

SOFTWARE

Open Access



# Legofit: estimating population history from genetic data

Alan R. Rogers

## Abstract

**Background:** Our current understanding of archaic admixture in humans relies on statistical methods with large biases, whose magnitudes depend on the sizes and separation times of ancestral populations. To avoid these biases, it is necessary to estimate these parameters simultaneously with those describing admixture. Genetic estimates of population histories also confront problems of statistical identifiability: different models or different combinations of parameter values may fit the data equally well. To deal with this problem, we need methods of model selection and model averaging, which are lacking from most existing software.

**Results:** The Legofit software package allows simultaneous estimation of parameters describing admixture, and the sizes and separation times of ancestral populations. It includes facilities for data manipulation, estimation, analysis of residuals, model selection, and model averaging.

**Conclusions:** Legofit uses genetic data to study the history of a subdivided population. It is unaffected by recent history and can therefore focus on the deep history of population size, subdivision, and admixture. It outperforms several statistical methods that have been widely used to study population history and should be useful in any species for which DNA sequence data is available from several populations.

**Keywords:** Population history, Coalescent theory, Genetics, Evolution

## Background

Genetic data now play a prominent role in research on human prehistory. In less than a decade, we have learned that modern humans carry DNA from Neanderthal ancestors [1] and also from a previously unknown “Denisovan” population [2, 3]; we have learned that the European Neolithic was primarily a movement of peoples [4, 5], but that farmers and foragers then lived side by side, exchanging genes for thousands of years [6]; we have learned that Indo-Europeans arrived in Europe about 5000 years ago as invaders from the Pontic Steppes [7]; and we have learned that some populations carry DNA from “superarchaics,” which separated from other humans perhaps a million years ago [8, 9].

There are reasons, however, to be skeptical of these new findings. First, many of the statistics used to estimate archaic admixture have large biases. For example, Rogers and Bohlender ([10], Fig. 4) document biases in

one statistic that range from 50 to 600%, depending on the separation time of Neanderthals and Denisovans. Petr et al. [11] show that similar bias in another statistic underlies an apparent (but artifactual) decline in the frequency of Neanderthal DNA in Europe during the past 45,000 years. To avoid these biases, one must simultaneously estimate the parameters that underlie them.

In addition to bias, there are also problems of statistical identifiability, which arise when several models fit the data equally well. Identifiability problems can lead us to prefer incorrect models of history, and they can make confidence intervals unrealistically narrow. Consequently, it is likely that some of the recent findings summarized above are incorrect.

The Legofit package [12, 13] introduces methods that address these problems. It reduces bias by allowing simultaneous estimation of the parameters that introduce bias into competing estimators. It uses model selection and model averaging to cope with identifiability problems, and it uses residual analysis to diagnose misspecified models. This article will not attempt a comprehensive review of genetic methods for estimation of population history.

Correspondence: [rogers@anthro.utah.edu](mailto:rogers@anthro.utah.edu)  
Department of Anthropology, University of Utah, Gardner Commons, Salt Lake City, USA



Instead, it will describe Legofit and compare it against several methods that are widely used in the study of archaic admixture.

### Implementation

#### Nucleotide site patterns

Legofit works with the frequencies of *nucleotide site patterns* [14, 15], which are defined below. The first step in any analysis involves tabulating site pattern frequencies from data. Legofit provides tools that tabulate these frequencies from standard data formats and also from several forms of simulation output.

Site patterns are illustrated in Fig. 1. A nucleotide site exhibits the *yn* site pattern if random nucleotides drawn from populations *Y* and *N* carry the derived allele, but those drawn from other populations carry the ancestral allele. They represent the special case of the site frequency spectrum [16] in which the sample consists of one haploid genome per population.

Many different gene trees—even trees with different topologies—may contribute to any given site pattern. Nonetheless, let us begin with a particular gene tree, which is shown in Fig. 1. There we see a population network and, embedded within it, the gene tree (or gene genealogy) of one particular locus (nucleotide site). A mutation on the red branch would generate *yn*, whereas one on the blue branch would generate *ynd*. Mutations elsewhere would generate other site patterns. Let  $B_i$  represent the length in generations of the branch generating site pattern *i*. For example,  $B_{yn}$  is the length of the red branch

in Fig. 1 and  $B_{ynd}$  is the length of the blue branch. The gene tree will vary from locus to locus, and in any given gene tree many of these lengths will be zero. For example,  $B_{xy} = 0$  in Fig. 1, because no single mutation on that gene tree could generate site pattern *xy*.

At a particular locus, and conditional on  $B_i$ , the number of mutations on the branch generating pattern *i* is Poisson with mean  $uB_i$ , where  $u$  is the mutation rate per nucleotide site per generation. We use the model of infinite sites [17], which assumes that  $u$  is small enough that we can ignore the possibility of multiple mutations on a given branch. To this standard of approximation, the unconditional probability of site pattern *i* on a random gene tree is  $uE[B_i]$ , where the expectation is with respect to the coalescent process constrained by the network of populations.

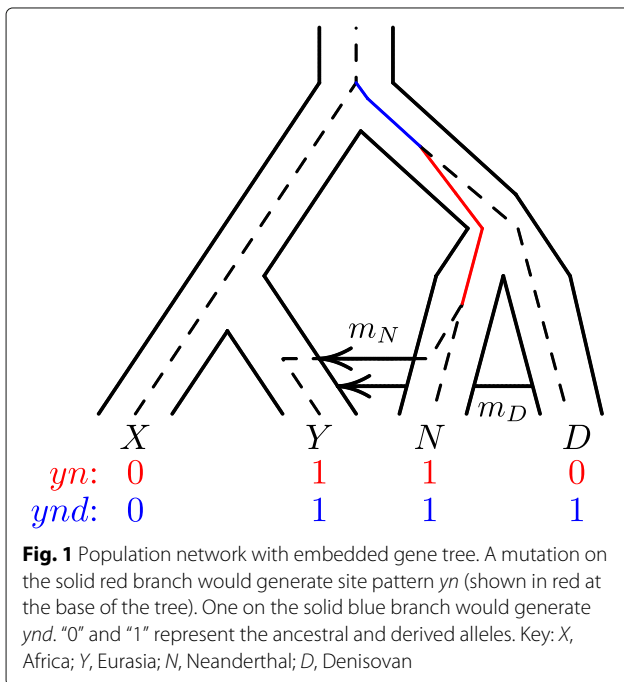
Let  $I_i$  represent the count of site pattern *i* across all sequenced nucleotide positions. Its expected value is  $E[I_i] = uLE[B_i]$ , where  $L$  is the number of nucleotide positions in the sequence. The probability that a particular polymorphic site exhibits pattern *i* is

$$P_i = \frac{E[B_i]}{\sum_{j \in \Omega} E[B_j]} \tag{1}$$

where  $\Omega$  is the set of site patterns under study.

In previous publications [10, 18] we and others have derived analytical expressions for  $E[B_i]$  under particular models of history. This analytical approach becomes difficult as models grow in complexity. Legofit relies instead on computer simulations, which make it feasible to deal with complex models of history. In each iteration of the simulation, the coalescent algorithm builds a gene genealogy analogous to the one in Fig. 1. From this genealogy, legofit<sup>1</sup> calculates branch lengths ( $B_i$ ). It estimates  $E[B_i]$  as the average of  $B_i$  across simulation replicates. Equation 1 then estimates  $P_i$ .

This approach simulates branch lengths but not mutations, and the simulations can be done in parallel. For a given level of accuracy, it is orders of magnitude faster than programs that simulate both mutation and recombination, as shown in the Additional file 1. This speed makes it possible to deal with the entire suite of site patterns and with complex models involving tens of populations. Nonetheless, this is still a computationally intensive approach. In a recent analysis [19], we studied nine different models. This took 10 days to do but would have taken 12 years without parallel processing. This 440-fold speed-up was possible because the calculations were parallelized not only across cores on each compute node, but also across nodes on the cluster at our local Center for High-Performance Computing. The legofit program



<sup>1</sup>We use lower case "legofit" to refer to the estimation program within the (capitalized) "Legofit" package.

parallelizes automatically across cores. Section 4 of the Additional file 1 describes methods for parallel processing on a cluster.

To validate our numerical approach to estimating probabilities, we compared it with theoretical results in models for which analytical theory is feasible [10]. We can also validate by comparing the expected values generated by our method to data simulated in other ways. This is done in Fig. 2, which shows that all three simulators generate distributions of site pattern frequencies that are centered around the expected values estimated by legofit. This verifies the reliability of our approach.

### Models of history

A model of population history is specified in a file whose name ends with “.lgo.” This file specifies the population network and the location of genetic samples within it. It uses a flexible syntax to describe population histories of arbitrary complexity. Populations can separate, combine, exchange migrants, and change in size. Changes in population size occur in discrete steps, and episodes of gene flow are modeled as discrete events, but there is no limit on the number of steps or episodes of gene flow. A model with  $K$  samples generates  $2^K - 2$  site patterns. For example, 10 samples would generate 1022 site patterns, which would provide a rich basis for estimating parameters.

Parameters fall into three categories: (1) *free* parameters are estimated by legofit; (2) *fixed* parameters have values that do not change; and (3) *constrained* parameters are specified as known functions of one or more other parameters. Constrained parameters model relationships among variables. We use them below to reexpress free variables in terms of principal components.

### Tabulating site patterns from data

The first stage of analysis involves tabulating site patterns from DNA sequence data. These data need not be

phased, but they should be free of ascertainment bias. In the discussion above, I assumed that one haploid genome is sampled from each population. Real samples are larger, and a given nucleotide site may contribute to several site patterns. The contribution to a given site pattern is the probability that a sub-sample, consisting of one haploid genome drawn at random from the larger sample of each population, would exhibit this site pattern. For example, consider a model with three populations,  $X$ ,  $Y$ , and  $N$ , and let  $p_{iX}$ ,  $p_{iY}$ , and  $p_{iN}$  represent derived allele frequencies at the  $i$ th polymorphic site in the samples from these populations. Then site pattern  $xy$  occurs at site  $i$  with probability  $z_i = p_{iX}p_{iY}(1 - p_{iN})$  ([1], p. S131). Aggregating over sites,  $I_{xy} = \sum_i z_i$  summarizes the information in the data about this site pattern. In general, for the  $j$ th site pattern, the analogous summary is  $I_j$ . In this formulation  $I_j$  is no longer a count. It is the expected count in a random subsample of the full sample.

The Legofit package includes programs for tabulating site patterns from data and from several publicly-available programs for coalescent simulation: *ms* [20], *msprime* [21], and *scrm* [22].


### Estimation

Legofit estimates parameters by maximizing the composite likelihood,

$$L(\theta) = \prod_{j \in \Omega} P_j^{I_j}(\theta) \tag{2}$$

where  $P_j$  is as given in Eq. 1,  $\Omega$  is the set of site patterns under study, and  $\theta$  is a vector of free parameters. This is not the full likelihood, because it ignores linkage disequilibrium and treats nucleotide sites as though they were independent.

Legofit uses a numerical algorithm—differential evolution (DE, [23])—to maximize  $L$ . DE maintains a swarm of points, which are initially distributed widely across the



**Fig. 2** Deviation from expected values in 50 data sets generated by each of three simulation programs: *ms* [20], *msprime* [21], and *scrm* [22]. All simulations assume the same model of history, which is illustrated in Fig. 1 and described fully in the Additional file 1. Expected values were calculated with legosim. Blue circles show 50 simulated data sets

parameter space. In each generation, these points mutate and recombine to form offspring, which then undergo selection to form the next generation. The objective functions of the points are evaluated in parallel, in separate threads of execution. This process involves several stages, beginning with an initial stage in which the objective function is evaluated with modest precision and progressing to a final stage, which typically uses two million simulation replicates per function evaluation. This provides much more precision than a sample of two million polymorphic nucleotide sites, because we are simulating branch lengths only—not mutation or recombination. (See the Additional file 1 for details).

### Bootstrap confidence intervals

The Legofit package uses a bootstrap [24] to measure uncertainty. Because linked loci are not statistically independent, we cannot use an ordinary bootstrap. Instead, Legofit uses a moving-blocks bootstrap [25], which resamples blocks of nucleotides. By default, each block consists of 500 polymorphic nucleotide sites.

Bootstrap replicates approximate independent samples from the stochastic process that produced the original data. By applying legofit to many bootstrap replicates, we obtain an approximation of the sampling distribution of the estimates. This distribution is used to estimate confidence intervals.

Each bootstrap replicate is analyzed by a separate instance of the legofit program. These instances can operate in parallel, on separate nodes of a compute cluster. Legofit is thus parallel in two senses: within each node, legofit uses multiple threads to parallelize across the points maintained by the DE algorithm. It also uses multiple nodes to parallelize across bootstrap replicates.

### Model selection

The study of population history requires that we choose among complex, non-nested models. Better fits can usually be achieved with more complex models, but this improvement may be illusory—the consequence of fitting noise rather than signal. Overfitting, as this is called, can produce incorrect inferences about population history [26]. We may report evidence of gene flow or of bottlenecks in population size where no such inference is warranted. Reliable inference requires that we protect against overfitting. This is not possible with the genetic methods currently used to study archaic admixture.

In other statistical contexts, such problems might be addressed via tools such as Akaike's information criterion (AIC, [27]), or the Bayesian information criterion (BIC, [28]), which penalize complex models in a principled way. These tools, however, require access to the full likelihood function, which is never available for genome-scale data sets. Because of the size and complexity of

the nuclear genome, all statistical methods simplify the problem in some way. Legofit uses *composite likelihood*, which ignores genetic linkage and treats nucleotide sites as though they were statistically independent. This produces unbiased estimates but does not allow us to use AIC or BIC.

Legofit provides two methods of model selection: the *bootstrap estimate of predictive error* (bepe, [24, 29]), and a *composite likelihood information criterion* (clik, [30]).

### Bootstrap estimate of predictive error (bepe)

Bepe is analogous to cross-validation, but uses bootstrap replicates instead of partitions of the data. The first step in the process uses legofit to fit a given model to each bootstrap replicate. These runs report the predicted frequency of each nucleotide site pattern. Legofit's "bepe" program then calculates the mean squared difference between these bootstrap-predicted frequencies and those in the real data and applies a small bias correction. The resulting estimate of predictive error compares favorably with cross-validation ([24], sec. 17.6). It is convenient, because we need bootstraps anyway for confidence intervals.

### Composite likelihood information criterion (clik)

Clik generalizes Akaike's information criterion (AIC, [27]) to the case of composite likelihood. Varin and Vidoni ([30], p. 523) define an information criterion that is the negative of

$$\text{clik} = -\ln L(\theta) - \text{tr}\{HC\}, \quad (3)$$

I have reversed the sign so that we can select models by minimizing (rather than maximizing) clik. In this expression,  $L$  is composite likelihood (Eq. 2),  $\theta$  is the vector of parameters,  $C$  is a matrix whose  $ij$ th entry is the sampling covariance between the  $i$ th and  $j$ th parameters, and  $H$  is the expectation of the negative of the Hessian matrix, and "tr" represents the matrix trace.

I estimate  $C$  from covariances across bootstrap or simulation replicates.  $H$  is a matrix of expectations of second-order partial derivatives of  $\ln L$  with respect to pairs of parameters. Rather than taking these expectations, I evaluate the derivatives at the maximum composite likelihood estimate,  $\hat{\theta}$  [31]. Within a small neighborhood near  $\hat{\theta}$ ,  $\ln L$  can be approximated by a quadratic surface,

$$\ln L(\theta) \approx \alpha + \sum_i \beta_i(\theta_i - \hat{\theta}_i) + \sum_{i \leq j} \gamma_{ij}(\theta_i - \hat{\theta}_i)(\theta_j - \hat{\theta}_j), \quad (4)$$

where  $\alpha$  is the  $Y$  intercept, and  $\beta_i$  and  $\gamma_{ij}$  are regression coefficients.

I estimate  $\alpha$ ,  $\beta_i$ , and  $\gamma_{ij}$  by ordinary least squares, using points in the neighborhood of the estimate,  $\hat{\theta}$ . Then  $H$  is assembled using the second-order derivatives of  $\ln L$ , as implied by Eq. 4. Finally,  $C$  and  $H$  are used with Eq. 3 to calculate clik.



### Bootstrap model averaging (booma)

Below, we will consider three models whose bepe values are  $2.17 \times 10^{-7}$ ,  $5.54 \times 10^{-7}$ , and  $6.17 \times 10^{-5}$ . The first model has the smallest value and is therefore preferred. But the other values are also small. Are we justified in ignoring them? To answer this question, let us consider the problem of model averaging.

When no model is clearly superior, it is better to average across several than to choose just one [32]. Otherwise, confidence intervals are misleadingly narrow because they ignore uncertainty about the model itself. In model averaging, individual models are assigned weights as discussed below. Parameters are estimated as the weighted average of estimates from individual models. Most authors rely on information criteria to provide the weights [33]. One could use clic in this way, but I prefer *bootstrap model averaging* [32], which works with either bepe or clic.

This method is implemented by the Legofit program “booma.” Some model selection criterion (bepe or clic) is calculated separately for the real data and for each bootstrap replicate. (To calculate bepe for a bootstrap replicate, we pretend that the replicate is real data and the real data are a bootstrap replicate.) If there are 50 bootstrap replicates, this process gives us 51 values of the model selection criterion for each model. For each of these 51 cases, booma asks which model “wins,” i.e., which has the lowest value of the criterion. The weight of the  $i$ th model is the fraction of cases in which it is the winning model.

Using these weights, booma averages across models to obtain a model-averaged estimate of each parameter. If a parameter is present in only a subset of the models, the weights are re-normalized so that they sum to unity across this subset. This averaging is applied not only to the real data but also to each bootstrap replicate. This allows us to estimate confidence intervals for model-averaged estimators.

If one model is clearly superior, its weight will be unity and those of the other models will be zero. This provides a simple criterion for choosing one model over its alternatives. For the three models mentioned at the top of this section, the weights were 1, 0, and 0. This implies that the differences among the bepe values are large compared to those expected in repeated sampling from the stochastic process that generated the original data. We are therefore justified in rejecting all models but the first. This analysis is described in more detail below.

### Identifiability and principal components

Figure 3 illustrates a problem of statistical identifiability, which arises frequently not only with Legofit, but with all methods that estimate complex population histories. Each panel in the figure is a bivariate scatterplot comparing two parameters. Each point indicates the estimated values of

the two parameters in one simulation replicate. In several panels, the points fall along straight lines, indicating that the parameters are tightly correlated. These associations represent ridges in the composite likelihood surface and imply that our statistical problem has fewer dimensions than parameters. This does not lead to incorrect inferences, but it does broaden the confidence intervals of the parameters involved.

These problems can be ameliorated by reducing the dimension of the parameter space. The Legofit package includes *pclgo*, a program that calculates principal components from the bootstrap replicates and then uses these to re-express the free variables in terms of principal components. Predictive error (as measured by bepe) can be improved by excluding principal components with small eigenvalues. This usually tightens confidence intervals.

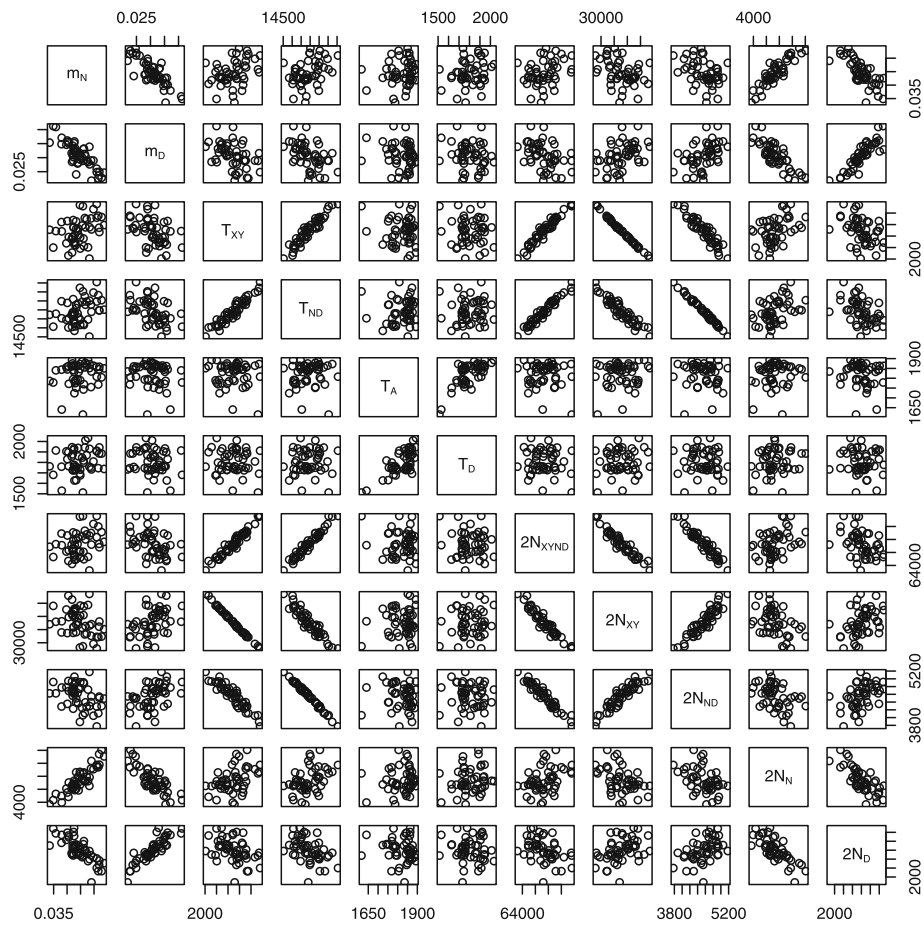
By default, *pclgo* merely re-expresses the free variables in terms of the principal components, and there is no reduction in dimension. To reduce dimensionality, the user must specify a tolerance criterion. The command `pclgo -tol 0.001` would include only those components that explain at least a fraction 0.001 of the variance. Different choices of this tolerance criterion constitute different models, and we can choose among them using bepe or clic, together with booma.

### Results

Rogers and Bohlender [10] document pronounced biases in the statistics that underlie our current understanding of archaic admixture. These biases are profound if there are multiple sources of admixture. To check for such bias in legofit, I simulate data under the model in Fig. 1, which allows gene flow into Eurasia ( $Y$ ) not only from Neanderthals ( $N$ ), but also from Denisovans ( $D$ ). Details of this model and of all the analyses below can be found in the Additional file 1. Here, I summarize results.

Figure 4 shows the true parameter values (red crosses) and sampling distributions (blue circles) estimated using legofit from 50 independent simulation replicates. I used *pclgo* to reduce dimensionality. This involves excluding dimensions that explain less than some arbitrarily-chosen fraction of the variance. I considered three models: one in terms of the original variables (without using *pclgo*), one using principal components with no reduction of dimension, and one excluding components that explain less than a fraction 0.001 of the variance. The weights of these three models are 0, 0.42, and 0.58 using bepe and 0, 0.12, and 0.88 using clic. Thus, *pclgo* seems to improve estimates, especially when some principal components are excluded. Figure 4 shows the bepe version of the model-averaged estimates.

All of the sampling distributions enclose the true parameter values, and several are reassuringly narrow. Nonetheless, some bias is evident in the distributions



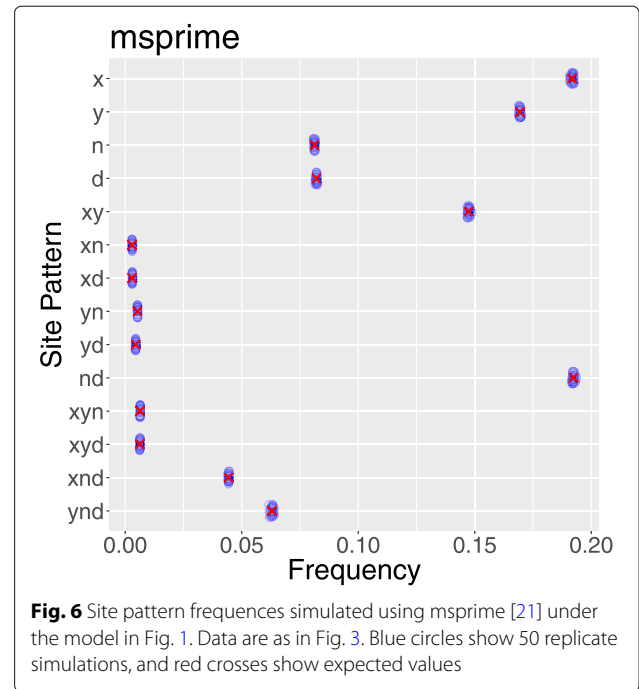
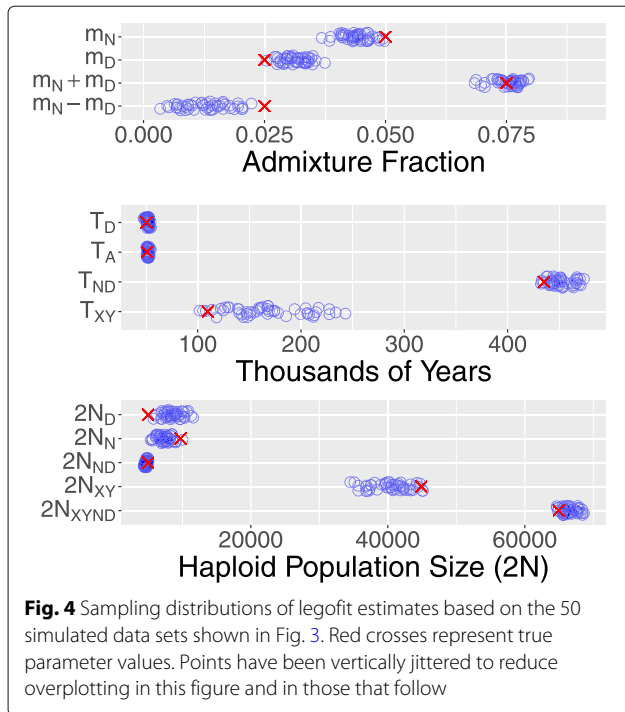
**Fig. 3** Associations between pairs of parameter estimates in 50 data sets simulated with msprime [21] under the model in Fig. 1. Key:  $m_N$ , fraction of admixture from  $N$  into  $Y$ ;  $m_D$ , fraction of admixture from  $D$  into  $Y$ ;  $T_{XY}$ , separation time of  $X$  and  $Y$ ;  $T_{ND}$ , separation time of  $N$  and  $D$ ;  $T_A$ , age of fossil genome from population  $N$ ;  $T_D$ , age of fossil from  $D$ ;  $N_{XYND}$ , size of ancestral population;  $N_{XY}$ , size of population ancestral to  $X$  and  $Y$ ;  $N_{ND}$ , size of population ancestral to  $N$  and  $D$ ;  $N_N$ , size of population  $N$ ;  $N_D$ , size of population  $N$ . The separation time,  $T_{XYND}$ , of  $XY$  and  $ND$  was fixed exogenously to calibrate the molecular clock

of Neanderthal admixture ( $m_N$ ) and Denisovan admixture ( $m_D$ ). The mean estimates of these parameters are closer together than are the true parameter values. This is because Neanderthals and Denisovans are sister populations, and it is hard to tell them apart. We get a better estimate of total archaic admixture,  $m_N + m_D$ , than of the difference,  $m_N - m_D$ .

For comparison with legofit’s estimates of the admixture fraction, Fig. 5 shows the behavior of three previously-published estimators [2, 3] that have been used to study archaic admixture in humans. Nea and den work by comparing the frequencies with which derived alleles are shared by pairs of samples from different populations. Nea has also been called  $R_{Neanderthal}$  [2]. Rogers and Bohlender [10] show that these estimators have large biases, especially when (as in the present model) a population receives gene flow from more than one source. Thus, it is no surprise that nea and den exhibit large biases in Fig. 5. Indeed,

the black triangles show that the observed bias is in good agreement with theoretical expectations.

Many studies have cited an estimate that about 6% of Papuan DNA derives from Denisovans. This result is due to Meyer et al. [3], who inferred it using TreeMix [34]. However, these authors suspected that the result was biased, because their analysis excluded Neanderthals ([3], supp. note 12). The TreeMix results in Fig. 5 should avoid this problem, because Neanderthals are included along with Denisovans and moderns from Africa and Eurasia. TreeMix was able to detect a signal of gene flow from Neanderthals into Eurasians. As the figure shows, however, its estimate of the admixture fraction was profoundly biased. TreeMix was unable to detect gene flow from Denisovans into Eurasians. This episode of gene flow did not appear in the output from any of the simulation replicates. Instead, TreeMix reported evidence of gene flow in various parts of the tree. These episodes of gene flow



were not consistent from replicate to replicate and did not exist in the simulation model.

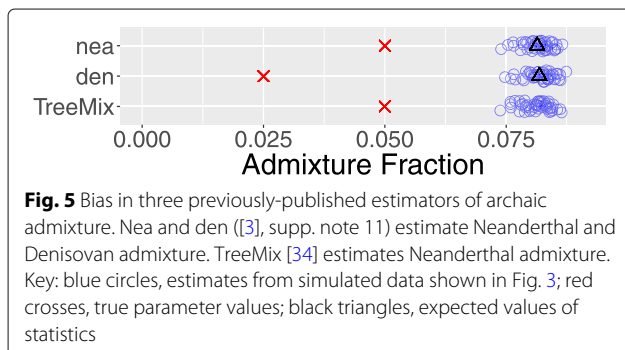
**A worked example**

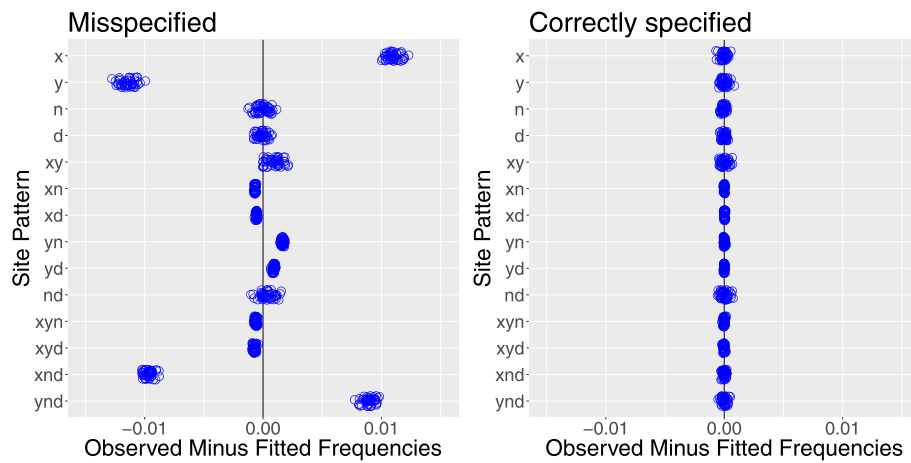
In Fig. 4, we had the advantage of working with the true model of history. This is never the case with real data. Let us therefore consider how the analysis might proceed if we did not know the true model in advance. We would start by examining site pattern frequencies, which are shown in Fig. 6. The most common patterns (apart from singletons) are *xy* and *nd*, reflecting the shared ancestry of populations *X* and *Y* and of *N* and *D*. Let us therefore fit a model with a tree of form  $((X, Y), (N, D))$ . This model is misspecified, because it omits gene flow. The residuals of this model are shown in Fig. 7 along with those of a correctly-specified model. The misspecified model generates many residuals that are far from zero, and these

discrepancies provide clues about what is wrong with the model. For example, note that the misspecified model has positive residuals for *yn* and *ynd* but a negative residual for *y*. This suggests that we should add  $N \rightarrow Y$  gene flow to the model, because such gene flow inflates the first two of these site patterns but deflates the third.

Table 1 compares the two models and shows that the one with  $N \rightarrow Y$  gene flow is unambiguously better than the one without gene flow. However, the residuals of this new model (not shown) still show discrepancies, which might lead us to consider adding  $D \rightarrow Y$  gene flow to the model. Table 2 shows that this third model is unambiguously better than the one with only one episode of gene flow. The residuals (right panel of Fig. 7) show that this model provides a good description of the data. In this example, the correct model was identifiable because the alternate models could not fully account for the pattern in the data.

There are also less tractable identifiability problems. Let us consider two. Figure 8 shows a model that is like that in the simulations (Fig. 1) but has an additional episode of gene flow from a “superarchaic” population (*S*) into Denisovans (*D*), as suggested by Prüfer et al [8]. When the superarchaic admixture fraction is zero, this model reduces to that used in our simulations. As expected, legofit’s estimate of this parameter was very close to zero in all simulation replicates, and all other parameters were also well estimated. Consequently, this model provides an excellent fit to the data, comparable to that in the right panel of Fig. 7. Nonetheless, I expected bepe and





**Fig. 7** Residuals from misspecified and correctly-specified models. Each circle represents one of the simulated data sets in Fig. 3. The misspecified model ignores the two episodes of gene flow seen in Fig. 1

clik to prefer the correct model because of its simplicity. Instead, bepe and clic gave appreciable weight to both models but preferred the more complex one, as shown in Table 3. This did not lead to incorrect inferences, because all parameters were well estimated.

Table 4 illustrates another identifiability problem. It compares the standard model (Fig. 1) with one in which the order of the two admixture events is reversed:  $D \rightarrow Y$  admixture precedes  $N \rightarrow Y$  admixture. This change has little effect on site pattern frequencies, and all parameters are well estimated. I expected bepe and clic to weight these models roughly equally. The table shows that they do give appreciable weight to both models but prefer the (incorrect) reversed model. In another experiment (not shown), using ms instead of msprime, bepe gave 94% of the weight to the true model. Bepe and clic both behave sensibly when dealing with models that are indistinguishable or nearly so. In such cases, they tend to give appreciable weight to several models. We cannot assume, however, that they will always prefer the correct model.

**Table 1** Booma weights for models with and without  $N \rightarrow Y$  gene flow

Weights		Model
bepe	clic	
0	0	No gene flow; full dimension
0	0	No gene flow; reduced dimension
0.04	0.5	$N \rightarrow Y$ gene flow; full dimension
0.96	0.5	$N \rightarrow Y$ gene flow; reduced dimension

All models re-express free variables in terms of principal components. Models with reduced dimension exclude principal components that explain less than a fraction 0.001 of the variance

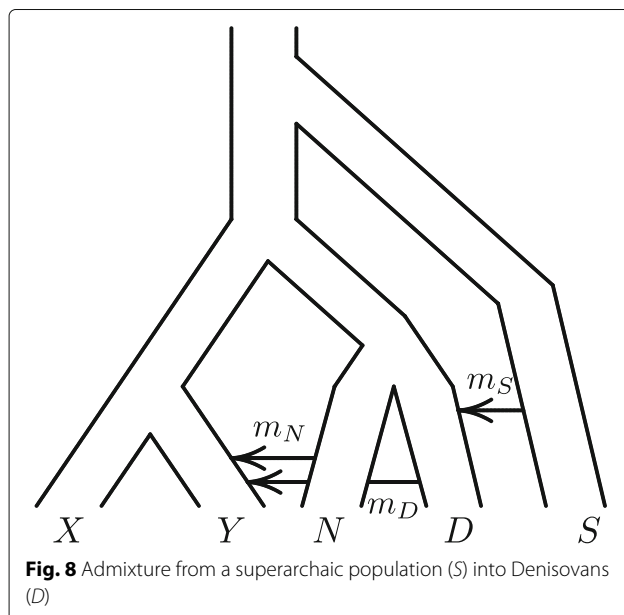
### Discussion

There are two reasons for studying site patterns rather than the full site frequency spectrum, the first of which involves statistical power at deep time scales. As we look backwards into the past, large samples coalesce rapidly to small collections of ancestors. For this reason, although large samples are essential for recent history, their value is limited in the distant past. Furthermore, the random-haploid samples used by legofit provide an advantage: they insulate the analysis from recent population history. If we had sampled several haploid genomes from population  $X$  in Fig. 1, then our model would need parameters describing changes in the size of  $X$  since its separation from  $Y$ . With legofit, these parameters aren't needed, because no coalescent events can occur until  $X$  and  $Y$  merge into their ancestral population. Thus, site pattern frequencies reduce the parameter count without losing much power at deep time scales. They allow us to study the deep history of multiple populations.

**Table 2** Booma weights for models with and without  $D \rightarrow Y$  gene flow

Weights		Model
bepe	clic	
0	0	No $D \rightarrow Y$ gene flow; full dimension
0	0	No $D \rightarrow Y$ gene flow; reduced dimension
0.42	0.12	$D \rightarrow Y$ gene flow; full dimension
0.58	0.88	$D \rightarrow Y$ gene flow; reduced dimension

All models include  $N \rightarrow Y$  gene flow and re-express free variables in terms of principal components. Models with reduced dimension exclude principal components that explain less than a fraction 0.001 of the variance



## Conclusions

The Legofit package provides computer programs for estimating population histories. It uses the frequencies of nucleotide site patterns to summarize genetic data. The package includes programs that tabulate these frequencies, calculate their expected values, and use them to estimate parameters describing population history. It includes facilities for model selection and model averaging. It uses principal components to reduce the complexity of high-dimensional models of history. Legofit outperforms several methods that have been widely used to study archaic admixture in humans and should be useful in any species for which DNA sequence data is available from several populations.

## Availability and requirements

*Project name* Legofit

*Project home page* <https://github.com/alanrogers/legofit>

*Operating system* Linux and macOS

**Table 3** Booma weights for models with and without superarchaic admixture

Weights		Model
bepe	clic	
0.24	0.04	No superarchaic admixture; full dimension
0.02	0.16	No superarchaic admixture; reduced dimension
0	0	Superarchaic admixture; full dimension
0.74	0.80	Superarchaic admixture; reduced dimension

All models include  $N \rightarrow Y$  and  $D \rightarrow Y$  gene flow and re-express free variables in terms of principal components. Models with reduced dimension exclude principal components that explain less than a fraction 0.001 of the variance

**Table 4** Booma weights for models with and without reversing the order of the two admixture events in Fig. 1

Weights		Model
bepe	clic	
0.18	0.02	True model; full dimension
0	0.22	True model; reduced dimension
0	0.02	Reversed model; full dimension
0.82	0.74	Reversed model; reduced dimension

All models include  $N \rightarrow Y$  and  $D \rightarrow Y$  gene flow and re-express free variables in terms of principal components. Models with reduced dimension exclude principal components that explain less than a fraction 0.001 of the variance

*Programming language* C and Python

*Requirements* pthreads and the Gnu Scientific Library

*License* Internet Systems Consortium License

*Any restrictions to use by non-academics* none

## Supplementary information

**Supplementary information** accompanies this paper at <https://doi.org/10.1186/s12859-019-3154-1>.

**Additional file 1:** Supplementary methods and results

## Abbreviations

AIC: Akieke's information criterion; bepe: Bootstrap estimate of predictive error; BIC: Bayesian information criterion booma: Bootstrap model averaging; clic: Composite likelihood information criterion; DE: Differential evolution

## Acknowledgements

I am grateful to Alan Achenbach, Kiela Gwin, Nathan Harris, Annie-Louise Holbrook, Mitchell Lokey, and Daniel Tabin, who have all used the software and provided feedback. Daniel Tabin helped write several programs within the package. Elizabeth Cashdan, Ilan Gronau, Timothy Webster provided useful comments on the text.

## Authors' contributions

Not applicable, because there is only one author. The author read and approved the final manuscript.

## Funding

This work was supported by NSF award BCS 1638840 and by the Center for High Performance Computing at the University of Utah.

## Availability of data and materials

The Legofit package is freely available at <https://github.com/alanrogers/legofit>. The pipelines and intermediate data files used in the analysis reported here are freely available at the Open Science Framework, <https://osf.io/j38eg>.

## Ethics approval and consent to participate

Not applicable, because all the data studied here were generated by computer simulation.

## Consent for publication

Not applicable, because this article includes no details, images, or videos relating to an individual person.

## Competing interests

The author declares that he has no competing interests.

Received: 25 February 2019 Accepted: 14 October 2019

Published online: 28 October 2019



## References

- Green RE, Krause J, Briggs AW, Maricic T, Stenzel U, Kircher M, Patterson N, Li H, Zhai W, Fritz M. H.-Y., Hansen NF, Durand EY, Malaspinas A.-S., Jensen JD, Marques-Bonet T, Alkan C, Prüfer K, Meyer M, Burbano HA, Good JM, Schultz R, Aximu-Petri A, Butthof A, Höber B, Höffner B, Siegemund M, Weihmann A, Nusbaum C, Lander ES, Russ C, Novod N, Affourtit J, Egholm M, Verna C, Rudan P, Brajkovic D, Kucan V., Gušić I, Doronichev VB, Golovanova LV, Lalueza-Fox C, de la Rasilla M, Fortea J, Rosas A, Schmitz RW, Johnson PLF, Eichler EE, Falush D, Birney E, Mullikin JC, Slatkin M, Nielsen R, Kelso J, Lachmann M, Reich D, Pääbo S. A draft sequence of the Neandertal genome. *Science*. 2010;328(5979):710–22. <https://doi.org/10.1126/science.1188021>.
- Reich D, Green RE, Kircher M, Krause J, Patterson N, Durand EY, Viola B, Briggs AW, Stenzel U, Johnson PLF, et al. Genetic history of an archaic hominin group from Denisova Cave in Siberia. *Nature*. 2010;468(7327):1053–60.
- Meyer M, Kircher M, Gansauge M.-T., Li H, Racimo F, Mallick S, Schraiber JG, Jay F, Prüfer K, de Filippo C, Sudmant PH, Alkan C, Fu Q, Do R, Rohland N, Tandon A, Siebauer M, Green RE, Bryc K, Briggs AW, Stenzel U, Dabney J, Shendure J, Kitzman J, Hammer MF, Shunkov MV, Dereviako AP, Patterson N, Andr as AM, Eichler EE, Slatkin M, Reich D, Kelso J, Pääbo S. A high-coverage genome sequence from an archaic Denisovan individual. *Science*. 2012;338(6104):222–6.
- Bollongino R, Nehlich O, Richards MP, Orschiedt J, Thomas MG, Sell C, Fajkošová Z, Powell A, Burger J. 2000 years of parallel societies in Stone Age central Europe. *Science*. 2013;342:479–81.
- Skoglund P, Malmström H, Raghavan M, Stor  J, Hall P, Willerslev E, Gilbert MTP, Götherström A, Jakobsson M. Origins and genetic legacy of Neolithic farmers and hunter-gatherers in Europe. *Science*. 2012;336(6080):466–9.
- Lipson M, Szécsényi-Nagy A, Mallick S, Pósa A, Stégm r B, Keerl V, Rohland N, Stewardson K, Ferry M, Michel M, Oppenheimer J, Broomandkoshbacht N, Harney E, Nordenfelt S, Llamas B, Guszt v Mende B, K hler K, Oross K, Bond r M, Marton T, Osz s A, Jakucs J, Paluch T, Horv th F, Csengeri P, Ko s J, Seb k K, Anders A, Raczky P, Regenye J, Barna JP, F bi n S, Serlegi G, Toldi Z, Gy ngyv r Nagy E, Dani J, Moln r E, P lfi G, M rk L, Melegh B, B nfai Z, Dombor czi L, Fern ndez-Eraso J, Antonio Mujika-Alustiza J, Alonso Fern ndez C, Jim nez Echevarr  J, Bollongino R, Orschiedt J, Schierhold K, Meller H, Cooper A, Burger J, B nffy E, Alt KW, Lalueza-Fox C, Haak W, Reich D. Parallel palaeogenomic transects reveal complex genetic history of early European farmers. *Nature*. 2017;551(7680):368–72. <https://doi.org/10.1038/nature24476>.
- Haak W, Lazaridis I, Patterson N, Rohland N, Mallick S, Llamas B, Brandt G, Nordenfelt S, Harney E, Stewardson K, Fu Q, Mittnik A, B nffy E, Economou C, Francken M, Friederich S, Pena RG, Hallgren F, Khartanovich V, Khokhlov A, Kunst M, Kuznetsov P, Meller H, Mochalov O, Moiseyev V, Nicklisch N, Pichler SL, Risch R, Rojo Guerra MA, Roth C, Sz cs nyi-Nagy A, Wahl J, Meyer M, Krause J, Brown D, Anthony D, Cooper A, Alt KW, Reich D. Massive migration from the steppe is a source for Indo-European languages in Europe. *Nature*. 2015;522(7555):207. <https://doi.org/10.1038/nature14317>.
- Pr ufer K, Racimo F, Patterson N, Jay F, Sankararaman S, Sawyer S, Heinze A, Renaud G, Sudmant PH, de Filippo C, Li H, Mallick S, Dannemann M, Fu Q, Kircher M, Kuhlwilm M, Lachmann M, Meyer M, Ongyerth M, Siebauer M, Theunert C, Tandon A, Moorjani P, Pickrell J, Mullikin JC, Vohr SH, Green RE, Hellmann I, Johnson PLF, Blanche H, Cann H, Kitzman JO, Shendure J, Eichler EE, Lein ES, Bakken TE, Golovanova LV, Doronichev VB, Shunkov MV, Dereviako AP, Viola B, Slatkin M, Reich D, Kelso J, P  bo S. The complete genome sequence of a Neanderthal from the Altai Mountains. *Nature*. 2014;505(7481):43–49. <https://doi.org/10.1038/nature12886>.
- Mendez FL, Watkins JC, Hammer MF. Global genetic variation at OAS1 provides evidence of archaic admixture in Melanesian populations. *Mol Biol Evol*. 2012;29(6):1513–20.
- Rogers AR, Bohlender RJ. Bias in estimators of archaic admixture. *Theor Popul Biol*. 2015;100:63–78. <https://doi.org/10.1016/j.tpb.2014.12.006>.
- Petr M, P  bo S, Kelso J, Vernot B. Limits of long-term selection against Neandertal introgression. *Proc Natl Acad Sci USA*. 2019;116(5):1639–44. <https://doi.org/10.1073/pnas.1814338116>.
- Rogers AR, Bohlender RJ, Huff CD. Early history of Neanderthals and Denisovans. *Proc Natl Acad Sci USA*. 2017;114(37):9859–63. <https://doi.org/10.1073/pnas.1706426114>.
- Rogers AR, Bohlender RJ, Huff CD. Reply to Mafessoni and Pr ufer: Inferences with and without singleton site patterns. *Proc Natl Acad Sci USA*. 2017;114(48):10258–60. <https://doi.org/10.1073/pnas.1717085114>.
- Yang Z, Rannala B. Molecular phylogenetics: Principles and practice. *Nat Rev Genet*. 2012;13(5):303.
- Eaton DA, Hipp AL, Gonz lez-Rodr guez A, Cavender-Bares J. Historical introgression among the American live oaks and the comparative nature of tests for introgression. *Evolution*. 2015;69(10):2587–601. <https://doi.org/10.1111/evo.12758>.
- Hudson RR. A new proof of the expected frequency spectrum under the standard neutral model. *PLoS*. 2015;10(1):0118087. <https://doi.org/10.1371/journal.pone.0118087>.
- Kimura M. The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutation. *Genetics*. 1969;61:893–903.
- Durand EY, Patterson N, Reich D, Slatkin M. Testing for ancient admixture between closely related populations. *Mol Biol Evol*. 2011;28(8):2239–52.
- Rogers AR, Harris NS, Achenbach AA. Neanderthal-Denisovan ancestors interbred with a distantly-related hominin. *bioRxiv*. 2019;657247. <https://doi.org/10.1101/657247>.
- Hudson RR. Generating samples under a Wright-Fisher neutral model of genetic variation. *Bioinformatics*. 2002;18:337–8.
- Kelleher J, Etheridge AM, McVean G. Efficient coalescent simulation and genealogical analysis for large sample sizes. *PLoS Comput Biol*. 2016;12(5):1–22. <https://doi.org/10.1371/journal.pcbi.1004842>.
- Staab PR, Zhu S, Metzler D, Lunter G. Scrm: Efficiently simulating long sequences using the approximated coalescent with recombination. *Bioinformatics*. 2015;31(10):1680–2.
- Price K, Storn RM, Lampinen JA. *Differential Evolution: A Practical Approach to Global Optimization*. Berlin: Springer; 2006.
- Efron B, Tibshirani RJ. *An Introduction to the Bootstrap*. New York: Chapman and Hall; 1993.
- Liu RY, Singh K. Moving blocks jackknife and bootstrap capture weak dependence. In: LePage R, Billard L, editors. *Exploring the "Limits" of the Bootstrap*. New York: Wiley; 1992. p. 225–48.
- Hawkins DM. The problem of overfitting. *J Chem Inf Comput Sci*. 2004;44(1):1–12.
- Akaike H. A new look at the statistical model identification. *IEEE Trans Autom Control*. 1974;19(6):716–23.
- Schwarz GE. Estimating the dimension of a model. *Ann Stat*. 1978;41(2):461–4.
- Efron B. Estimating the error rate of a prediction rule: Improvement on cross-validation. *J Am Stat Assoc*. 1983;78(382):316–31.
- Varin C, Vidoni P. A note on composite likelihood inference and model selection. *Biometrika*. 2005;92(3):519–28.
- Efron B, Hinkley DV. Assessing the accuracy of the maximum likelihood estimator: Observed versus expected Fisher information. *Biometrika*. 1978;65(3):457–82. <https://doi.org/10.2307/2335893>.
- Buckland ST, Burnham KP, Augustin NH. Model selection: an integral part of inference. *Biometrics*. 1997;53(2):603–18.
- Claeskens G, Hjort NL. *Model Selection and Model Averaging*. Cambridge: Cambridge University Press; 2008.
- Pickrell JK, Patterson N, Barbieri C, Berthold F, Gerlach L, G ldemann T, Kure B, Mpoloka SW, Nakagawa H, Naumann C, Lipson M, Loh P.-R., Lachance J, Mountain J, Bustamante CD, Berger B, Tishkoff SA, Henn BM, Stoneking M, Reich D, Pakendorf B. The genetic prehistory of southern Africa. *Nat Commun*. 2012;3:1143.

## Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.