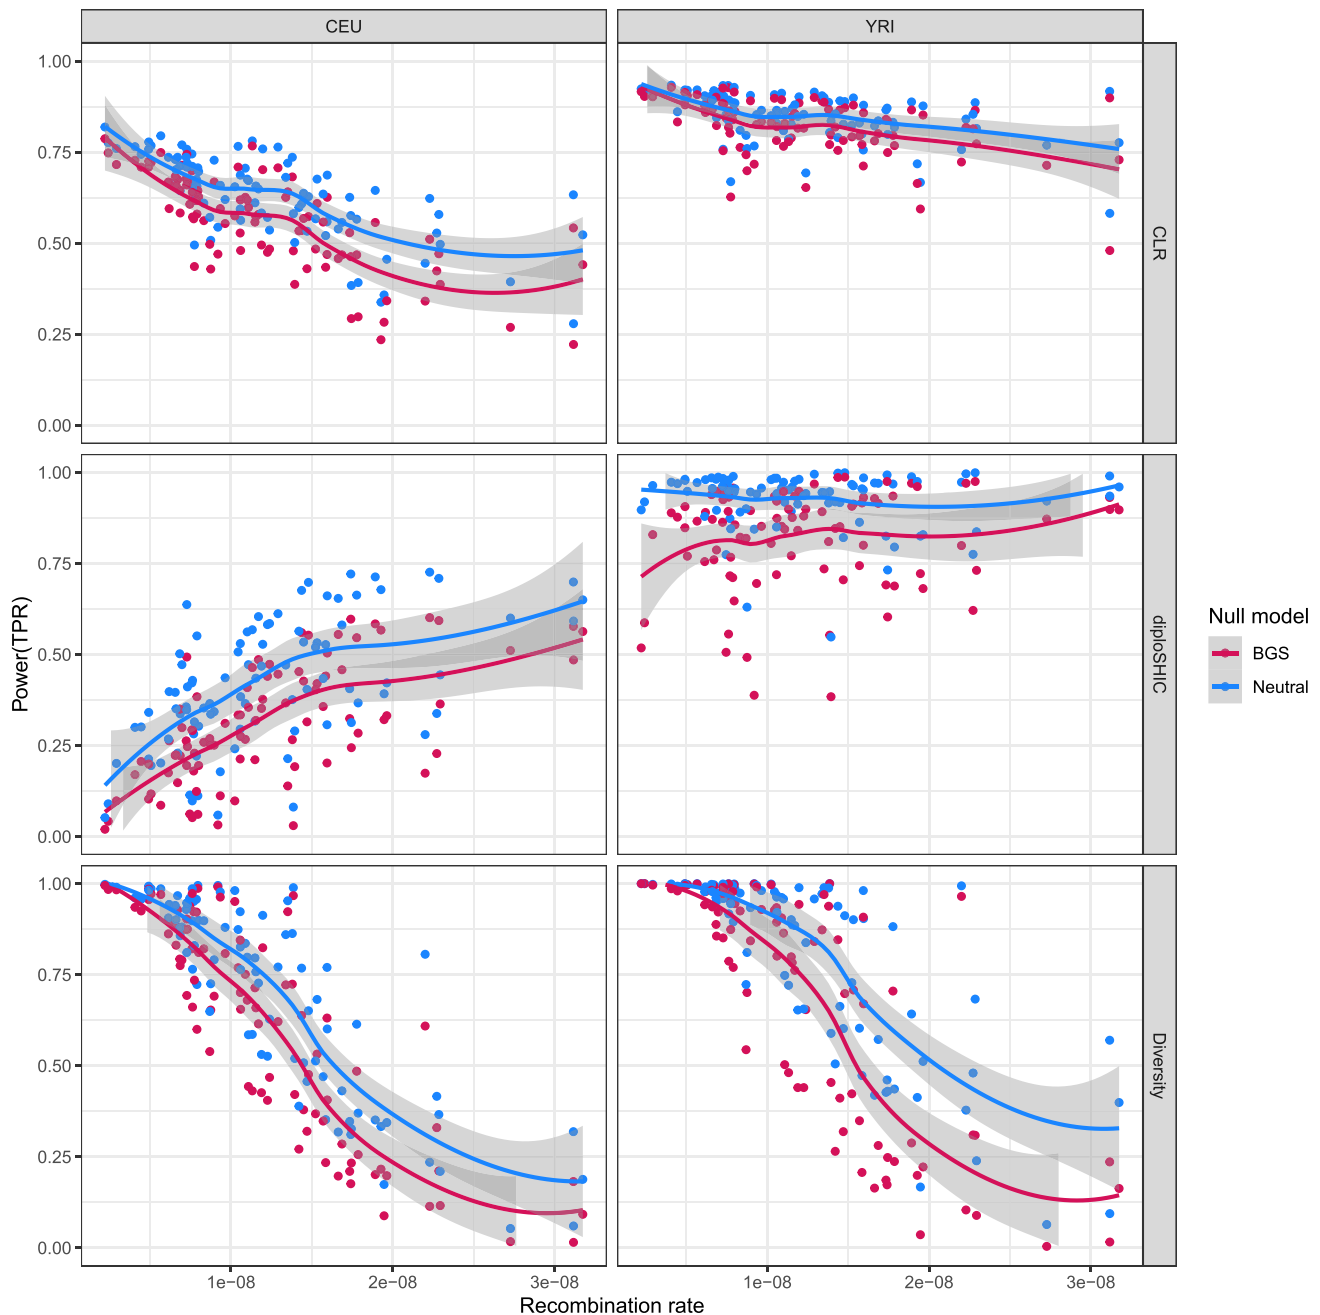
The exploration of performance among methods done here is not meant to be thorough benchmarking. Rather, we include these analyses as examples to highlight the utility of `stdpopsim` for both benchmarking and other tasks where researchers would benefit from the ability to perform



**Fig. 5.** Power to detect selective sweeps along human chromosome 1. Genomic segments were simulated with sweeps under a three population out-of-Africa model and with background selection from deleterious mutations in exons. Three methods for detecting sweeps were applied to simulated data: `sweepfinder2` (top row—labeled CLR), `diploshic` (middle row), and reduced diversity ( $\pi$ ) (bottom row). Power (true positive rate) is shown for these methods for the CEU and YRI samples (left and right, respectively). The thresholds of the test statistics were set to control for 5% false positive rate under a neutral null model (blue) and a null model with background selection from deleterious mutations in exons (red). Also shown are heatmaps of local recombination rate and exon density along the chromosome (bottom).

simulations using parts of published population genetic models in a streamlined and quality-controlled fashion. Having said that, there are a number of interesting results that merit discussion. The first pertains to estimation of demographic history ( $N_e(t)$ ) in the context of linked purifying selection. Our simulations of human and vaquita genomes model

deleterious mutations only at exonic positions of the genome. As shown in [Figs. S1](#) and [S2](#), this combination only produces sizable decreases in polymorphism at exonic positions, which represent a small minority of sites within the genome, and does not lead to large-scale perturbations in genome-wide diversity. Thus, it is not surprising that estimates of  $N_e(t)$  from genome-



**Fig. 6.** Power to detect selective sweeps as a function of local recombination rate. This figure shows the same power estimates shown in Fig. 5, but with the genomic segments plotted against their average recombination rates instead of position along chromosome 1. Genomic segments were simulated with sweeps under a three population out-of-Africa model and with background selection from deleterious mutations in exons. Three methods for detecting sweeps were applied to simulated data: *sweepfinder2* (top row—labeled CLR), *diploSHIC* (middle row), and reduced diversity ( $\pi$ ) (bottom row). Power (true positive rate) is shown for these methods for the CEU and YRI samples (left and right, respectively). The thresholds of the test statistics were set to obtain a 5% false positive rate under a neutral null model (blue) and a null model with background selection from deleterious mutations in exons (red). Fitted lines represent loess smoothed regressions.

wide data from which exons have been masked should be accurate (e.g. Figs. 2 and S3). Indeed, our results echo those of Marsh and Johri (2024) who found similar results on realistic human simulations that included selection. However, the question remains to what extent this would be true in empirical data where selection is not constrained to coding positions. For instance, we know that many genomes harbor large numbers of conserved noncoding elements (e.g. Siepel et al. 2005) with extremely strong selection acting on them (Katzman et al. 2007; McLean and Bejerano 2008). We thus caution regarding extrapolation from our results. Future

studies might build upon the annotations we used here to include selection at other portions of the genome as well as mixtures of DFEs across sites.

A second important finding is that inference of the DFE is sensitive to the dominance of deleterious mutations. All three of the DFE-inference methods applied here assume that deleterious mutations act in an additive fashion ( $b = 0.5$ ) by default. However, the simulated DFE for the vaquita genome included highly recessive mutations. Specifically, these simulations used a DFE where mutations with  $s < -0.01$  are quite recessive, with  $b \leq 0.01$ , and where mutations with

$-0.01 < s < -0.001$  have  $h = 0.1$ . We found that the three DFE-inference methods underestimated  $E[s]$  and overestimated the shape parameter (Fig. 4a and b). This performance is not surprising since decreasing  $h$  for a given distribution of  $s$  will lead to a decrease in  $hs$ , or an apparent decrease in the amount of selection. Thus, fitting a mis-specified DFE that assumes  $h = 0.5$  leads to underestimates of  $E[s]$ . Our analysis suggests that the distributions of  $s$  estimated by these methods provide a reasonable approximation for the true distribution of  $hs$  (Fig. 4c). This points to a notable limitation of most existing DFE inference methods, which assume that fitness effects ( $hs$ ) are gamma-distributed. As shown in Fig. 4c, while a gamma distribution might fit the marginal distribution of  $s$ , a relationship between  $h$  and  $s$  might result in fitness effects ( $hs$ ), which are not gamma-distributed.

Overall, these results are consistent with prior work showing that it is challenging to separately infer the DFE of  $s$  along with  $h$  (Wright 1937; Simons et al. 2014; Veeramah et al. 2014; Fuller et al. 2019; Balick et al. 2022; Kyriazis and Lohmueller 2024; reviewed in Fuller et al. 2019) and that a large proportion (>5%) of recessive deleterious mutations can confound inferences of  $s$  (Wade et al. 2023). Given the plethora of evidence for recessive deleterious mutations (Mukai et al. 1972; Agrawal and Whitlock 2011; Huber et al. 2018; Di and Lohmueller 2024), these results suggest the need for care in DFE inference. Kyriazis and Lohmueller (2024) inferred the DFE of  $s$  for nonsynonymous mutations in humans when considering different values of  $h$ . Intriguingly, additive models showed a satisfactory fit to the SFS of nonsynonymous variants in humans, and more complex models of dominance did not substantially improve the fit. Nevertheless, models including recessive deleterious mutations, like those simulated here for the vaquita genome, show a similar fit to the SFS and may be more biologically realistic than fully additive models, given the evidence for the presence of recessive deleterious mutations. More broadly, new computational approaches are needed to precisely quantify the DFE and the covariance of  $h$  and  $s$ .

Further, inference of the DFE was done here using exon annotations of the genome in question. A key distinction between our analyses and those done in empirical work is that we have complete knowledge of which sites are under selection in the simulation. Our inferences therefore represent a best case scenario. Empirical inference of the DFE is a much harder problem, since some mutations assumed to be free of selection might in reality be under selection, and the strength of selection can change over time or across environments.

Regarding sweep detection, the inclusion of selection models into `stdpopsim` allowed us to readily compare the effectiveness of different methods using data that reflects the demographic history, recombination rate variation, and genomic locations of negatively selected sites present in human populations. These results clearly demonstrate the impact of recombination rate variation, and underscore the importance of training methods on data with variable recombination rates (e.g. Schrider and Kern 2017). In addition, our results illustrate a strength of likelihood methods like `sweepfinder2` that do not require a predefined window size. Window-based methods, such as `diploshic`, tend to lag in performance when the genomic region affected by the sweep is larger than the chosen window size, as can occur with very strong sweeps or low levels of recombination. Moreover, while all the methods examined here retain adequate power for sweep identification across scenarios,

including the inclusion of background selection, our results underscore the importance of developing methods with adequate statistical power even in the face of unknown demographic histories. Indeed, larger differences in power were observed when applying a method to genomes simulated with different demographic histories (i.e. YRI versus CEU) than when calibrating a method using a null model with and without background selection.

Although we have shown how the new `stdpopsim` API can be used to benchmark the performance of different methods, the comparisons shown here merely scratch the surface of what can be done with the new API. We envision that, moving forward, many researchers will use the new API to simulate data under a wide range of models of selection, demography, and genome architecture, and to develop new methods for the analysis of genetic data in the presence of selection. The additions described here further expand the utility of `stdpopsim` for both characterizing the patterns of diversity expected under a wide range of population genetic models, and for benchmarking methods for the analysis of genetic data on common ground in a way that was not possible before.

## Methods

### Annotations

Currently available annotation data in `stdpopsim` were downloaded from Ensembl. Coding sequence and exon annotations are available for *Arabidopsis thaliana* (Cheng et al. 2017), *Drosophila melanogaster* (Hoskins et al. 2015), *Homo sapiens* (Hunt et al. 2018), and *Phocoena sinus* (vaquita porpoise, Morin et al. 2021).

### Implementing a DFE Model Using `stdpopsim`'s DFE Class

The DFE class has the following attributes for modeling natural selection in a given set of annotated genomic segments. Similarly to other `stdpopsim` classes, the DFE class has fields for specifying the model identifier, description strings for documentation, and a citation object for the publication containing information on the inference on which the DFE model is based. A DFE model is specified using a mixture of mutation types, each associated with a frequency and a distribution of the selection coefficient ( $s$ ). Thus, each DFE object contains a list of `MutationType` objects and a separate list (of identical length) of positive frequencies (that sum up to 1). Each `MutationType` object is associated with a distribution type, distribution parameters, and dominance coefficients. The distribution types that are currently supported by `stdpopsim` are:

- Fixed ( $f$ ), with one parameter specifying the fixed selection coefficient.
- Exponential ( $e$ ), with one parameter specifying the mean of the exponential distribution.
- Gamma ( $g$ ), with two parameters specifying the mean and shape of the gamma distribution.
- Normal ( $n$ ), with two parameters specifying the mean and standard deviation of the normal distribution.
- Weibull ( $w$ ), with two parameters specifying the scale and shape of the Weibull distribution.
- Uniform ( $u$ ), with two parameters specifying the minimum and maximum of the interval.

- Positive log-normal (1p), with two parameters specifying the mean and standard deviation of the normal distribution in log scale.
- Negative log-normal (1n), with two parameters specifying the mean and standard deviation of the normal distribution in log scale.

Typically, one would associate a `MutationType` object with a single dominance coefficient (using the `dominance_coeff` attribute):  $b = 0$  for recessive mutations,  $b = \frac{1}{2}$  corresponding to an additive model (default value), and  $b = 1$  for completely dominant mutations. The dominance coefficient,  $b$ , can take any real value, including negative values (for underdominance) and values greater than 1 (overdominance). There is also an option to model a relationship between the selection and dominance coefficients by associating different ranges of selection coefficients with different dominance coefficients. This is done using two attributes of the `MutationType` class: `dominance_coeff_list` is used to specify a list of  $n$  dominance coefficients, and `dominance_coeff_breaks` is used to specify  $n - 1$  selection coefficient break points. The  $i$ th dominance coefficient in the list is associated with selection coefficients in the range between the  $(i - 1)$ th and the  $i$ th break points. For an explicit example, see the description of the vavita porpoise genome simulations below.

During simulation, when a mutation is generated in an annotated region associated with a given DFE object, a mutation type is sampled according to the frequencies of the different types. Then, a selection coefficient is sampled for that mutation based on the distribution associated with the sampled mutation type. Finally, a dominance coefficient is associated with the mutation either based on the fixed dominance coefficient of the mutation type (using attribute `dominance_coeff`), or based on the sampled selection coefficient (using attributes `dominance_coeff_list` and `dominance_coeff_breaks`). For additional implementation details on the DFE class, see the `stdpopsim` docs at <https://popsim-consortium.github.io/stdpopsim-docs/main/introduction.html>.

### Selective Sweeps and Rejection Sampling

Selective sweeps are implemented in the SLiM engine by introducing a mutation with a fitness effect of  $hs$  in heterozygotes and  $s$  in homozygotes, where  $b$  is the dominance coefficient (as for the DFE model) and  $s \geq 0$  is the selection coefficient. The allele is introduced in a given population at a specified time and genomic position, and may be either globally or locally adaptive (so that the fitness effect is confined to the population of origin) over a fixed duration. The start of the sweep may be offset from the introduction of the allele, to model selection on standing genetic variation (so-called “soft sweeps”).

Crucially, the advantageous allele must remain extant over the course of the sweep; this conditioning is accomplished via rejection sampling, whereby the simulation is restarted from the time of introduction of the mutation if the allele is lost in the population of origin. The conditioning may optionally be extended to reject trajectories where the frequency at the beginning and end of the sweep period is below a specified threshold. If there are competing advantageous alleles (simultaneous sweeps at distinct loci) and one of these does not satisfy the acceptance conditions, then the simulation is restarted from the onset of the first sweep; this avoids conditioning on particular (partial) allele frequency trajectories for the

remaining alleles. A downside of this strategy is that rejection becomes more likely with increasing numbers of simultaneous sweeps or with hard-to-satisfy constraints; but it allows for greater flexibility in sweep specification and for interoperability with the `stdpopsim` DFE model.

### Quality Control and Software Development

Following the protocol established previously for other aspects of `stdpopsim` (Adrion et al. 2020), new DFEs are subject to a rigorous quality control (QC) process, by which a different individual re-implements the same DFE, and the two implementations are programmatically tested for equality. This QC process has proven its effectiveness on numerous occasions in detecting implementation bugs and resolving modeling ambiguities. For example, this process detected an error in coding the log-normal distribution of the *Drosophila melanogaster* DFE inferred by Booker et al. (2021, `stdpopsim` label `LognormalPlusPositive_R16`). This faulty DFE model was used in a study conducted by one of our groups before it underwent the QC process, where we noticed that deleterious selection coefficients were several orders of magnitude larger than expected. Subsequently, the QC process was able to find and fix this bug. In addition to QC, all changes and additions to the code must pass a comprehensive set of unit tests: at the time of writing, the library includes over 11,000 lines of testing code, comprising 2,862 unit tests, including comparisons to QC implementations and cross-checks for agreement, including between SLiM and `msprime` implementations of the same models. For instance, for testing we make use of SLiM’s scriptability to separately record as metadata the simulation parameters (current population sizes, DFEs, etc.), which are then programmatically compared to the models used to generate the SLiM script that ran the simulation. We make use of automatic test “code coverage” tools to help ensure that all portions of the code base are validated by these unit tests.

### Whole Genome Simulations

In this study, we performed simulations of the human and vavita porpoise genomes using `stdpopsim`. We ran replicate simulations for each genome/parameter combination, in which chromosomes were simulated separately from one another. All simulations were executed using the SLiM engine with a burn-in of  $10N$  generations (where  $N$  is the initial population size), and a scaling factor of 1 (no scaling). We note that aggressive scaling of forward-in-time simulations can lead to quite different outcomes when selection is involved (Uricchio and Hernandez 2014; Ferrari et al. 2024; Dabi and Schrider 2025), and we caution users against using high degrees of scaling for drawing final inferences, although we recommend it during development to save on computational costs. All simulations were implemented using a `snakemake` workflow, where the code utilized is similar to that shown in Fig. 1b. The complete simulation workflow is available at <https://github.com/popsim-consortium/analysis2/blob/main/workflows/simulation.snake> (see Data availability below).

### Simulations of Complete Human Genomes

Simulation of human genomes used a genetic map from the HapMap Project (International HapMap Consortium 2007, `stdpopsim` label `HapMapII_GRCh38`) and the out of Africa demographic model with archaic admixture from Ragsdale and Gravel (2019, `stdpopsim` label

OutOfAfricaArchaicAdmixture\_5R19). Simulations with selection additionally applied a DFE model with gamma-distributed selection coefficients inferred by Kim et al. (2017, `stdpopsim` label `Gamma_K17`) to exons annotated by the HAVANA group release 104 (Hunt et al. 2018, `stdpopsim` label `ensembl_havana_104_exons`), which cover 2.1% of the genome. Each simulation generated 100 diploid genomes from each of three extant human populations (labeled YRI, CEU, and CHB), with each genome consisting of 22 autosomes. Using these settings, we generated three replicate datasets without selection and three replicate datasets with selection. Three additional datasets, used in Fig. S4, were simulated with the same model of selection and a simple demographic model with one population with a constant size of 10,000. Each of these simulations generated 100 diploid genomes. The configuration file used in combination with our `snakemake` workflow for these simulations can be found at [https://github.com/popsim-consortium/analysis2/blob/main/workflows/config/snakemake/production\\_config\\_HomSap.yml](https://github.com/popsim-consortium/analysis2/blob/main/workflows/config/snakemake/production_config_HomSap.yml). The running times of these simulations depend on chromosome length, the demographic model, and incorporation of selection (Fig. S8).

### Simulations of Complete Vaquita Porpoise Genomes

Simulation of vaquita porpoise genomes used the two-epoch demographic model from Robinson et al. (2022, `stdpopsim` label `Vaquita2Epoch_1R22`). Since no genetic map is currently available for this species, we assumed a constant recombination rate of  $10^{-8}$  across all chromosomes (Morin et al. 2021). Simulations with selection additionally applied a DFE model with gamma-distributed selection coefficients inferred by Robinson et al. (2022, `stdpopsim` label `Gamma_R22`) to exons annotated from the vaquita porpoise genome assembly (Morin et al. 2021, `stdpopsim` label `Phocoena_sinus.mPhoSin1.pri.110_exons`), which cover 2.2% of the genome. This DFE model implements a relationship between selection and dominance as follows: very deleterious mutations (with  $s < -0.1$ ) are recessive ( $h = 0$ ), mutations with  $s \in [-0.1, -0.01]$  have a dominance coefficient of  $h = 0.01$ , mutations with  $s \in [-0.01, -0.001]$  have a dominance coefficient of  $h = 0.1$ , and very mild mutations (with  $s \in [-0.001, 0]$ ) are nearly additive ( $h = 0.4$ ). This discretized relationship matches that used by Robinson et al. (2022) and is specified using the following attributes in the DFE class: `dominance_coeff_list=(0, 0.01, 0.1, 0.4)`, and `dominance_coeff_breaks=(-0.1, -0.01, -0.001)`.

Each simulation generates 100 diploid genomes, with each genome consisting of 21 autosomes. Using these settings, we generated three replicate datasets without selection and three replicate datasets with selection. Three additional datasets, used in Fig. S5, were simulated with the same model of selection and a simple demographic model with one population with a constant size of 3,500. The configuration file used in combination with our `snakemake` workflow for these simulations can be found at [https://github.com/popsim-consortium/analysis2/blob/main/workflows/config/snakemake/production\\_config\\_PhSin.yml](https://github.com/popsim-consortium/analysis2/blob/main/workflows/config/snakemake/production_config_PhSin.yml).

### Inference of $N_e(t)$

Inference of  $N_e(t)$  was conducted on simulated datasets after masking nonrecombining portions of the genome (to avoid bias caused by extreme recombination rate variation) and

exons (to avoid the direct effect of natural selection). Thus, inference from datasets simulated with selection was not influenced by direct selection acting on exons, but it was influenced by selection acting on linked sites (background selection). We assessed the accuracy of demographic inferences from four software packages: `msmc2` (Schiffels and Wang 2020), `stairwayplot` (Liu and Fu 2020), `GONE` (Santiago et al. 2020), and `SMC++` (Terhorst et al. 2017). For `msmc2`, we used a random sample of  $N = 6$  simulated genomes with 20 iterations of the EM algorithm. For `GONE` we set `max_snps` to 500,000, the number of generations to 2,000, and the number of bins to 400, all representing default settings. For `stairwayplot`, we did one run of  $\theta$  estimation for each replicate dataset, and used a `dimFactor` of 5,000, both of which represent default settings. For `SMC++`, we used the unfolded SFS (`--unfold`), along with the true mutation rate under which the simulations were generated. The complete `snakemake` workflow for demographic inference can be found at [https://github.com/popsim-consortium/analysis2/blob/main/workflows/n\\_t.snake](https://github.com/popsim-consortium/analysis2/blob/main/workflows/n_t.snake).

### DFE Inference

We used three software packages to infer the distribution of fitness effects (DFE) from genetic data: `dadi-cli` (Kim et al. 2017; Huang et al. 2023), `polyDFE` (Tataru and Bataillon 2020), and `GRAPES` (Galtier 2016). Each software package was run using the same set of simulations. DFE inference using `dadi-cli` was conducted by first inferring a demographic model from neutral sites, using a two epoch demographic model with 100 repeated runs of the optimization procedure. `dadi-cli` then proceeds by using a cached strategy wherein we used 2000 points in the gamma distribution. Finally, the DFE was inferred using a ratio of nonsynonymous to synonymous mutations of 2.31, again with 100 repeated runs of the optimization procedure. For `polyDFE`, we used model “C” which uses a mixture of a gamma distribution and an exponential distribution, where the gamma component is used for negatively selected mutations. For optimization we used the default BFGS optimizer with the `-e` flag to automatically estimate the initial parameters values. For `GRAPES`, we inferred the DFE under the “GammaZero” model, which assumes that selection coefficients are drawn from a single gamma distribution. We used the `-no-div-param` flag to disable the use of divergence-based parameter estimation. All other settings were left at their default values. All methods used the same exon annotations as generated in the simulations. The total number of coding sites in the human and vaquita genomes used were 61,644,801 and 31,039,513, respectively.

To facilitate direct comparison between inference methods, we scale each program’s output to obtain a selection coefficient ( $s$ ) consistent with the one used by `SLiM`. `GRAPES` outputs the mean and shape of the inferred distribution of  $4N_e s$ . We thus scale the mean by a factor of  $\frac{\mu}{\theta}$ , where  $\mu$  is the assumed mutation rate and  $\theta = 4N_e \mu$  is the population-scaled mutation rate, as estimated by `GRAPES`. `polyDFE` is similarly scaled by  $\frac{2\mu}{\theta}$ , reflecting the fact that it defines the selection coefficient such that a homozygote has fitness  $1 + 2s$ . `dadi-cli` defines the selection coefficient similarly to `polyDFE` and infers a gamma distribution for  $2N_e s$ , parameterized by the shape ( $\alpha$ ) and rate ( $\nu$ ). It also infers the total population-scaled mutation rate for selected sites over the considered region  $\theta = 4N_e \mu L p$ , where  $L$  is the total number of sites available

for the analysis, and  $p$  is the proportion of sites under selection. Thus, the mean value of  $s$  in SLiM terms is obtained by  $4\mu avLp/\theta$ . Complete details for the DFE inference analysis can be found in the `snakemake` workflow at <https://github.com/popsim-consortium/analysis2/blob/main/workflows/dfc.snake>.

### Simulations with Selective Sweeps

To test methods for detecting selective sweeps, we considered 100 evenly distributed 5 Mb segments of human chromosome 1. We included a 2.5 cM buffer on either side of each segment to minimize the effect of chromosome ends on patterns of genetic variation. Seven segments in which the central 1 Mb had an average recombination rate below  $10^{-11}$  were excluded from the analysis. Each segment was simulated in each of the following three settings: (i) neutral, (ii) background selection (BGS), and (iii) hard sweep with BGS. In all three settings, we applied a genetic map from the HapMap Project ([International HapMap Consortium 2007](#), `stdpopsim` label `HapMapII_GRCh38`) and the three-population out of Africa (OOA) demographic model from [Gutenkunst et al. \(2009\)](#), `stdpopsim` label `OutOfAfrica_3G09`). In the two settings with BGS, we applied the DFE model with gamma-distributed selection coefficients inferred by [Kim et al. \(2017\)](#), `stdpopsim` label `Gamma_K17`) to exons annotated by the HAVANA group release 104 ([Hunt et al. 2018](#), `stdpopsim` label `ensembl_havana_104_exons`). Hard sweeps were modeled by generating a single beneficial mutation with selection coefficient of  $s = 0.03$  in the center of the 5 cM genomic segment and discarding simulations in which the beneficial mutation did not reach (or exceed) a terminal frequency of 95%. For each of the 93 genomic segments of chromosome 1, we generated 1,000 replicate datasets under each of the three settings. We used the SLiM engine with a burn-in of  $2N$  generations (where  $N$  is the initial population size), and a scaling factor of 2. Each simulation generated ten diploid genomic segments for each of the two sampled populations: YRI and CEU. The sequences simulated with a hard sweep and BGS were used for sweep detection, and sequences simulated under the other two settings were used to calibrate the thresholds of the test statistics used by the three sweep-detection methods (see below). See [Fig. S9](#) for the distribution of runtimes for these simulations.

### Selective Sweeps Detection

Each of the three sweep-detection methods is based on a test statistic. Reduced nucleotide diversity is based on nucleotide diversity ( $\pi$ ) computed directly from the underlying tree sequences using `tskit` ([Ralph et al. 2020](#)) over ten equally sized and nonoverlapping subwindows. `sweepfinder2` ([DeGiorgio et al. 2016](#)) is based on a composite likelihood ratio (CLR) computed at 21 uniformly distributed positions along the simulated region. `diploshic` ([Kern and Schrider 2018](#)) uses supervised learning to train a sweep classifier. Its training phase used `discoal` ([Kern and Schrider 2016](#)) to simulate 20 haploid sequences of length 1.21 Mbp with a constant population size of  $N_e = 10,000$ , a mutation rate of  $2 \times 10^{-8}$ , and a recombination rate of  $2 \times 10^{-8}$ . For `discoal` simulations that included selective sweeps, selection coefficients were drawn from a uniform distribution between 0.001 and 0.10, and the time of the sweep was chosen from a uniform distribution between 0 and  $0.01N_e$  generations in the past.

For soft sweeps, the frequency at which selection was introduced was drawn from a uniform distribution between 0.0001 and 0.1. We note that since the training set of `diploshic` was generated under a simple model with constant mutation and recombination rates and constant effective population size, we expect it to perform suboptimally on the simulated segments of human chromosome 1, which were generated using more realistic simulations (see Simulations with selective sweeps above).

The three sweep-detection methods were associated with thresholds for their respective test statistics ( $\pi$ , `sweepfinder2`'s CLR, or `diploshic`'s probability), which were determined using an empirical distribution representing a null model without sweeps. We considered two null models: a neutral evolution model (neutral), and a model with background selection from deleterious mutations in exons (BGS). Each null model was represented by a different set of 93,000 simulations generated according to the appropriate model (see Simulations with selective sweeps above). Thus, for each test statistic we determined two thresholds corresponding to these two different null models (neutral and BGS) using the extreme 5% quantile in each empirical distribution. Then, power (or true positive rate) to detect a sweep by a given method was determined for positions on chromosome 1 as follows: each sweep-detection method was applied to the 1,000 simulated datasets generated with a hard sweep and BGS for that position, and the power was set to be the proportion of simulations (out of 1,000) in which the test statistic exceeded the threshold determined by each of the two null models (neutral and BGS). We note that using an empirical distribution to determine the threshold for calling a selective sweep with `diploshic` deviates from the original implementation of the method, however for consistency among methods we chose to take this approach. The complete `snakemake` workflow for sweep detection (including simulations) can be found at [https://github.com/popsim-consortium/analysis2/blob/main/workflows/sweep\\_simulate.snake](https://github.com/popsim-consortium/analysis2/blob/main/workflows/sweep_simulate.snake).

### Supplementary Material

Supplementary material is available at [Molecular Biology and Evolution](#) online.

### Acknowledgments

We thank Tom Booker for discussions and encouragement to incorporate realistic BGS models into `stdpopsim`. We thank the numerous users of the `stdpopsim` package, as well as workshop participants who provided feedback and suggestions for improvements.

### Funding

M.R., S.T., N.P., and A.K. were supported in part by NIH award R35GM148253. R.G. and L.T. were supported in part by NIH award R35GM149235. D.S. was supported in part by NIH award R35GM138286. K.L. was supported by NIH award R35GM119856. P.R., B.H., J.K., and A.K. were supported in part by NIH award R01HG012473. G.G. and Y.Z. were supported in part by BBSRC (BBS/E/RL/230001A, BBS/E/RL/230001C, BB/T014067/1), Chinese Scholarship Council (CSC) and Genus. C.H. was supported by NIH award R35GM146886.

## Data Availability

The release of `stdpopsim` is available at <https://github.com/popsim-consortium/stdpopsim>. This article coincides with version 0.3.0 of `stdpopsim`. The entire analysis pipeline associated with this article is available as a snakemake workflow at <https://github.com/popsim-consortium/analysis2>.

## References

- Adrión JR *et al.* A community-maintained standard library of population genetic models. *Elife*. 2020;9:e54967. <https://doi.org/10.7554/eLife.54967>.
- Agrawal AF, Whitlock MC. Inferences about the distribution of dominance drawn from yeast gene knockout data. *Genetics*. 2011;187:553–566. <https://doi.org/10.1534/genetics.110.124560>.
- Akey JM *et al.* Population history and natural selection shape patterns of genetic variation in 132 genes. *PLoS Biol*. 2004;2:e286. <https://doi.org/10.1371/journal.pbio.0020286>.
- Balick DJ, Jordan DM, Sunyaev S, Do R. Overcoming constraints on the detection of recessive selection in human genes from population frequency data. *Am J Hum Genet*. 2022;109:33–49. <https://doi.org/10.1016/j.ajhg.2021.12.001>.
- Baumdicker F *et al.* Efficient ancestry and mutation simulation with `msprime` 1.0. *Genetics*. 2021;220:iyab229. <https://doi.org/10.1093/genetics/iyab229>.
- Beaumont MA, Zhang W, Balding DJ. Approximate Bayesian computation in population genetics. *Genetics*. 2002;162:2025–2035. <https://doi.org/10.1093/genetics/162.4.2025>.
- Booker TR, Jackson BC, Craig RJ, Charlesworth B, Keightley PD. 2021. Selective sweeps influence diversity over large regions of the mouse genome [preprint]. bioRxiv 447924. <https://doi.org/10.1101/2021.06.10.447924>.
- Braverman JM, Hudson RR, Kaplan NL, Langley CH, Stephan W. The hitchhiking effect on the site frequency spectrum of DNA polymorphisms. *Genetics*. 1995;140:783–796. <https://doi.org/10.1093/genetics/140.2.783>.
- Browning SR, Waples RK, Browning BL. Fast, accurate local ancestry inference with `flare`. *Am J Hum Genet*. 2023;110:326–335. <https://doi.org/10.1016/j.ajhg.2022.12.010>.
- Castellano D, Macià MC, Tataru P, Bataillon T, Munch K. Comparison of the full distribution of fitness effects of new amino acid mutations across great apes. *Genetics*. 2019;213:953–966. <https://doi.org/10.1534/genetics.119.302494>.
- Charlesworth B, Morgan M, Charlesworth D. The effect of deleterious mutations on neutral molecular variation. *Genetics*. 1993;134:1289–1303. <https://doi.org/10.1093/genetics/134.4.1289>.
- Charlesworth D, Charlesworth B, Morgan M. The pattern of neutral molecular variation under the background selection model. *Genetics*. 1995;141:1619–1632. <https://doi.org/10.1093/genetics/141.4.1619>.
- Cheng C-Y *et al.* `Araport11`: a complete reannotation of the *Arabidopsis thaliana* reference genome. *Plant J*. 2017;89:789–804. <https://doi.org/10.1111/tpj.2017.89.issue-4>.
- Cousins T, Tabin D, Patterson N, Reich D, Durvasula A. 2024. Accurate inference of population history in the presence of background selection [preprint]. bioRxiv 576291. <https://doi.org/10.1101/2024.01.18.576291>.
- Dabi A, Schrider DR. Population size rescaling significantly biases outcomes of forward-in-time population genetic simulations. *Genetics*. 2025;229:iyae180. <https://doi.org/10.1093/genetics/iyae180>.
- DeGiorgio M, Huber CD, Hubisz MJ, Hellmann I, Nielsen R. `Sweepfinder2`: increased sensitivity, robustness and flexibility. *Bioinformatics*. 2016;32:1895–1897. <https://doi.org/10.1093/bioinformatics/btw051>.
- DeGiorgio M, Szpiech ZA. A spatially aware likelihood test to detect sweeps from haplotype distributions. *PLoS Genet*. 2022;18:e1010134. <https://doi.org/10.1371/journal.pgen.1010134>.
- Di C, Lohmueller KE. Revisiting dominance in population genetics. *Genome Biol Evol*. 2024;16:evae147. <https://doi.org/10.1093/gbe/evae147>.
- Ewing GB, Jensen JD. The consequences of not accounting for background selection in demographic inference. *Mol Ecol*. 2016;25:135–141. <https://doi.org/10.1111/mec.2016.25.issue-1>.
- Fay JC, Wu C-I. Hitchhiking under positive Darwinian selection. *Genetics*. 2000;155:1405–1413. <https://doi.org/10.1093/genetics/155.3.1405>.
- Ferrari T, Feng S, Zhang X, Mooney J. 2024. Towards simulation optimization: an examination of the impact of scaling on coalescent and forward simulations [preprint]. bioRxiv 591463. pages 2024–04. <https://doi.org/10.1101/2024.04.27.591463>.
- Fuller ZL, Berg JJ, Mostafavi H, Sella G, Przeworski M. Measuring intolerance to mutation in human genetics. *Nat Genet*. 2019;51:772–776. <https://doi.org/10.1038/s41588-019-0383-1>.
- Galloway J, Cresko WA, Ralph P. A few stickleback suffice for the transport of alleles to new lakes. *G3 (Bethesda)*. 2020;10:505–514. <https://doi.org/10.1534/g3.119.400564>.
- Galtier N. Adaptive protein evolution in animals and the effective population size hypothesis. *PLoS Genet*. 2016;12:e1005774. <https://doi.org/10.1371/journal.pgen.1005774>.
- Gaynor RC, Gorjanc G, Hickey JM. `AlphaSimR`: an R-package for breeding program simulations. *G3: Genes—Genomes—Genetics*. 2021;11. <https://doi.org/10.1093/g3journal/gkaa017>.
- Gillespie JH. The molecular clock may be an episodic clock. *Proc Natl Acad Sci USA*. 1984;81:8009–8013. <https://doi.org/10.1073/pnas.81.24.8009>.
- Gillespie JH. *The causes of molecular evolution*. Vol. 2. Oxford University Press; 1991.
- Gutenkunst RN, Hernandez RD, Williamson SH, Bustamante CD. Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genet*. 2009;5:e1000695. <https://doi.org/10.1371/journal.pgen.1000695>.
- Haag J, Jordan AI, Stamatakis A. `Pandora`: a tool to estimate dimensionality reduction stability of genotype data. *Bioinform Adv*. 2025;vbaf040. <https://doi.org/10.1093/bioadv/vbaf040>.
- Haller BC, Messer PW. `SLiM 3`: forward genetic simulations beyond the Wright–Fisher model. *Mol Biol Evol*. 2019;36:632–637. <https://doi.org/10.1093/molbev/msy228>.
- Hartfield M, Glémin S. Polygenic selection to a changing optimum under self-fertilisation. *PLoS Genet*. 2024;20:e1011312. <https://doi.org/10.1371/journal.pgen.1011312>.
- Hernandez RD. A flexible forward simulator for populations subject to selection and demography. *Bioinformatics*. 2008;24:2786–2787. <https://doi.org/10.1093/bioinformatics/btn522>.
- Hoskins RA *et al.* The release 6 reference sequence of the *Drosophila melanogaster* genome. *Genome Res*. 2015;25:445–458. <https://doi.org/10.1101/gr.185579.114>.
- Huang X, Struck TJ, Davey SW, Gutenkunst RN. 2023. `dadi-cli`: automated and distributed population genetic model inference from allele frequency spectra [preprint]. bioRxiv 545182. <https://doi.org/10.1101/2023.06.15.545182>.
- Huber CD, Durvasula A, Hancock AM, Lohmueller KE. Gene expression drives the evolution of dominance. *Nat Commun*. 2018;9:2750.
- Huber CD, Kim BY, Marsden CD, Lohmueller KE. Determining the factors driving selective effects of new nonsynonymous mutations. *Proc Natl Acad Sci USA*. 2017;114:4465–4470. <https://doi.org/10.1073/pnas.1619508114>.
- Hudson RR. Testing the constant-rate neutral allele model with protein sequence data. *Evolution*. 1983;203–217. <https://doi.org/10.2307/2408186>.
- Hudson RR. Gene genealogies and the coalescent process. *Oxf Surv Evol Biol*. 1990;7:44.
- Hudson RR, Bailey K, Skarecky D, Kwiatowski J, Ayala FJ. Evidence for positive selection in the superoxide dismutase (*sod*) region of *Drosophila melanogaster*. *Genetics*. 1994;136:1329–1340. <https://doi.org/10.1093/genetics/136.4.1329>.
- Hudson RR, Boos DD, Kaplan NL. A statistical test for detecting geographic subdivision. *Mol Biol Evol*. 1992;9:138–151. <https://doi.org/10.1093/oxfordjournals.molbev.a040703>.

- Hudson RR, Kaplan NL. Deleterious background selection with recombination. *Genetics*. 1995;141:1605–1617. <https://doi.org/10.1093/genetics/141.4.1605>.
- Hunt SE *et al.* Ensembl variation resources. *Database*. 2018:2018: bay119. <https://doi.org/10.1093/database/bay119>.
- International HapMap Consortium. A second generation human haplotype map of over 3.1 million SNPs. *Nature*. 2007;449:851. <https://doi.org/10.1038/nature06258>.
- Johri P *et al.* The impact of purifying and background selection on the inference of population history: problems and prospects. *Mol Biol Evol*. 2021;38:2986–3003. <https://doi.org/10.1093/molbev/msab050>.
- Johri P, Charlesworth B, Jensen JD. Toward an evolutionarily appropriate null model: jointly inferring demography and purifying selection. *Genetics*. 2020;215:173–192. <https://doi.org/10.1534/genetics.119.303002>.
- Katzman S *et al.* Human genome ultraconserved elements are ultraselected. *Science*. 2007;317:915–915. <https://doi.org/10.1126/science.1142430>.
- Kern AD, Hahn MW. The neutral theory in light of natural selection. *Mol Biol Evol*. 2018;35:1366–1371. <https://doi.org/10.1093/molbev/msy092>.
- Kern AD, Schrider DR. Discoal: flexible coalescent simulations with selection. *Bioinformatics*. 2016;32:3839–3841. <https://doi.org/10.1093/bioinformatics/btw556>.
- Kern AD, Schrider DR. diploS/HIC: an updated approach to classifying selective sweeps. *G3 (Bethesda)*. 2018;8:1959–1970. <https://doi.org/10.1534/g3.118.200262>.
- Kim BY, Huber CD, Lohmueller KE. Inference of the distribution of selection coefficients for new nonsynonymous mutations using large samples. *Genetics*. 2017;206:345–361. <https://doi.org/10.1534/genetics.116.197145>.
- Kimura M. *The neutral theory of molecular evolution*. Cambridge University Press; 1985.
- Kingman JF. On the genealogy of large populations. *J Appl Probab*. 1982;19:27–43. <https://doi.org/10.2307/3213548>.
- Kyriazis CC, Lohmueller KE. Constraining models of dominance for nonsynonymous mutations in the human genome. *PLoS Genet*. 2024;20:e1011198. <https://doi.org/10.1371/journal.pgen.1011198>.
- Kyriazis CC, Robinson JA, Lohmueller KE. Using computational simulations to model deleterious variation and genetic load in natural populations. *Am Nat*. 2023;202:737–752. <https://doi.org/10.1086/726736>.
- Lauterbur ME *et al.* Expanding the stdpopsim species catalog, and lessons learned for realistic genome simulations. *Elife*. 2023;12: RP84874. <https://doi.org/10.7554/eLife.84874>.
- Lewontin RC, Krakauer J. Distribution of gene frequency as a test of the theory of the selective neutrality of polymorphisms. *Genetics*. 1973;74:175–195. <https://doi.org/10.1093/genetics/74.1.175>.
- Li H, Durbin R. Inference of human population history from individual whole-genome sequences. *Nature*. 2011;475:493. <https://doi.org/10.1038/nature10231>.
- Lin K, Li H, Schlotterer C, Futschik A. Distinguishing positive selection from neutral evolution: boosting the performance of summary statistics. *Genetics*. 2011;187:229–244. <https://doi.org/10.1534/genetics.110.122614>.
- Liu X, Fu Y-X. Stairway plot 2: demographic history inference with folded SNP frequency spectra. *Genome Biol*. 2020;21:280. <https://doi.org/10.1186/s13059-020-02196-9>.
- MacLeod IM, Hayes BJ, Goddard ME. The effects of demography and long-term selection on the accuracy of genomic prediction with sequence data. *Genetics*. 2014;198:1671–1684. <https://doi.org/10.1534/genetics.114.168344>.
- Marsh JI, Johri P. Biases in arg-based inference of historical population size in populations experiencing selection. *Mol Biol Evol*. 2024;41: msae118. <https://doi.org/10.1093/molbev/msae118>.
- McLean C, Bejerano G. Dispensability of mammalian DNA. *Genome Res*. 2008;18:1743–1751. <https://doi.org/10.1101/gr.080184.108>.
- Morin PA *et al.* Reference genome and demographic history of the most endangered marine mammal, the vaquita. *Mol Ecol Resour*. 2021;21:1008–1020. <https://doi.org/10.1111/men.v21.4>.
- Mughal MR, DeGiorgio M. Localizing and classifying adaptive targets with trend filtered regression. *Mol Biol Evol*. 2019;36:252–270. <https://doi.org/10.1093/molbev/msy205>.
- Mukai T, Chigusa SI, Mettler L, Crow JF. Mutation rate and dominance of genes affecting viability in *Drosophila melanogaster*. *Genetics*. 1972;72:335–355. <https://doi.org/10.1093/genetics/72.2.335>.
- Nielsen R *et al.* Genomic scans for selective sweeps using SNP data. *Genome Res*. 2005;15:1566–1575. <https://doi.org/10.1101/gr.4252305>.
- Ohta T. Slightly deleterious mutant substitutions in evolution. *Nature*. 1973;246:96–98. <https://doi.org/10.1038/246096a0>.
- Peischl S, Dupanloup I, Kirkpatrick M, Excoffier L. On the accumulation of deleterious mutations during range expansions. *Mol Ecol*. 2013;22:5972–5982. <https://doi.org/10.1111/mec.2013.22.issue-24>.
- Przeworski M. The signature of positive selection at randomly chosen loci. *Genetics*. 2002;160:1179–1189. <https://doi.org/10.1093/genetics/160.3.1179>.
- Przeworski M, Coop G, Wall JD. The signature of positive selection on standing genetic variation. *Evolution*. 2005;59:2312–2323. <https://doi.org/10.1554/05-273.1>.
- Ragsdale AP, Coffman AJ, Hsieh P, Struck TJ, Gutenkunst RN. Triallelic population genomics for inferring correlated fitness effects of same site nonsynonymous mutations. *Genetics*. 2016;203: 513–523. <https://doi.org/10.1534/genetics.115.184812>.
- Ragsdale AP, Gravel S. Models of archaic admixture and recent history from two-locus statistics. *PLoS Genet*. 2019;15:e1008204. <https://doi.org/10.1371/journal.pgen.1008204>.
- Ragsdale AP, Nelson D, Gravel S, Kelleher J. Lessons learned from bugs in models of human history. *Am J Hum Genet*. 2020;107:583–588. <https://doi.org/10.1016/j.ajhg.2020.08.017>.
- Ralph P, Thornton K, Kelleher J. Efficiently summarizing relationships in large samples: a general duality between statistics of genealogies and genomes. *Genetics*. 2020;215:779–797. <https://doi.org/10.1534/genetics.120.303253>.
- Robinson JA *et al.* The critically endangered vaquita is not doomed to extinction by inbreeding depression. *Science*. 2022;376:635–639. <https://doi.org/10.1126/science.abm1742>.
- Rodrigues MF, Kern AD, Ralph PL. Shared evolutionary processes shape landscapes of genomic variation in the great apes. *Genetics*. 2024;226:iyae006. <https://doi.org/10.1093/genetics/iyae006>.
- Sabeti PC *et al.* Detecting recent positive selection in the human genome from haplotype structure. *Nature*. 2002;419:832–837. <https://doi.org/10.1038/nature01140>.
- Sanchez T, Cury J, Charpiat G, Jay F. Deep learning for population size history inference: design, comparison and combination with approximate Bayesian computation. *Mol Ecol Resour*. 2021;21: 2645–2660. <https://doi.org/10.1111/men.v21.8>.
- Santiago E *et al.* Recent demographic history inferred by high-resolution analysis of linkage disequilibrium. *Mol Biol Evol*. 2020;37: 3642–3653. <https://doi.org/10.1093/molbev/msaa169>.
- Schiffels S, Wang K. MSMC and MSMC2: the multiple sequentially Markovian coalescent. Springer; 2020. p. 147–166.
- Schrider DR. Background selection does not mimic the patterns of genetic diversity produced by selective sweeps. *Genetics*. 2020;216: 499–519. <https://doi.org/10.1534/genetics.120.303469>.
- Schrider DR, Kern AD. Soft sweeps are the dominant mode of adaptation in the human genome. *Mol Biol Evol*. 2017;34:1863–1877. <https://doi.org/10.1093/molbev/msx154>.
- Schrider DR, Mendes FK, Hahn MW, Kern AD. Soft shoulders ahead: spurious signatures of soft and partial selective sweeps result from linked hard sweeps. *Genetics*. 2015;200:267–284. <https://doi.org/10.1534/genetics.115.174912>.
- Schrider DR, Shanku AG, Kern AD. Effects of linked selective sweeps on demographic inference and model selection. *Genetics*. 2016;204: 1207–1223. <https://doi.org/10.1534/genetics.116.190223>.

- Schweiger R, Durbin R. Ultrafast genome-wide inference of pairwise coalescence times. *Genome Res.* 2023;33:1023–1031. <https://doi.org/10.1101/gr.277665.123>.
- Sheehan S, Song YS. Deep learning for population genetic inference. *PLoS Comput Biol.* 2016;12:e1004845. <https://doi.org/10.1371/journal.pcbi.1004845>.
- Siepel A *et al.* Evolutionarily conserved elements in vertebrate, insect, worm, and yeast genomes. *Genome Res.* 2005;15:1034–1050. <https://doi.org/10.1101/gr.3715005>.
- Simons YB, Turchin MC, Pritchard JK, Sella G. The deleterious mutation load is insensitive to recent population history. *Nat Genet.* 2014;46:220–224. <https://doi.org/10.1038/ng.2896>.
- Simonsen KL, Churchill GA, Aquadro CF. Properties of statistical tests of neutrality for DNA polymorphism data. *Genetics.* 1995;141:413–429. <https://doi.org/10.1093/genetics/141.1.413>.
- Smith CC, Kern AD. dispersenn2: a neural network for estimating dispersal distance from georeferenced polymorphism data. *BMC Bioinformatics.* 2023;24:385. <https://doi.org/10.1186/s12859-023-05522-7>.
- Speidel L *et al.* Inferring population histories for ancient genomes using genome-wide genealogies. *Mol Biol Evol.* 2021;38:3497–3511. <https://doi.org/10.1093/molbev/msab174>.
- Stephan W. Genetic hitchhiking versus background selection: the controversy and its implications. *Philos Trans R Soc B Biol Sci.* 2010;365:1245–1253. <https://doi.org/10.1098/rstb.2009.0278>.
- Tajima F. Statistical method for testing the neutral mutation hypothesis by DNA polymorphism. *Genetics.* 1989;123:585–595. <https://doi.org/10.1093/genetics/123.3.585>.
- Tataru P, Bataillon T. polydf: inferring the distribution of fitness effects and properties of beneficial mutations from polymorphism data. *Stat Popul Genomics.* 2020;125–146. <https://doi.org/10.1007/978-1-0716-0199-0>.
- Temple SD, Waples RK, Browning SR. Modeling recent positive selection using identity-by-descent segments. *Am J Hum Genet.* 2024;111:2510–2529. <https://doi.org/10.1016/j.ajhg.2024.08.023>.
- Terhorst J, Kamm JA, Song YS. Robust and scalable inference of population history from hundreds of unphased whole genomes. *Nat Genet.* 2017;49:303. <https://doi.org/10.1038/ng.3748>.
- Thornton KR. A C++ template library for efficient forward-time population genetic simulation of large populations. *Genetics.* 2014;198:157–166. <https://doi.org/10.1534/genetics.114.165019>.
- Thornton KR. Polygenic adaptation to an environmental shift: temporal dynamics of variation under Gaussian stabilizing selection and additive effects on a single trait. *Genetics.* 2019;213:1513–1530. <https://doi.org/10.1534/genetics.119.302662>.
- Torres R, Stetter MG, Hernandez RD, Ross-Ibarra J. The temporal dynamics of background selection in nonequilibrium populations. *Genetics.* 2020;214:1019–1030. <https://doi.org/10.1534/genetics.119.302892>.
- Uricchio LH, Hernandez RD. Robust forward simulations of recurrent hitchhiking. *Genetics.* 2014;197:221–236. <https://doi.org/10.1534/genetics.113.156935>.
- Veeramah KR, Gutenkunst RN, Woerner AE, Watkins JC, Hammer MF. Evidence for increased levels of positive and negative selection on the x chromosome versus autosomes in humans. *Mol Biol Evol.* 2014;31:2267–2282. <https://doi.org/10.1093/molbev/msu166>.
- Wade EE, Kyriazis CC, Cavassim MIA, Lohmueller KE. Quantifying the fraction of new mutations that are recessive lethal. *Evolution.* 2023;77:1539–1549. <https://doi.org/10.1093/evolut/qpaa061>.
- Wang Z *et al.* Automatic inference of demographic parameters using generative adversarial networks. *Mol Ecol Resour.* 2021;21:2689–2705. <https://doi.org/10.1111/men.v21.8>.
- Williamson S, Orive ME. The genealogy of a sequence subject to purifying selection at multiple sites. *Mol Biol Evol.* 2002;19:1376–1384. <https://doi.org/10.1093/oxfordjournals.molbev.a004199>.
- Wright S. The distribution of gene frequencies in populations. *Proc Natl Acad Sci USA.* 1937;23:307–320. <https://doi.org/10.1073/pnas.23.6.307>.
- Yang M *et al.* Trustgwas: a full-process workflow for encrypted GWAS using multi-key homomorphic encryption and pseudorandom number perturbation. *Cell Syst.* 2022;13:752–767. <https://doi.org/10.1016/j.cels.2022.08.001>.
- Zhang X *et al.* Maladapt reveals novel targets of adaptive introgression from Neanderthals and Denisovans in worldwide human populations. *Mol Biol Evol.* 2023;40:msad001. <https://doi.org/10.1093/molbev/msad001>.
- Zhang X, Kim B, Lohmueller KE, Huerta-Sánchez E. The impact of recessive deleterious variation on signals of adaptive introgression in human populations. *Genetics.* 2020;215:799–812. <https://doi.org/10.1534/genetics.120.303081>.
- Zhen Y, Huber CD, Davies RW, Lohmueller KE. Greater strength of selection and higher proportion of beneficial amino acid changing mutations in humans compared with mice and *Drosophila melanogaster*. *Genome Res.* 2021;31:110–120. <https://doi.org/10.1101/gr.256636.119>.