

1 Estimation of site frequency spectra from
2 low-coverage sequencing data using stochastic EM
3 reduces overfitting, runtime, and memory usage

4 Malthe Sebro Rasmussen¹, Genís Garcia-Erill¹, Thorfinn Sand
5 Korneliussen², Carsten Wiuf³, and Anders Albrechtsen¹

6 ¹Department of Biology, University of Copenhagen

7 ²GLOBE Institute, University of Copenhagen

8 ³Department of Mathematical Sciences, University of Copenhagen

9 **Abstract**

10 The site frequency spectrum (SFS) is an important summary statistic
11 in population genetics used for inference on demographic history and
12 selection. However, estimation of the SFS from called genotypes intro-
13 duce bias when working with low-coverage sequencing data. Methods
14 exist for addressing this issue, but sometimes suffer from two problems.
15 First, they can have very high computational demands, to the point
16 that it may not be possible to run estimation for genome-scale data.
17 Second, existing methods are prone to overfitting, especially for multi-
18 dimensional SFS estimation. In this article, we present a stochastic
19 expectation-maximisation algorithm for inferring the SFS from NGS
20 data that addresses these challenges. We show that this algorithm
21 greatly reduces runtime and enables estimation with constant, trivial
22 RAM usage. Further, the algorithm reduces overfitting and thereby
23 improves downstream inference. An implementation is available at
24 github.com/malthesr/winsfs .

25 **1 Introduction**

26 The site frequency spectrum (SFS) is the joint distribution of allele frequencies
27 among one or more populations, and it serves as an important summary stat-
28 istic in population genetics. For instance, the SFS is sufficient for computing
29 nucleotide diversity [1], F_{ST} [2], and f -statistics [3]. Furthermore, the SFS
30 may be used for inferring demographic history [4–6] and selection [7–9].

31 When working with high-quality data, it is usually straightforward to
32 estimate the SFS from called genotypes. However, when genotype calls are
33 uncertain, standard methods lead to significant bias in the estimated SFS
34 [10], which propagates to downstream inference [11]. In particular, this
35 situation arises when working with next-generation sequencing (NGS) data
36 at low coverage and may be compounded by additional data-quality issues.
37 Low-coverage NGS data is sometimes the only available option, for instance

38 when working with ancient DNA [12–14]. Sequencing at low coverage is also
39 a popular choice to reduce sequencing costs, since most of the key population
40 genetics analysis remain possible with such data [15].

41 To estimate the SFS from low-coverage data, several methods have been
42 proposed which account for the genotype uncertainty in estimation of the
43 SFS [10, 16]. These are based on finding the SFS that maximises the data
44 likelihood using numeric optimisation. Two factors combine to create a
45 computational challenge for such methods. First, in order to achieve an
46 accurate estimate of the SFS, these methods usually require many iterations,
47 each of which requires a full pass over the input data. Second, unlike most
48 genetics analyses, the SFS cannot be based on only the small subset of the
49 variable sites, but must consider all sites. Taken together, this means that
50 some summary of the full data must be held in RAM and iterated over many
51 times. For genome-scale NGS data from more than a few dozen samples, or
52 in more than one dimension, this is often not computationally feasible, as
53 tens of hours of runtime and hundreds of gigabytes of RAM may be required.
54 Current approaches for dealing with this issue restrict the analysis to fewer
55 individuals and/or smaller regions of the genome [17], leading to less accurate
56 results.

57 An additional problem with current methods is that they are prone to
58 overfitting. In the multi-dimensional setting in particular, there is often very
59 little information available for many of the entries in the frequency spectrum.
60 Therefore, by considering the full data set, existing algorithms risk fitting
61 noise, leading to estimates with poor generalisability.

62 In this paper, we present a novel version of the stochastic expectation-
63 maximisation (EM) algorithm for estimation of the SFS from NGS data. In
64 each pass through the data, this algorithm updates the SFS estimate multiple
65 times in smaller blocks of sites. We show that for low-coverage whole-genome
66 sequencing (WGS) data, this algorithm requires only a few full passes over the
67 data. This considerably decreases running time, and means that it is possible
68 to estimate the SFS using constant, negligible RAM usage by streaming data
69 from disk. Moreover, by only considering smaller subsets of the data at a
70 time, we show that this method reduces overfitting, which in turns leads to
71 improved downstream inference.

72 2 Methods

73 Estimation of the SFS from low-coverage sequencing data requires pre-
74 computing site allele frequency likelihoods for each site, and these are based
75 on genotype likelihoods. We begin by briefly reviewing these concepts.

76 **Genotype likelihoods** Assume we have NGS data X sampled from K
77 different populations (indexed by k), with N_k individuals in the k th population.
78 Further, say that we have M diallelic sites (indexed by m), so that $G_{mkn} \in$
79 $\{0, 1, 2\}$ is the genotype of a diploid individual n at site m in population k ,
80 coding genotypes as the number of derived alleles. In the same way, we use
81 X_{mkn} to refer to the sequencing data at this location.

82 We define the genotype likelihood $p(X_{mkn} | G_{mkn})$ as the probability of
 83 the data given a particular genotype. Genotype likelihoods form the basis
 84 of genotype calling and are calculated from aligned sequencing reads by
 85 various bioinformatic tools including **bcftools/samtools** [18, 19], GATK [20],
 86 and ANGSD [21], using slightly different models. For clarity, we outline the
 87 basic GATK model below, though the choice of model is not important for our
 88 purposes.

89 For D sequencing reads aligned to position m for individual n in population
 90 k , let b_d be the base call of the d th read. Assuming independence of base
 91 calls, we have

$$p(X_{mkn} | G_{mkn} = g) = \prod_{d=1}^D p(b_d | G_{mkn} = g). \quad (1)$$

92 If we consider the genotype as two alleles $a_1, a_2 \in \{0, 1\}$ such that $G_{mkn} =$
 93 $a_1 + a_2$, then by random sampling of the parental alleles,

$$p(b_d | G_{mkn} = g) = \frac{1}{2}p(b_d | a_1) + \frac{1}{2}p(b_d | a_2). \quad (2)$$

94 In turn, this probability is modelled by

$$p(b_d | a) = \begin{cases} \epsilon_d/3 & \text{if } b_d \neq a \\ 1 - \epsilon_d & \text{else} \end{cases}, \quad (3)$$

95 where ϵ_d is the sequencing error probability associated with the d th base.

96 **Site allele frequency likelihoods** Using genotype likelihoods, we can cal-
 97 culate site allele frequency (SAF) likelihoods, also sometimes known as sample
 98 allele frequency likelihoods. It is possible to think of the SAF likelihoods as
 99 the generalisation of genotype likelihoods from individuals to populations:
 100 instead of asking about the probability of the data for one individual given a
 101 genotype, we ask about the probability of the data for a population given the
 102 sum of their derived alleles.

103 More formally, define the sum of derived alleles for population k at site
 104 m ,

$$Z_{mk} = \sum_{n=1}^{N_k} G_{mkn}, \quad (4)$$

105 with $Z_{mk} \in \{0, 1, \dots, 2N_k\}$ each corresponding to possible sample frequencies
 106 $\{0, 1/2N_k, \dots, 1\}$. Now define the SAF likelihood for a single population k ,

$$p(X_{mk} | Z_{mk} = j_k) = \sum_{g \in \{0,1,2\}^{N_k}} p(g | Z_{mk} = j_k) \prod_{n=1}^{N_k} p(X_{mkn} | G_{mkn} = g_n), \quad (5)$$

107 where X_{mk} is the data for all individuals sampled in population k at site
 108 m , $p(g | Z_{mk} = j_k)$ is the combinatorial probability of the genotype vector
 109 $g = (g_1, \dots, g_{N_k})$ conditional on the sum of the genotypes being j_k , and
 110 $p(X_{mkn} | G_{mkn} = g_n)$ is a standard genotype likelihood. Using a dynamic

111 programming algorithm, SAF likelihoods can be calculated from the genotype
 112 likelihoods of N individuals in $O(N^2)$ time per site [22], and a linear time
 113 approximation has also been given [23].

114 To extend this to the multi-dimensional SFS with K populations, let $\mathcal{J} =$
 115 $\times_{k=1}^K \{0, 1, \dots, 2N_k\}$ be the set of possible derived allele count combinations
 116 across populations, let X_m be the data across all individuals in all populations
 117 at site m , and define $Z_m = (Z_{m1}, \dots, Z_{mK}) \in \mathcal{J}$. Then

$$p(X_m | Z_m) = \prod_{k=1}^K p(X_{mk} | Z_{mk}), \quad (6)$$

118 is the joint SAF likelihood for K populations.

119 **Site frequency spectrum** Using the definition of \mathcal{J} above, we define the
 120 SFS as a parameter $\phi = \{\phi_j : j \in \mathcal{J}\}$ such that ϕ_j is the probability that
 121 $Z_m = j$. That is,

$$\phi_j = p(Z_m = j | \phi), \quad (7)$$

122 for site m . That is, the SFS is the probability of a particular vector of derived
 123 allele sums at a site chosen at random.

124 When genotypes are available, the SFS can be estimated simply by
 125 counting observed allele count combinations. When genotypes cannot be
 126 called, the standard approach is maximum-likelihood estimation.

127 Assuming independence of sites, we write the likelihood function

$$\begin{aligned} p(X | \phi) &= \prod_{m=1}^M p(X_m | \phi) \\ &= \prod_{m=1}^M \sum_{j \in \mathcal{J}} p(X_m | \phi, Z_m = j) p(Z_m = j | \phi) \\ &= \prod_{m=1}^M \sum_{j \in \mathcal{J}} p(X_m | Z_m = j) \phi_j, \end{aligned} \quad (8)$$

128 where X_m refers to all sequencing data for site m . Note that the likelihood
 129 can be expressed solely in terms of joint SAF likelihoods.

130 The maximum likelihood estimate $\hat{\phi} = \arg \max_{\phi} p(X | \phi)$ cannot be found
 131 analytically. Instead, $\hat{\phi}$ is typically estimated using some iterative procedure
 132 such as BFGS [22] or an EM algorithm [16, 21]), of which the latter has
 133 become the standard choice. An overview of the this algorithm is given below.
 134 For details and proof, see supplementary text S1.

135 **Standard EM algorithm** Before optimization, we pre-compute the SAF
 136 likelihoods for all sites, populations, and possible sample frequencies. In
 137 addition, we make an arbitrary initial guess of the SFS $\hat{\phi}^{(0)}$. The EM
 138 algorithm then alternates between an E-step, and an M-step.

¹³⁹ The E-step consists of computing posterior probabilities of derived allele
¹⁴⁰ counts conditional on the current SFS estimate,

$$\begin{aligned} q_{mj}^{(t)} &= p(Z_m = j \mid X_m, \hat{\phi}_j^{(t)}) \\ &= \frac{p(X_m \mid Z_m = j) \hat{\phi}_j^{(t)}}{\sum_{j' \in \mathcal{J}} p(X_m \mid Z_m = j') \hat{\phi}_{j'}^{(t)}}, \end{aligned} \quad (9)$$

¹⁴¹ for all sites $m \in \{1, \dots, M\}$ and possible derived allele counts $j \in \mathcal{J}$. Note
¹⁴² that this conditional posterior depends only on the current SFS estimate and
¹⁴³ the (joint) SAF likelihoods.

¹⁴⁴ Using the result of the E-step, the M-step updates the estimate by setting

$$\begin{aligned} \hat{\phi}_j^{(t+1)} &= \frac{\sum_{m=1}^M q_{mj}^{(t)}}{\sum_{m=1}^M \sum_{j' \in \mathcal{J}} q_{mj'}^{(t)}} \\ &= \frac{1}{M} \sum_{m=1}^M q_{mj}^{(t)}, \end{aligned} \quad (10)$$

¹⁴⁵ for all $j \in \mathcal{J}$.

¹⁴⁶ The EM algorithm guarantees a monotonically increasing likelihood of
¹⁴⁷ successive values of $\hat{\phi}^{(t)}$. The runtime of the algorithm is linear in the
¹⁴⁸ number of iterations required before convergence, with each iteration taking
¹⁴⁹ $O(M \prod_{k=1}^K N_k)$ time. In practice, the standard implementation is `realSFS`
¹⁵⁰ [22] from the software suite `ANGSD` [21] which uses a generic EM acceleration
¹⁵¹ scheme [24]. The details of this acceleration will not be important in this
¹⁵² context, so we omit the details.

¹⁵³ **Window EM algorithm** As in standard EM, we pre-compute all SAF
¹⁵⁴ likelihoods and make an arbitrary initial guess $\hat{\phi}^{(0)}$ of the SFS. In addition,
¹⁵⁵ we choose two hyperparameters B (the number of blocks) and W (the window
¹⁵⁶ size). Before starting optimization, all sites indices are randomly assigned
¹⁵⁷ to one of B blocks $\mathcal{B} = (\mathcal{B}_1, \dots, \mathcal{B}_B)$ with $|\mathcal{B}_b| = \lfloor M/B \rfloor$ for $b < B$, and
¹⁵⁸ $|\mathcal{B}_B| = M \bmod B$. The reason for doing so is simply to break patterns of
¹⁵⁹ linkage disequilibrium in particular blocks of input data, which will make
¹⁶⁰ the SFS within each block more similar to the global SFS. Blocks are non-
¹⁶¹ overlapping and exhaustive, so that $\bigcup_{b=1}^B \mathcal{B}_b = \{1, \dots, M\}$ and $\bigcap_{b=1}^B \mathcal{B}_b = \emptyset$.

¹⁶² After this initialisation, the window EM algorithm is defined as an iterative
¹⁶³ procedure that alternates between an E-step and an M-step, where the M-step
¹⁶⁴ in turn is split into an M1-step and an M2-step.

¹⁶⁵ The E-step of the algorithm involves computing posteriors conditional on
¹⁶⁶ the current estimate of the SFS, much like standard EM. The difference is
¹⁶⁷ that we only process a single block of sites. Let $f(t) = (t-1) \bmod B + 1$, so
¹⁶⁸ that $f(1+xB) = 1, f(2+xB) = 2, \dots$ for $x \geq 0$. Then, at time step t , we
¹⁶⁹ compute $q_{mj}^{(t)}$ for all $m \in \mathcal{B}_{f(t+1)}$ and all possible derived allele counts $j \in \mathcal{J}$
¹⁷⁰ using eq. (9).

¹⁷¹ In the M1-step, the qs for the current block are used to give a block SFS
¹⁷² estimate $\hat{\psi}^{(t)}$. This is analogous to the standard M-step eq. (10), so that for

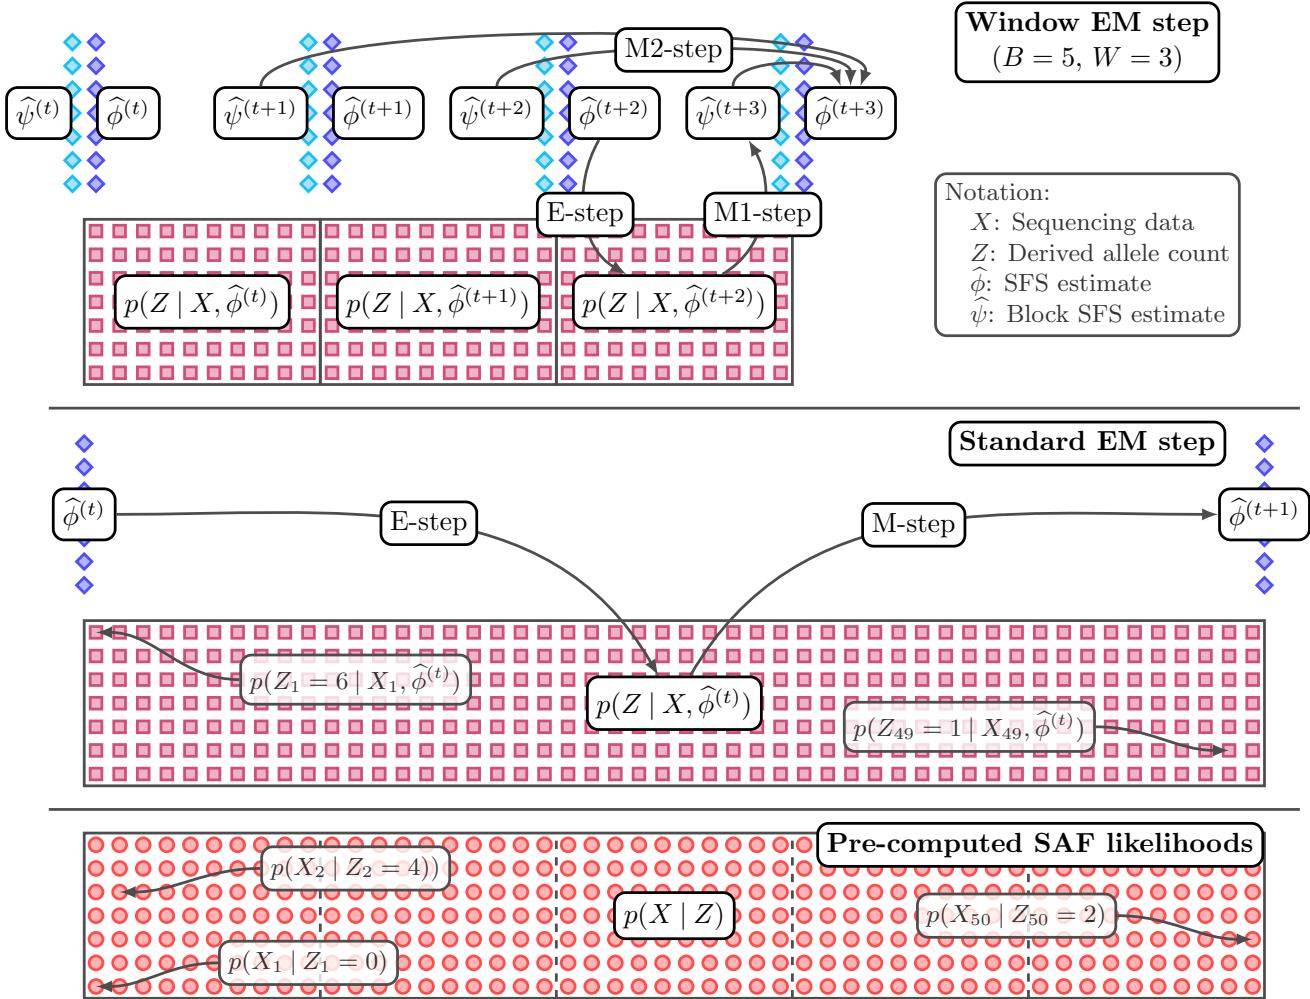


Figure 1: Schematic illustration of the standard and window EM algorithms for input consisting of a single population with $N = 3$ individuals and $M = 50$ sites. Sites are shown horizontally, derived allele frequencies are shown vertically. The pre-computed SAF likelihoods are illustrated at the bottom with blocks indicated by dashed lines. Standard EM computes the conditional posterior derived allele counts over all sites (E-step) and uses these to update the SFS estimate (M-step). Window EM computes the conditional posteriors for a small blocks of sites (E-step), computes a block SFS estimate after each block (M1-step), and updates the overall estimates as sliding window average (M2-step) of the W past block estimates. In this example, the sites have been split into $B = 5$ blocks with 10 sites each, and the sliding window covers $W = 3$ blocks.

Algorithm 1: Window EM algorithm

Input : (1) SAF likelihoods $p(X_{mk} | Z_{mk} = j_k)$ for sites $m \in \{1, \dots, M\}$ and N_k individuals in each of populations $k \in \{1, \dots, K\}$, with $j \in \mathcal{J} = \times_{k=1}^K \{0, 1, \dots, 2N_k\}$. (2) Random, non-overlapping assignment of sites indices from 1 to M into B blocks $(\mathcal{B}_1, \dots, \mathcal{B}_B)$. (3) Initial SFS estimate $\hat{\phi}^{(0)}$.

Output : Estimate $\hat{\phi}$ of the K -dimensional SFS.

Parameters: Number of blocks B , number of blocks per window W .

```

 $t \leftarrow 0$ 
while not converged do
   $b_t \leftarrow t \bmod B + 1$  // Block index
  for  $m \in \mathcal{B}_{b_t}$  do
    for  $j \in \mathcal{J}$  do
       $q_{mj}^{(t)} \leftarrow \frac{p(X_m | Z_m = j) \hat{\phi}_j^{(t)}}{\sum_{j' \in \mathcal{J}} p(X_m | Z_m = j') \hat{\phi}_{j'}^{(t)}}$  // E-step
     $\mathcal{W}_t \leftarrow \{(t-w) \bmod B + 1 \mid w \in \{0, \dots, \min(t, W-1)\}\}$  // Window indices
    for  $j \in \mathcal{J}$  do
       $\hat{\psi}_j^{(t+1)} \leftarrow \frac{1}{|\mathcal{B}_{b_t}|} \sum_{m \in \mathcal{B}_{b_t}} q_{mj}^{(t)}$  // M1-step
       $\hat{\phi}_j^{(t+1)} \leftarrow \frac{1}{\sum_{w \in \mathcal{W}_t} |\mathcal{B}_w|} \sum_{w \in \mathcal{W}_t} \hat{\psi}_j^{(w)} |\mathcal{B}_w|$  // M2-step
     $t \leftarrow t + 1$ 
return  $\hat{\phi}^{(t)}$ 

```

173 each $j \in \mathcal{J}$

$$\hat{\psi}_j^{(t+1)} = \frac{1}{|\mathcal{B}_{f(t+1)}|} \sum_{m \in \mathcal{B}_{f(t+1)}} q_{mj}^{(t)}, \quad (11)$$

174 These block estimates are then used in the M2-step to update the overall
 175 SFS estimate for each $j \in \mathcal{J}$,

$$\begin{aligned} \hat{\phi}_j^{(t+1)} &= \frac{1}{\sum_{w \in \mathcal{W}_t} |\mathcal{B}_w|} \sum_{w \in \mathcal{W}_t} \hat{\psi}_j^{(w)} |\mathcal{B}_w| \\ &\stackrel{*}{=} \frac{1}{W} \sum_{w \in \mathcal{W}_t} \hat{\psi}_j^{(w)}. \end{aligned} \quad (12)$$

176 where $\mathcal{W}_t = \{f(t+1-w) \mid w \in \{0, \dots, \min(t, W-1)\}\}$ is the window of
 177 the W latest block indices at time t . We use $\stackrel{*}{=}$ to express equality under the
 178 common special case when either $M/B = 0$ or $B \notin \mathcal{W}_t$, so that there are no
 179 issues with blocks of unequal sizes in the current window. In this case, the
 180 M2-step simplifies to the mean of the past W block estimates.

181 Pseudo-code for window EM is given in algorithm 1, and an illustration
 182 comparing window EM to standard EM is shown in figure 1.

183 In the below, we are interested in comparing standard EM and window
 184 EM. For clarity, we will use the term ‘epoch’ to refer to a full pass through

185 the data for either algorithm. In the case of standard EM, an epoch is simply
186 a single iteration; for window EM, an epoch corresponds to B iterations.

187 **Convergence** In the standard EM algorithm, the data log-likelihood (8)
188 can typically be evaluated with little computational overhead during the
189 E-step. Therefore, a common convergence criterion is based on the difference
190 between the log-likelihood values of successive epochs. That is, let

$$L_t = \frac{1}{M} \sum_{m=1}^M \log p(X | \hat{\phi}^{(t)}), \quad (13)$$

191 and convergence is reached when $L_{t+1} - L_t < \delta$, for some tolerance δ decided
192 ahead of time.

193 For window EM, the same does not apply, since no full E-step is ever
194 taken. However, the likelihood for each block can be calculated cheaply during
195 each block E-step. Therefore, we define for epoch $e \in \{1, 2, \dots\}$,

$$L'_e = \sum_{b=1}^B \frac{1}{|\mathcal{B}_b|} \sum_{m \in \mathcal{B}_b} \log p(X_m | \hat{\phi}^{(eB-b)}), \quad (14)$$

196 that is, the sum of log-likelihoods of SFS estimates used over the past epoch,
197 each evaluated in the block for which they were used in a block E-step, norm-
198 alised by block size for convenience. We then propose the simple convergence
199 criterion for window EM such that convergence is defined as $L'_{e+1} - L'_e < \delta$.

200 3 Results

201 To test the window EM algorithm, we implemented it in the `winsfs` pro-
202 gram, available at github.com/malthesr/winsfs. We compare `winsfs` to
203 `realSFS`, which implements the standard EM algorithm and serves as the
204 current state of art. We adopt two complementary approaches for evaluating
205 performance of `winsfs`. First, we use two different real-world WGS data
206 sets to compare `winsfs` to `realSFS`, which implements the standard EM
207 algorithm and serves as the current state of the art. `realSFS` has already
208 been validated on simulated data [21, 23], and use split training and test
209 data sets to evaluate any observed differences. Second, we use simulated data
210 to validate `winsfs` under conditions of known truth across a range of data
211 qualities and sample sizes.

212 **Real-world data sets** We tested `winsfs` and `realSFS` on two real-world
213 WGS data sets of very different quality as described below. An overview is
214 shown in table 1.

215 We first analyse 10 random individuals from each of the YRI (Yoruba
216 Nigerian) and CEU (Europeans in Utah) populations from the 1000 Genomes
217 Project [25]. This human data was sequenced to 3x–8x coverage and mapped
218 to the high quality human reference genome. We created SAF files using `ANGSD`
219 [21] requiring minimum base and mapping quality 30 and polarising the

Population	Individuals	Sites	Median depth (range)	Contigs ≥ 100 kb	n_{50}	F_{ST}
Human						
YRI	10	$1.17 \cdot 10^9$	$5.0 \times (3.2 \times - 7.2 \times)$		52	$1.5 \cdot 10^8$ 0.13
CEU	10	$1.17 \cdot 10^9$	$6.2 \times (2.9 \times - 8.0 \times)$			
Impala						
Masai Mara	12	$6.34 \cdot 10^8$	$2.8 \times (1.4 \times - 10.2 \times)$		7811	$3.4 \cdot 10^5$ 0.24
Shangani	8	$6.34 \cdot 10^8$	$2.9 \times (2.6 \times - 16.8 \times)$			

Table 1: Overview of the input training data.

220 spectrum using the chimpanzee as an outgroup. We then split this input
 221 data into test and training data, such that the first half of each autosome
 222 was assigned to the training set, and the second half to the test set. The
 223 resulting training data set contains $1.17 \cdot 10^9$ sites for both YRI and CEU,
 224 while the test data set contains $1.35 \cdot 10^9$ sites for both. Training set depth
 225 distributions for each individual are shown in supplementary figure 1.

226 We also analyse a data set of much lower quality from 12 and 8 individuals
 227 from two impala populations that we refer to as ‘Maasai Mara’ and ‘Shangani’,
 228 respectively, based on their sampling locations. These populations were
 229 sequenced to only $1 \times - 3 \times$ with the addition of a single high-depth sample in
 230 each population (see supplementary figure 2). The data was mapped to a
 231 very fragmented assembly, and then we split the data into training and test
 232 sets just as for the human data. However, due to the low quality assembly
 233 we analysed only sites on contigs larger than 100 kb, and filtering sites based
 234 on depth outliers, excess heterozygosity, mappability, and repeat regions. We
 235 polarised using the impala reference itself. This process is meant to mirror
 236 a realistic workflow for working with low-quality data from a non-model
 237 organism. The impala input data ends up somewhat smaller than the human
 238 data set, with approximately $6.3 \cdot 10^8$ sites in both test and training data
 239 sets.

240 Broadly, the human data is meant to exemplify medium-quality data with
 241 coverage towards the lower end, but with no other significant issues. The
 242 impala data, on the other hand, represents low-quality data: not only is the
 243 coverage low and fewer sites are available, but the impala reference genome is
 244 poor quality with 7811 contigs greater than 100 kb and $n_{50} = 3.4 \cdot 10^{-5}$ (that
 245 is, 50 % of the assembly bases lie on contigs of this size or greater). This serves
 246 to introduce further noise in the mapping process, which amplifies the overall
 247 data uncertainty. Finally, the impala populations are more distinct, with
 248 $F_{ST} \approx 0.24$ compared to 0.13 between the human populations. As we will see
 249 below, this creates additional challenges for estimation of the two-dimensional
 250 SFS.

251 **Estimation** Using the training data sets, we estimated the one-dimensional
 252 SFS for YRI and Maasai Mara, as well as the two-dimensional SFS for
 253 CEU/YRI and Shangani/Maasai Mara. We ran `winsfs` for 500 epochs using
 254 a fixed number of blocks $B = 500$ and window sizes $W \in \{100, 250, 500\}$.

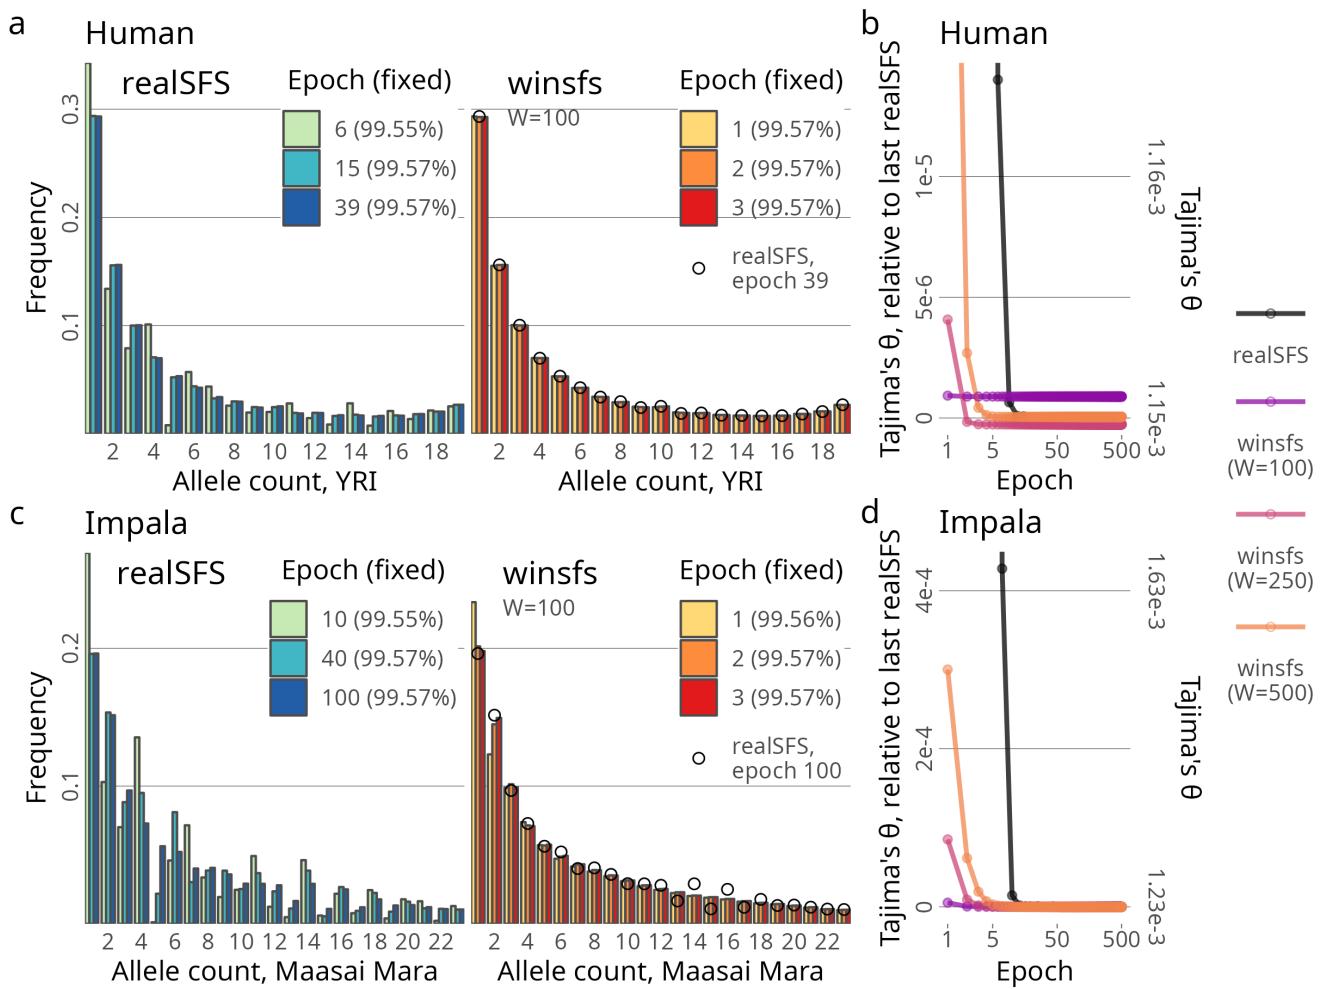


Figure 2: One-dimensional SFS estimation. **(a):** YRI SFS estimates from **realSFS** and **winsfs**₁₀₀ after various epochs. Only variable sites are shown, proportion of fixed sites is shown in the legend. The final **realSFS** estimate is overlaid with dots on the **winsfs** plot for comparison. **(b):** YRI Tajima's θ estimates calculated from **realSFS** and **winsfs** over epochs. **(c):** Maasai Mara SFS estimates from **realSFS** and **winsfs**₁₀₀ after various epochs. Only variable sites are shown, proportion of fixed sites is shown in the legend. The final **realSFS** estimate is overlaid with dots on the **winsfs** plot for comparison. **(d):** Maasai Mara Tajima's θ estimates calculated from **realSFS** and **winsfs** over epochs.

255 We will focus on the setting with window size $W = 100$. For convenience,
 256 we introduce the notation **winsfs**₁₀₀ to refer to **winsfs** with hyperparameter
 257 settings $B = 500$, $W = 100$. We return to the topic of hyperparameter
 258 settings below.

259 To compare, we ran **realSFS** using default settings, except allowing it
 260 to run for a maximum of 500 epochs rather than the default 100. We will
 261 still take the 100 epochs cut-off to mark convergence, if it has not occurred by
 262 other criteria before then, but results past 100 will be shown in places.

263 In each case, we evaluated the full log-likelihood (eq. (8)) of the estimates
 264 after each epoch on both the training and test data sets. In addition, we
 265 computed various summary statistics from the estimates after each epoch.
 266 For details, see supplementary text S2.

267 **One-dimensional SFS** Main results for the one-dimensional estimates are
268 shown in figure 2.

269 For the human YRI population, we find that a single epoch of `winsfs100`
270 produces an estimate of the SFS that is visually indistinguishable from the
271 converged estimate of `realSFS` at 39 epochs (figure 2a). Train and test set
272 log-likelihoods (supplementary figure 3) confirm that the likelihood at this
273 point is only very marginally lower for `winsfs100` than the last `realSFS`. By
274 increasing the window size to 250 or 500, we get test log-likelihood values
275 equal to or above those achieved by `realSFS`, and still within the first 5
276 epochs.

277 As an example of a summary statistic derived from the one-dimensional
278 SFS, figure 2b shows that `winsfs100` finds an estimate of Tajima's θ that
279 is very near to the final `realSFS`, with a difference on the order of $1 \cdot 10^{-6}$.
280 Increasing the window size removes this difference at the cost of a few more
281 epochs.

282 In the case of Maasai Mara, `realSFS` runs for the 500 epochs, so we take
283 epoch 100 to mark convergence. On this data, `winsfs100` requires two epochs
284 to give a good estimate of the SFS, as shown in figure 2c. Some subtle
285 differences relative to the `realSFS` results remain, however, especially at the
286 middle frequencies: the `realSFS` estimate exhibits a 'wobble' such that even
287 bins are consistently higher than odd bins. Such a pattern is not biologically
288 plausible, and is not seen in the `winsfs` spectrum.

289 Supplementary figure 4 shows train and test log-likelihood data for Maasai
290 Mara, which again support the conclusions drawn from looking at the estimates
291 themselves. In theory, we expect that the test log-likelihood should be
292 adversely impacted by the `realSFS` 'wobble' pattern. In practice, however,
293 with more than 99.5 % fixed sites, the fixed end of the spectrum dominate
294 the likelihood to the extent that the effect is not visible. We return to this
295 point below.

296 Finally, Figure 2d shows that Tajima's θ is likewise well-estimated by one
297 or two epochs of `winsfs100` on the impala data.

298 **Two-dimensional SFS** Overall results for the joint spectra are seen in
299 figure 3.

300 On the human data, `winsfs100` takes a single epoch for an estimate of
301 the SFS that is near-identical to `realSFS` at convergence after 93 epochs.
302 Looking at the log-likelihood results, it is notable that while `realSFS` does
303 better than `winsfs` when evaluated on the training data (figure 3b), the
304 picture is reversed when evaluated on the test data (figure 3c). In fact, all
305 `winsfs` hyperparameter settings achieved better test log-likelihood values
306 in the first 10 epochs than achieved by `realSFS` at convergence. This is
307 likely caused by a faint 'checkerboard' pattern in the `realSFS` estimate due
308 to overfitting, as we expect the spectrum to be smooth. We note that both
309 `realSFS` and `winsfs` preserve an excess of sites where all individuals are
310 heterozygous, corresponding to the peak in the centre of the spectrum. This
311 is a known issue with this data set [26], likely caused by paralogs in the
312 mapping process. It is an artefact which can be removed by filtering the data

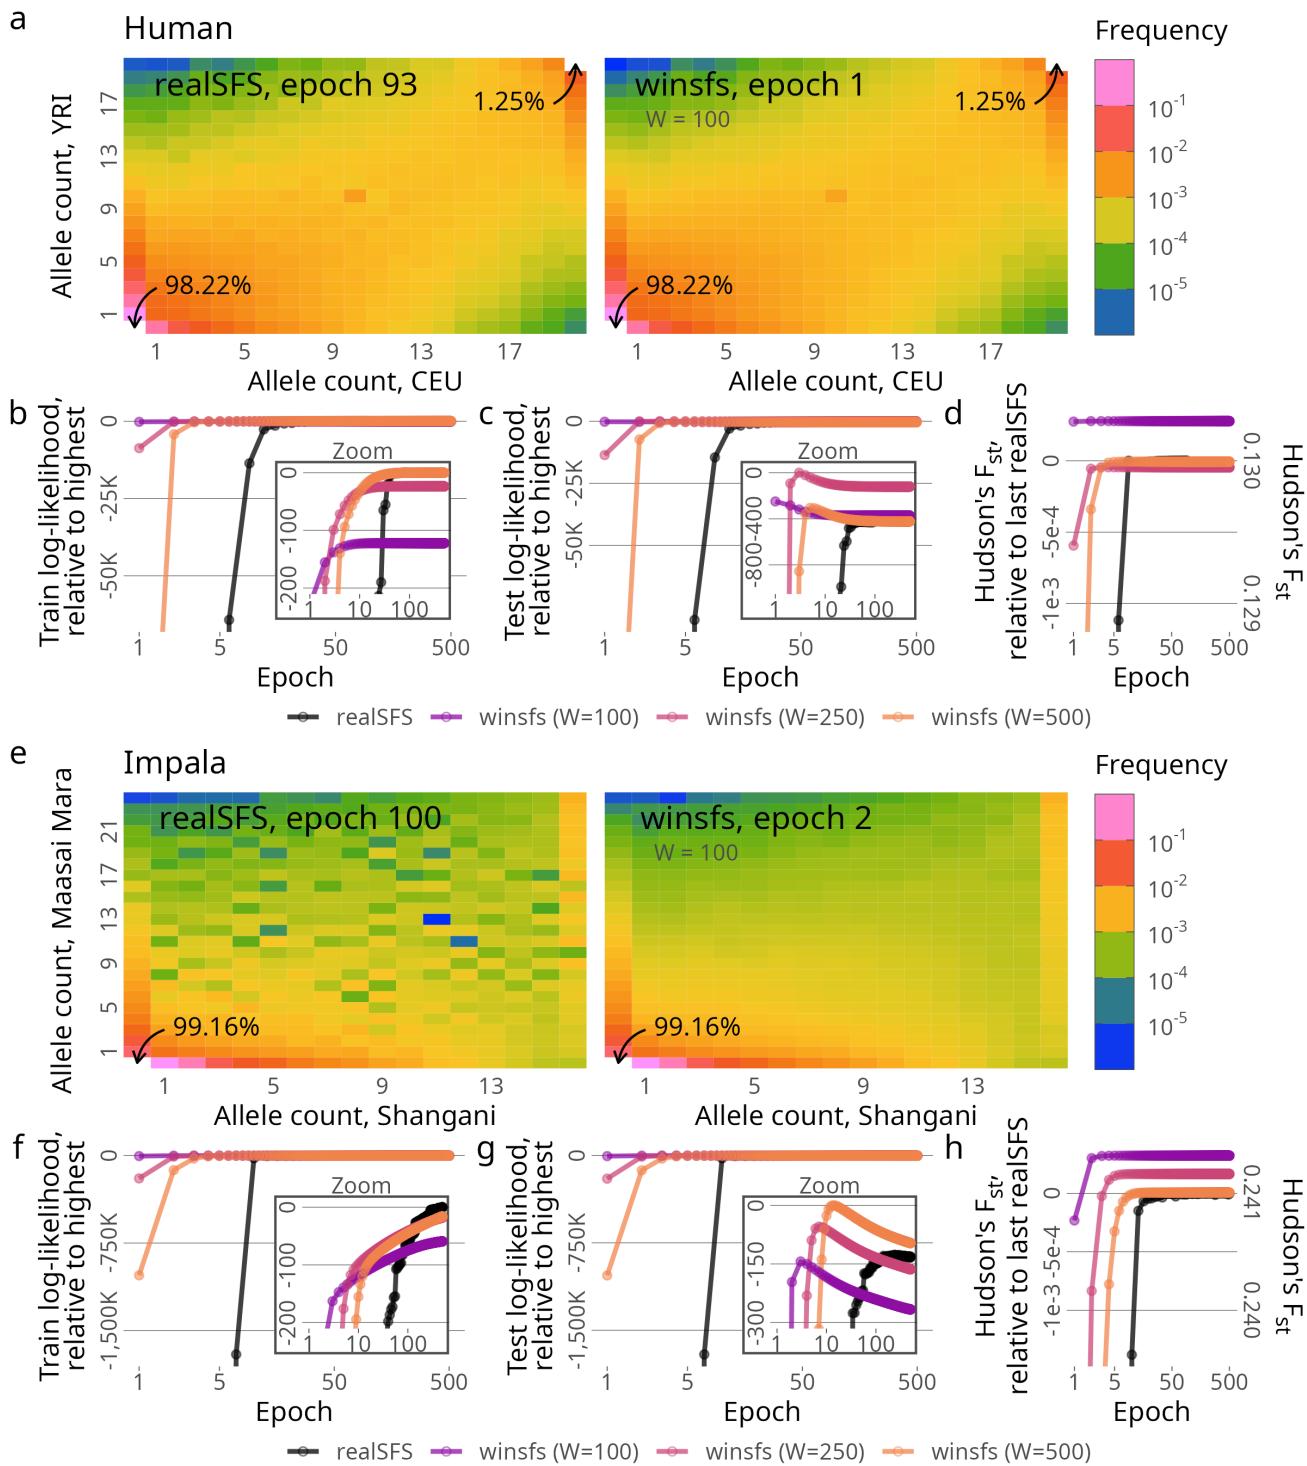


Figure 3: Two-dimensional SFS estimation. **(a):** CEU/YRI SFS estimates from **realSFS** after 93 epochs (converged) and from **winsfs₁₀₀** after a single epoch. Fixed sites not shown for scale, total proportion indicated by arrows. **(b), (c):** CEU/YRI SFS train and test log-likelihood over epochs for **realSFS** and **winsfs**. **(d):** CEU/YRI Hudson's F_{ST} estimates calculated from **realSFS** and **winsfs** over epochs **(e):** Shangani/Maasai Mara SFS estimates from **realSFS** after 100 epochs (converged) and from **winsfs₁₀₀** after a single epoch. Fixed reference sites not shown for scale, proportions indicated by arrows. **(f), (g):** Shangani/Maasai Mara SFS train and test log-likelihood over epochs for **realSFS** and **winsfs**. **(h):** Shangani/Maasai Mara Hudson's F_{ST} estimates calculated from **realSFS** and **winsfs** over epochs.

313 before SAF calculation, which we have not done here. Given this choice, it is
314 to be expected that this peak remains.

315 In two dimensions, we compute both Hudson's F_{ST} (figure 3d) and the
316 f_2 -statistic (supplementary figure 5) from SFS estimates after all epochs,
317 and we note similar patterns for these as we have seen before: one epoch of
318 `winsfs100` gives an estimate of the summary statistic that is almost identical
319 to the final `realSFS` estimate.

320 For the impalas, `winsfs100` requires two epochs for a good estimate of
321 the spectrum, while `realSFS` again does not report convergence within the
322 first 100. What is immediately striking about the impala results, however,
323 is that the checkerboard pattern is very pronounced for `realSFS`, and again
324 absent for `winsfs` (figure 3e). The problem for `realSFS` is likely exacerbated
325 by two factors: first, the sequencing depth is lower, increasing the uncertainty;
326 second, the relatively high divergence of the impala populations push most
327 of the mass in the spectrum towards the edges. Together, this means that
328 very little information is available for most of the estimated parameters. It
329 appears that `realSFS` therefore ends up overfitting to the particularities of
330 the training data at these bins.

331 This is also reflected in the difference between train and test log-likelihood
332 (figures 3f and 3g). Like in the case of the human data, the SFS estimated
333 by `winsfs` performs better on the test data compared to `realSFS`, while
334 `realSFS` performs the based on the training data. On the test data, all
335 `winsfs` settings again reach log-likelihood values comparable to or better
336 than `realSFS` in few epochs. However, the differences between `realSFS` and
337 `winsfs` remain relatively small in terms of log-likelihood, even on the test
338 set. This is somewhat surprising, given the marked checkerboarding in the
339 spectrum itself. Again, we attribute this to the fact that the log-likelihood is
340 dominated by all the mass lying in or around the zero-zero bin. We expect,
341 therefore, that methods that rely on the 'interior' of the SFS should do better
342 when using `winsfs`, compared to `realSFS`.

343 Before turning to test this prediction, we briefly note that F_{ST} (figure 3h)
344 and the f_2 -statistic (supplementary figure 5) are also adequately estimated
345 for the impalas by `winsfs100` in one epoch.

346 **Demographic inference** All the SFS-derived summary statistics con-
347 sidered so far are heavily influenced by the bins with the fixed allele bins
348 (that is, count 0 or $2N_k$ in all populations), or they are sums of alternating
349 frequency bins. In either case, this serves to mask issues with checkerboard
350 areas of the SFS in the lower-frequency bins. However, this will not be the
351 case for downstream methods that rely on the shape of the spectrum in more
352 detail.

353 To illustrate, we present a small case-study of inferring the demographic
354 history of the impala populations using the *ðaði* [5] software with the es-
355 timated impala spectra shown in figure 3e, though folded due to the lack
356 of an outgroup for proper polarisation. Briefly, based on an estimated SFS
357 and a user-specified demographic model, *ðaði* fits a model SFS based on the
358 demographic parameters so as to maximise the likelihood of these parameters.

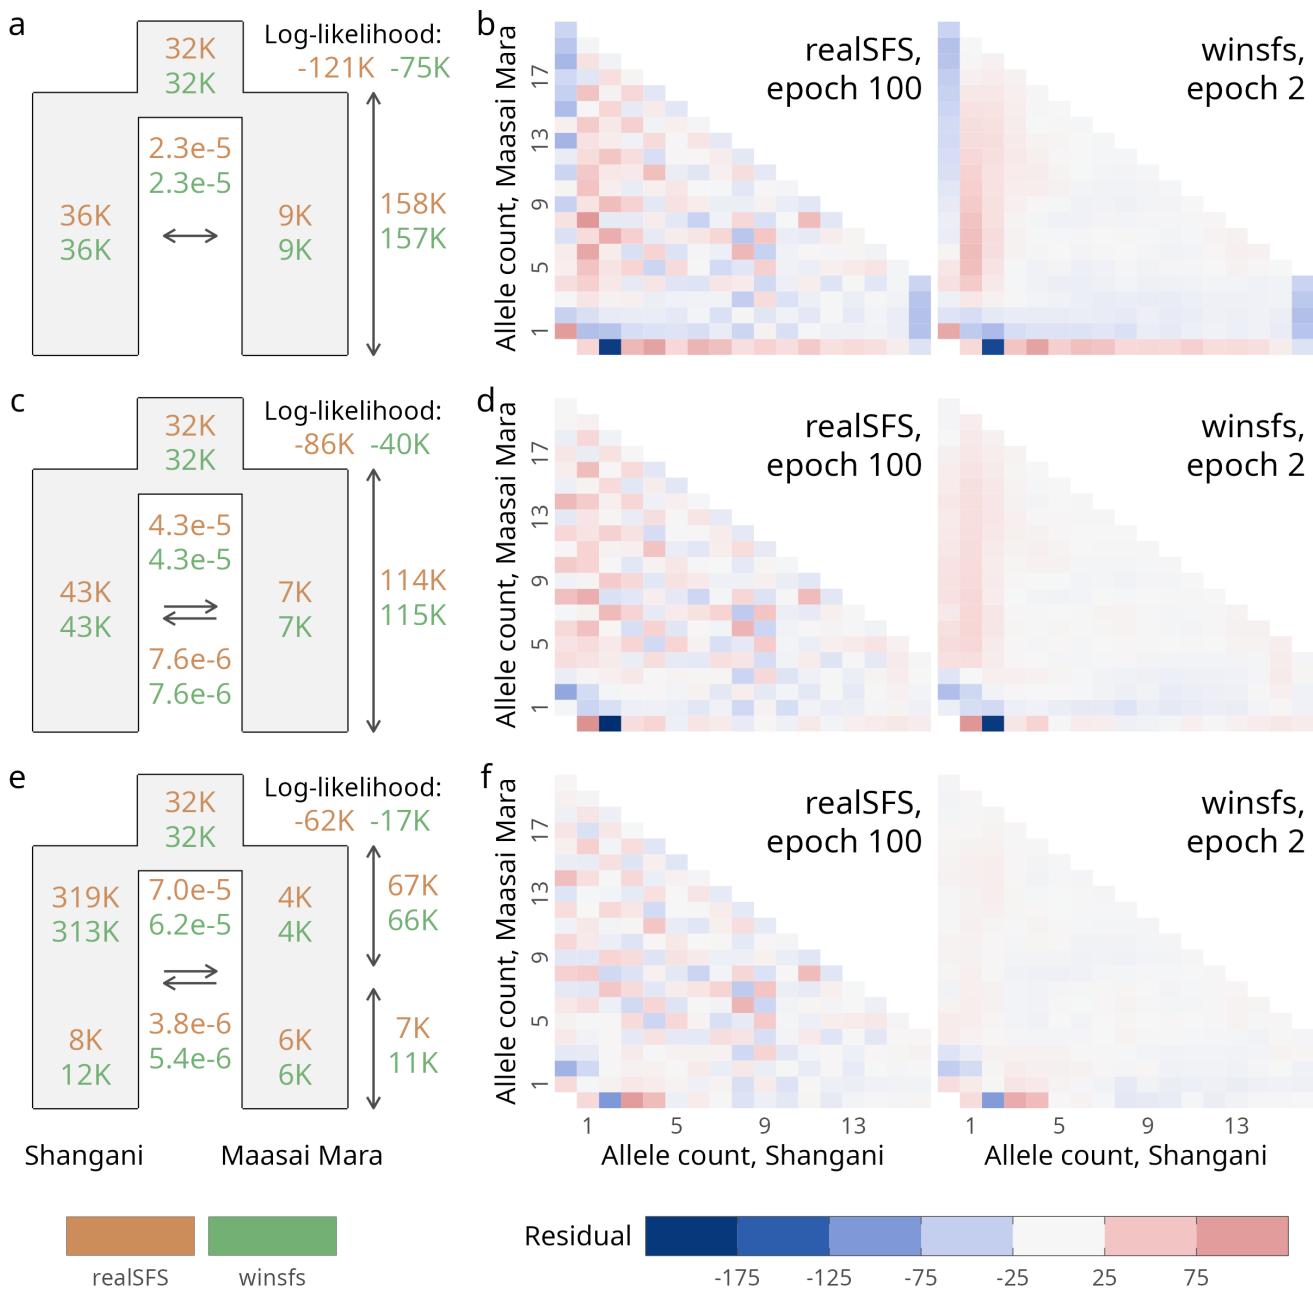


Figure 4: Demographic inference results. Each row corresponds to a demographic model fitted using $\partial a\partial i$. On the left, a schematic of the model is shown including parameter estimates using SFS estimates from **realSFS** after 100 epochs or from **winsfs₁₀₀** after two epochs. Time is given in years, population sizes in number of individuals, and migration rates is per chromosome per generation. All parameters were scaled assuming a mutation rate of $1.41 \cdot 10^{-8}$ per site per generation and a generation time of 5.7 years. On the right, the residuals of the SFS fitted by $\partial a\partial i$. Note that $\partial a\partial i$ folds the input SFS, hence the residuals are likewise folded. The fixed category is omitted to avoid distorting the scale. **(a), (b):** Model with symmetric migration and constant population size. **(c), (d):** Model with asymmetric migration and constant population size. **(e), (f):** Model with asymmetric migration and a single, instantaneous population size change.

359 Our approach was to fit a simple demographic model for the Shangani and
360 Maasai Mara populations, and then gradually add parameters to the model
361 as required based on the residuals of the input and model spectra. We take
362 this to be representative of a typical workflow for demographic inference.

363 For each successive demographic model [27], we ran $\partial a \partial i$ on the folded
364 spectra by performing 100 independent optimisation runs from random start-
365 ing parameters, and checking for convergence by requiring the top three
366 results to be within five likelihoods units of each other. If the optimisation did
367 not converge, we did additional optimisation runs until either they converged
368 or 500 independent runs were reached without likelihood convergence. In
369 that case, we inspected the results for the top runs, to assess whether they
370 were reliably reaching similar estimates and likelihoods. Results are shown in
371 figure 4.

372 The first, basic model assumes that the populations have had constant
373 populations sizes and a symmetric migration rate since diverging. The
374 parameter estimates based on **realSFS** and **winsfs** are similar, though the
375 **winsfs** model fit has significantly higher log-likelihood (figure 4a). However,
376 when inspecting the residuals in figure 4b, the **realSFS** residuals suffer from a
377 heavy checkerboard pattern, making it hard to distinguish noise from model
378 misspecification. In contrast, the **winsfs** residuals clearly show areas of the
379 spectrum where the model poorly fits the data.

380 In particular, the residuals along the very edge of the spectrum suggest
381 that a symmetric migration rate is not appropriate. Therefore, we fit a second
382 model with asymmetric migration (figure 4c) Now $\partial a \partial i$ finds migration rates
383 from Shangani to Maasai Mara an order of magnitude higher than *vice versa*.
384 The results for **winsfs** (figure 4d) show improved residuals, while the **realSFS**
385 residuals remain hard to interpret.

386 Finally, an area of positive residuals in the fixed and rare-variant end of
387 the Shangani spectrum suggests that this population has recently undergone a
388 significant bottleneck. Therefore, the third model allows for an instantaneous
389 size change in each of the impala populations (figure 4e). At this point, the
390 **winsfs** residuals (figure 4f) are negligible, suggesting that no more parameters
391 should be added to the model. Once again, though, the **realSFS** residuals
392 leave us uncertain whether further model extensions are required.

393 When looking at the final model fits, the $\partial a \partial i$ parameter estimates from
394 **realSFS** and **winsfs** also start to differ slightly. In several instances, estimates
395 disagree by about 50 %, and the log-likelihood remains much higher for **winsfs**,
396 with a difference of 45 000 log-likelihood units to **realSFS**. In addition, we
397 confirmed that the log-likelihood of the original test data set given the SFS
398 fitted by $\partial a \partial i$ is higher for **winsfs** ($-8.08 \cdot 10^8$) than for **realSFS** ($-8.38 \cdot 10^8$).
399 We stress, however, that we would have likely never found the appropriate
400 model without using **winsfs**, since the interpretation of the **realSFS** results
401 is difficult. In relation to this point, we note that the final model results in
402 considerably different estimates for parameters of biological interest, such
403 as split times and recent population sizes, relative to the initial model. We
404 also find that the last model is supported by the literature: previous genetic
405 and fossil evidence suggests extant common impala populations derive from
406 a refugia in Southern Africa that subsequently colonised East Africa in

407 the middle-to-late Pleistocene [28–30]. This is broadly consistent with the
 408 estimated split time, and the reduction in population size in East African
 409 populations as they colonised the new habitat. The difference in effective
 410 population size between the southern Shangani population and the eastern
 411 Maasai Mara was previously also found using microsatellite data [28].

412 **Simulations** To validate these findings in conditions with a known SFS,
 413 we ran simulations using `msprime` [31] and `tskit` [32]. Briefly, we simulated
 414 two populations, which we simply refer to as A and B. Populations A and
 415 B diverged 10 000 generations ago and both have effective populations sizes
 416 of 10 000 individuals, except for a period of 1000 generations after the split,
 417 during which time B went through a bottleneck of size 1000. We simulated 22
 418 independent chromosomes of 10 Mb for a total genome size of 220 Mb, using
 419 a mutation rate of $2.5 \cdot 10^{-8}$ and a uniform recombination rate of $1 \cdot 10^{-8}$.
 420 To explore the consequences of varying sample sizes, we sampled 5, 10, or 20
 421 individuals from the two populations. For each of these three scenarios, we
 422 calculated the true SFS from the resulting genotypes (shown in supplementary
 423 figure 6).

424 Using the true genotypes as input, we simulated the effects of NGS
 425 sequencing with error for both the variable and invariable sites. At every
 426 position in the genome, including the monomorphic sites, we sample $D \sim$
 427 Poisson(λ) bases and introduce errors with a constant rate of $\varepsilon = 0.002$
 428 independently for each base. We calculate genotype likelihoods according to
 429 the GATK model outlined in equations (1) to (3) and output GLF files. Using
 430 these, we create SAF files for A and B with no further filtering using `ANGSD`.
 431 The mean depth λ is set to either 2, 4, or 8 to investigate the performance
 432 of `winsfs` at different sequencing depths. This results in a grid of 3×3
 433 simulated NGS data sets with three different sample sizes and three different
 434 mean depth values.

435 From the simulated SAF files, we ran `winsfs` and `realSFS` as above to
 436 generate the two-dimensional SFS, except for a maximum of 100 epochs.
 437 For each method and each epoch e until convergence, we calculated the
 438 log-likelihood for the corresponding SFS $\hat{\phi}^{(e)}$,

$$\begin{aligned} \log p(\phi \mid \hat{\phi}^{(e)}) &= \log \prod_{j \in \mathcal{J}} \phi_j^{M\hat{\phi}_j^{(e)}} \\ &= \sum_{j \in \mathcal{J}} M\hat{\phi}_j^{(e)} \log \phi_j \end{aligned} \quad (15)$$

439 where ϕ is the observed true SFS and M is the total number of sites. Fig-
 440 ure 5 shows how the log-likelihood evolves over epochs for `winsfs` ($W \in$
 441 $\{100, 250, 500\}$) and `realSFS` for sample sizes $N_k \in \{5, 10, 20\}$ and simulated
 442 mean depths $\lambda \in \{2, 4, 8\}$. We observe that at a mean depth of 2, `winsfs`₁₀₀
 443 outperforms `realSFS` by a significant margin both in terms of speed and the
 444 final log-likelihood. At mean depth 4, the `winsfs` remains much faster and
 445 still achieves meaningfully better log-likelihoods, especially at higher sample
 446 sizes. Finally, at mean depth 8, `winsfs`₁₀₀ still converges 5–10 times faster

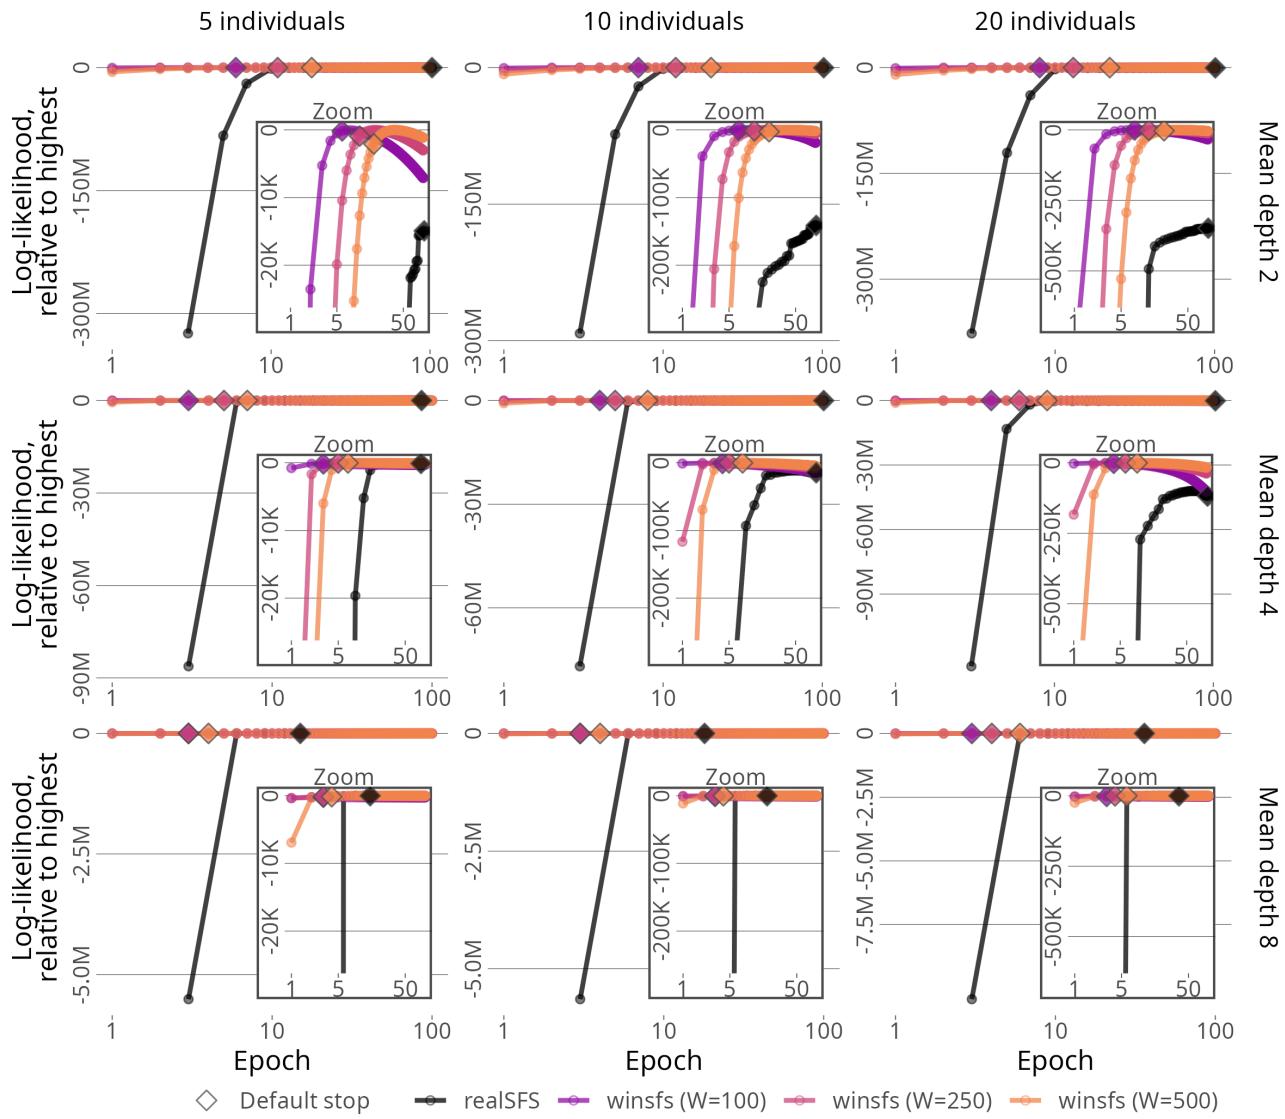


Figure 5: Log-likelihood over epochs of the true observed SFS given the two-dimensional SFS estimated by `winsfs` ($W \in \{100, 250, 500\}$) and `realSFS`. Different simulated scenarios (mean depth 2, 4, or 8; sample size 5, 10, or 20) shown. For each method, the epoch at which the default stopping criterion is triggered is shown. Note that the y -scale varies across sample sizes and depths in order to show the full range of data (main plot) and the difference between `realSFS` and `winsfs` (zoom plot). For each column of plots, corresponding to a simulated sample size, the y -scale in the zoom plot is held constant to allow for comparison across depths.

447 than **realSFS** (measured in epochs), but the methods provide estimates of
448 similar quality.

449 The estimated spectra for **realSFS** and **winsfs₁₀₀** at their default stopping
450 points are shown in supplementary figure 7 and supplementary figure 8 and
451 respectively. These confirm that the spectra on the whole are well-estimated by
452 **winsfs₁₀₀** as compared to the true SFS (supplementary figure 6). Moreover,
453 we again observe that **realSFS** introduces a checkerboard pattern in the
454 low-information part of the spectrum at 2x–4x, which is not present in the
455 true spectrum, and which is not inferred by **winsfs**. The pattern is more
456 pronounced at higher sample sizes. This supports the hypothesis that **realSFS**
457 tends to overfit in situations where many parameters must be inferred with
458 little information.

459 **Peak simulations** The averaging of block estimates in the window EM
460 algorithm appears to induce a certain ‘smoothing’ of the spectrum at low
461 depth. This smoothing effect is implicit in the sense of being nowhere explicitly
462 modelled, and each parameter is estimated independently. Nevertheless, this
463 observation may give rise to a concern that **winsfs**, unlike the maximum
464 likelihood estimate from **realSFS**, might remove true abrupt peaks in the
465 SFS.

466 To investigate, we modified the demographic simulation with sample size 20
467 described above in the following way. In each of seven arbitrarily chosen bins
468 near to the centre of the SFS, we artificially spiked 10 000 counts into the true
469 spectrum after running the demographic simulations (supplementary figure 9).
470 This represents a 30–40-fold increase relative to the original count and the
471 neighbouring cells. Based on this altered spectrum, we simulated sequencing
472 data for depth 2x, 4x, and 8x, created SAF files, and ran **realSFS** and
473 **winsfs₁₀₀** as before. The residuals of the **realSFS** and **winsfs** estimates are
474 shown in supplementary figure 10 and supplementary figure 11, respectively.
475 In this fairly extreme scenario, the spectra inferred by both **winsfs** and
476 **realSFS** appear to have a small but noticeable downwards bias in the peak
477 region at 2x and 4x. However, compared to **realSFS**, **winsfs** has smaller
478 residuals in all scenarios, and the apparent bias is inversely correlated with
479 depth. These results confirm that usage the window EM algorithm does not
480 lead to excess flattening of SFS peaks compared with the maximum likelihood
481 estimate from the standard EM algorithm.

482 **Hyperparameters** The window EM algorithm requires hyperparameter
483 settings for B and W . Moreover, it requires a choice of stopping criterion.
484 For ease of use, the **winsfs** software ships with defaults for these settings,
485 and we briefly describe these.

486 We expect that the choice of B is less important than the term W/B ,
487 which governs the fraction of data that is directly considered in any one update
488 step. Having analysed input data varying in size from 220 Mb (simulations)
489 to 1.17 Gb (human data), we find that fixing $B = 500$ works fine as a default
490 across a wide range of input sizes. Therefore, the more interesting question is
491 how to set the window size. In theory, there should be a trade-off between

492 speed of convergence and accuracy of results, where lower window size favours
493 the former and higher window size the latter. However, in practice, based
494 on our results, we have not seen evidence that using $W = 500$ over $W = 100$
495 leads to significantly better inference. On the other hand, the lower window
496 size has significantly faster convergence. Based on this, we feel that window
497 size of 100 makes for the best general default. By default, the **winsfs** software
498 uses $B = 500$ blocks and a window size $W = 100$.

499 As for stopping, **winsfs** implements the criterion based differences δ in L'_e
500 (eq. (14)) over successive epochs. Based on the initial analysis of the human
501 and impala data, we chose $\delta = 10^{-4}$ (see supplementary figure 12) as the
502 default value and used the simulations to validate this choice. Figure 5 shows
503 the point at which stopping occurs, which is generally around the maximum
504 log-likelihood as desired.

505 **Streaming** In the main usage mode, pre-calculated SAF likelihoods are
506 read into RAM, as in **realSFS**. However, it is also possible to run **winsfs**
507 while keeping the data on disk and streaming through the intersecting sites
508 in the SAF files. We refer to this as ‘streaming mode’.

509 Since the window EM algorithm requires randomly shuffling the input
510 data, a preparation step is required in which SAF likelihoods are (jointly)
511 shuffled into a new file. We wish to avoid loading the data into RAM in order
512 to perform a shuffle, and we also do not want multiple intermediate writes
513 to disk. To our knowledge, it is not possible to perform a true shuffle of the
514 input data within these constraints. Instead, since we are only interested
515 in shuffling for the purposes of breaking up blocks of LD, we perform a
516 pseudo-shuffle according to the following scheme. We pre-allocate a file with
517 space for exactly M intersecting sites in the input data. This file is then
518 split into S contiguous sections of roughly equal size, and we then assign
519 input site with index $m \in \{1, \dots, M\}$ to position $\lfloor (m+1)/S \rfloor + 1$ in section
520 $(m+1) \% S + 1$, where $\%$ is the remainder operation. That is, the first S
521 sites in the input end up in the first positions of each section, and the next S
522 sites in the input end up in the second positions of each section, and so on.
523 This operation can be performed with constant memory, without intermediate
524 writes to disk, and has the benefit of being reversible.

525 After preparing the pseudo-shuffled file, **winsfs** can be run exactly as
526 in the main mode. To confirm that this pseudo-shuffle is sufficient for the
527 purposes of the window EM algorithm, we ran 10 epochs of **winsfs** in
528 streaming mode for the impala and human data sets in both one and two
529 dimensions. After each epoch, we calculated the log-likelihood of the resulting
530 SFS and compared them to the log-likelihood obtained by running in main
531 mode above. The results are shown in supplementary figure 13 and show that
532 streaming mode yields comparable results to the main, in-RAM usage: the
533 likelihood differs slightly, but is neither systematically better or worse.

534 **Benchmark** To assess its performance characteristics, we benchmarked
535 **winsfs** in both the main mode and streaming mode as well as **realSFS** on
536 the impala data. For each of the three, we ran estimation until convergence,

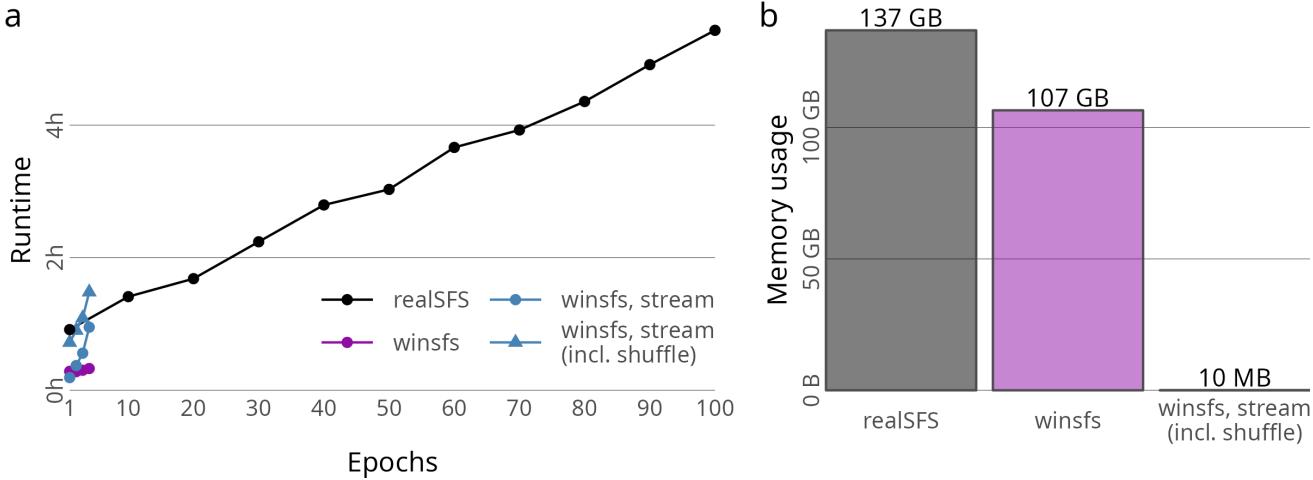


Figure 6: Computational resource usage of `winsfs` and `realSFS` for the joint estimation of the Shangani and Maasai Mara impala populations `winsfs` can be run while loading input data into RAM, or streaming through it on disk. In the latter case, data must be shuffled on disk before hand. (a): Runtime required with 20 threads for various numbers of epochs. Results for `winsfs` are shown for in-memory usage and streaming mode. For streaming modes, times are given with and without the extra time taken to shuffle data on disk before running. (b): Peak memory usage (maximum resident set size).

as well as until various epochs before then, collecting benchmark results using Snakemake [33]. Both `realSFS` and `winsfs` were given 20 cores. Results are shown in figure 6. In terms of run-time, we find that running `winsfs` in RAM is significantly faster than `realSFS` (figure 6a). This is true in part because `winsfs` requires fewer epochs, but also since `winsfs` runs faster than `realSFS` epoch-by-epoch. As expected, when switching `winsfs` to streaming mode, run-time suffers as epochs increase. However, taking the number of epochs required for convergence into account, streaming `winsfs` remains competitive with `realSFS`, even when including the initial overhead to shuffle SAF likelihoods on disk.

Looking at memory consumption, streaming `winsfs` has a trivial peak memory usage of 10 MB, including the initial pseudo-shuffle. In comparison, when reading data into RAM, `realSFS` and `winsfs` require 137 GB and 107 GB, respectively, even on the fairly small impala data set.

The benchmarking results for the one-dimensional Maasai Mara estimation are shown in supplementary figure 14 and support similar conclusions.

4 Discussion

We have presented the window EM algorithm for inferring the SFS from low-depth data, as well as the `winsfs` implementation of this algorithm. The window EM algorithm updates SFS estimates in smaller blocks of sites, and averages these block estimates in larger windows. We have argued that this approach has three related advantages relative to current methods. First, by updating more often, convergence happens one to two orders of magnitude faster. Due to the window averaging, this improvement in convergence times does not occur at the cost of stability. Second, due to the fast convergence, it

562 is feasible to run the window EM algorithm out of memory. This brings the
563 memory requirements of the algorithm from hundreds of gigabytes of RAM to
564 virtually nothing. Third, by optimising over different subsets of the data in
565 each iteration, the algorithm is prevented from overfitting to the input data.
566 In practice, this means we get biologically more plausible spectra.

567 On this last point, it is worth emphasising that while `winsfs` appears to
568 have the effect of smoothing the spectrum in a beneficial way, this smoothing
569 effect is entirely implicit. That is, it is nowhere explicitly modelled that each
570 estimated bin should be similar to neighbouring bins to avoid checkerboard
571 patterns. Rather, the apparent smoothing emerges because `winsfs` mitigates
572 some of the issues with overfitting that may otherwise manifest as a
573 checkerboard pattern. As shown in the simulations, `winsfs` does not remove
574 true peaks in the SFS. In the broader setting of stochastic optimization, window
575 EM is in this way related to forms of Polyak-Ruppert iterate averaging
576 schemes as used in stochastic gradient methods [34, 35], variants of which
577 have also been shown to control variance and induce regularisation [36, 37],
578 similar to what we have observed here.

579 Within the EM literature, window EM is *prima facie* quite similar in
580 spirit to other versions of the stochastic EM algorithm [38–42]. They too
581 work on smaller blocks, and seek some way of controlling stability in how the
582 block estimate $\hat{\psi}$ is incorporated in the overall estimate $\hat{\phi}$. Typically, this
583 involves an update of the form $\gamma_t \hat{\psi} + (1 - \gamma_t) \hat{\phi}$ for some weight γ_t decaying as
584 a function of iteration t . During initial experimentation, we empirically found
585 that such methods tended to increase the noise in the spectrum, rather than
586 reduce it. This problem likely arises because estimating the multidimensional
587 SFS requires estimating many parameters for which very little information is
588 available in any one batch. Therefore, by having an update step involving
589 only the current estimate and a single, small batch of sites, significant noise is
590 introduced in the low-density part of the spectrum. In contrast, the window
591 EM approach still optimises over smaller batches for speed, but actually
592 considers large amounts of data in the update step by summing the entire
593 window of batch estimates, thereby decreasing the noise.

594 For SFS inference specifically, prior work exists to improve estimation
595 for low-depth sequencing data. For example, it has been proposed to ‘band’
596 SAF likelihoods to make estimation scale better in the number of sampled
597 individuals [23, 43]. Briefly, the idea is that at each site, all the mass in the
598 SAF likelihood tends to be concentrated in a small band around the most
599 likely sample frequency, and downstream inference can be adequately carried
600 out by only propagating this band and setting all others to zero. By doing
601 so, run-time and RAM can be saved by simply ignoring all the zero bins
602 outside the chosen band. We note that such ideas are orthogonal to the work
603 presented here, since they are concerned with the representation of the input
604 data, and thereby indirectly modify all downstream optimisation methods.
605 Future work on `winsfs` may involve the ability to run from banded SAF
606 likelihoods. This will be important with large sample sizes, in the hundreds
607 of individuals.

608 Others have focused on the implementation details of the EM algorithm,
609 for instance using GPU acceleration [44]. Such efforts still have the typical

610 high memory requirements, and do not address the overfitting displayed by
611 the standard EM algorithm. Moreover, we find that the presented algorithmic
612 improvements, combined with an efficient implementation, serve to make
613 `winsfs` more than competitive with such efforts in terms of runtime. Indeed,
614 with `winsfs` converging in-memory in less than an hour on genome-scale data,
615 runtime is no longer a significant bottleneck for SFS estimation.

616 We emphasise, however, that the window EM algorithm and `winsfs` are
617 unlikely to yield any meaningful benefits with sequencing data at above
618 around $10 \times - 12 \times$ coverage. With such data, better inference of the SFS will be
619 obtained by estimation directly from genotype calls with appropriate filters.
620 Nevertheless, efficient and robust methods remain important for low-coverage
621 data. This is partly because low-coverage data may sometimes be the only
622 option, for example when working with ancient DNA. Also, such methods
623 allow intentionally sequencing at lower coverage, decreasing the sequencing
624 cost per individual.

625 In addition, we do not expect `winsfs` to perform better than `realSFS`
626 when data is not available for many sites (e.g. < 100 Mb) due to the fact that
627 `winsfs` only uses parts of the available data directly in the final estimation.

628 Finally, improvements in the SFS estimates by `winsfs` are unlikely to
629 be significant for simple summary statistics like θ , F_{ST} , or f -statistics. For
630 such purposes, `winsfs` simply produces results similar to `realSFS`, although
631 much faster. However, as the number of dimensions and samples increase,
632 and as sequencing depth decreases, overfitting will start to influence the low-
633 frequency bins of the spectrum. Where this information is used downstream,
634 `winsfs` will lead to better and more interpretable results, and can potentially
635 help solve commonly known biases in parameter estimates arising from model
636 misspecification [45]. We have seen this in the $\partial a \partial i$ case study, but we
637 believe the same would be true of other popular demographic inference
638 frameworks including `fastsimcoal` [6, 46], `moments` [47], and `momi` [48]. It
639 may also be significant for other methods for complex inference from the
640 multidimensional spectrum, including inference of fitness effects using `fit\partial a \partial i`
641 [49, 50] or introgression using `DFS` [51], though we have not explored these
642 methods.

643 5 Code and data availability

644 The human data analysed is part of the 1000 Genomes [25] phase 3 low depth
645 sequencing data. Alignments have been made available by the 1000G project
646 and can be accessed at `ftp.1000genomes.ebi.ac.uk/vol1/ftp/phase3/`
647 . The impala data has been made available via the SRA with accession
648 `PRJNA862915`. Analysis and plotting code, as well as the cleaned data corres-
649 ponding to the final results, is available at `github.com/malthesr/window`
650 and the `winsfs` software itself at `github.com/malthesr/winsfs` .

651 6 Acknowledgements

652 AA, CW, and MSR are supported by the Independent Research Fund Den-
653 mark (grant numbers: 8021-00360B, 0135-00211B) and the University of
654 Copenhagen through the Data+ initiative. GGE is supported by the In-
655 dependent Research Fund Denmark (grant number: 8049-00098B). TSK is
656 funded by a Carlsberg Foundation Young Researcher Fellowship awarded by
657 the Carlsberg Foundation in 2019 (CF19-0712).

658 References

- 659 1. Korneliussen, T. S., Moltke, I., Albrechtsen, A. & Nielsen, R. Calculation
660 of Tajima's D and Other Neutrality Test Statistics from Low Depth
661 Next-Generation Sequencing Data. *BMC Bioinformatics* **14** (2013).
- 662 2. Bhatia, G., Patterson, N., Sankararaman, S. & Price, A. L. Estimating
663 and Interpreting F_{ST} : The Impact of Rare Variants. *Genome Research*
664 **23**, 1514–1521 (2013).
- 665 3. Peter, B. M. Admixture, Population Structure, and F-Statistics. *Genetics*
666 **202**, 1485–1501 (2016).
- 667 4. Marth, G. T., Czabarka, E., Murvai, J. & Sherry, S. T. The Allele
668 Frequency Spectrum in Genome-Wide Human Variation Data Reveals
669 Signals of Differential Demographic History in Three Large World Pop-
670 ulations. *Genetics* **166**, 351–372 (2004).
- 671 5. Gutenkunst, R. N., Hernandez, R. D., Williamson, S. H. & Bustamante,
672 C. D. Inferring the Joint Demographic History of Multiple Populations
673 from Multidimensional SNP Frequency Data. *PLoS Genetics* **5** (ed
674 McVean, G.) e1000695 (2009).
- 675 6. Excoffier, L., Dupanloup, I., Huerta-Sánchez, E., Sousa, V. C. & Foll, M.
676 Robust Demographic Inference from Genomic and SNP Data. *PLoS
677 Genetics* **9** (ed Akey, J. M.) e1003905 (2013).
- 678 7. Tajima, F. Statistical Method for Testing the Neutral Mutation Hypo-
679 thesis by DNA Polymorphism. *Genetics* **123**, 585–595 (1989).
- 680 8. Fay, J. C. & Wu, C.-I. Hitchhiking Under Positive Darwinian Selection.
681 *Genetics* **155**, 1405–1413 (2000).
- 682 9. Nielsen, R. *et al.* A Scan for Positively Selected Genes in the Genomes
683 of Humans and Chimpanzees. *PLoS Biology* **3** (ed Tyler-Smith, C.) e170
684 (2005).
- 685 10. Nielsen, R., Paul, J. S., Albrechtsen, A. & Song, Y. S. Genotype and
686 SNP Calling from Next-Generation Sequencing Data. *Nature Reviews
687 Genetics* **12**, 443–451 (2011).
- 688 11. Han, E., Sinsheimer, J. S. & Novembre, J. Characterizing Bias in Popula-
689 tion Genetic Inferences from Low-Coverage Sequencing Data. *Molecular
690 Biology and Evolution* **31**, 723–735 (2013).
- 691 12. Olalde, I. *et al.* The Genomic History of the Iberian Peninsula Over the
692 Past 8000 Years. *Science* **363**, 1230–1234 (2019).
- 693 13. Margaryan, A. *et al.* Population Genomics of the Viking World. *Nature*
694 **585**, 390–396 (2020).
- 695 14. Van der Valk, T. *et al.* Million-Year-Old DNA Sheds Light on the
696 Genomic History of Mammoths. *Nature* **591**, 265–269 (2021).
- 697 15. Lou, R. N., Jacobs, A., Wilder, A. P. & Therkildsen, N. O. A Beginner's
698 Guide to Low-Coverage Whole Genome Sequencing for Population
699 Genomics. *Molecular Ecology* **30**, 5966–5993 (Aug. 2021).

700 16. Li, H. A Statistical Framework for SNP Calling, Mutation Discovery,
701 Association Mapping and Population Genetical Parameter Estimation
702 from Sequencing Data. *Bioinformatics* **27**, 2987–2993 (2011).

703 17. Sánchez-Barreiro, F. *et al.* Historical Population Declines Prompted
704 Significant Genomic Erosion in the Northern and Southern White Rhinoceros
705 (*Ceratotherium simum*). *Molecular Ecology* **30**, 6355–6369 (2021).

706 18. Li, H. *et al.* The Sequence Alignment/Map format and SAMtools. *Bioin-*
707 *formatics* **25**, 2078–2079 (2009).

708 19. Danecek, P. *et al.* Twelve Years of SAMtools and BCFtools. *GigaScience*
709 **10** (2021).

710 20. McKenna, A. *et al.* The Genome Analysis Toolkit: A MapReduce Frame-
711 work for Analyzing Next-Generation DNA Sequencing Data. *Genome*
712 *Research* **20**, 1297–1303 (2010).

713 21. Korneliussen, T. S., Albrechtsen, A. & Nielsen, R. ANGSD: Analysis of
714 Next Generation Sequencing Data. *BMC Bioinformatics* **15** (2014).

715 22. Nielsen, R., Korneliussen, T., Albrechtsen, A., Li, Y. & Wang, J. SNP
716 Calling, Genotype Calling, and Sample Allele Frequency Estimation
717 from New-Generation Sequencing Data. *PLoS ONE* **7** (ed Awadalla, P.)
718 e37558 (2012).

719 23. Han, E., Sinsheimer, J. S. & Novembre, J. Fast and accurate site fre-
720 quency spectrum estimation from low coverage sequence data. *Bioin-*
721 *formatics* **31**, 720–727 (2014).

722 24. Varadhan, R. & Roland, C. Simple and Globally Convergent Methods
723 for Accelerating the Convergence of Any EM Algorithm. *Scandinavian*
724 *Journal of Statistics* **35**, 335–353 (2008).

725 25. The 1000 Genomes Project Consortium. A Global Reference for Human
726 Genetic Variation. *Nature* **526**, 68–74 (2015).

727 26. Meisner, J. & Albrechtsen, A. Testing for Hardy–Weinberg Equilibrium
728 in Structured Populations Using Genotype or Low-Depth Next Gen-
729 eration Sequencing Data. *Molecular Ecology Resources* **19**, 1144–1152
730 (June 2019).

731 27. Portik, D. M. *et al.* Evaluating mechanisms of diversification in a Guineo-
732 Congolian tropical forest frog using demographic model selection. *Mol-
733 ecular Ecology* **26**, 5245–5263 (Aug. 2017).

734 28. Lorenzen, E. D., Arctander, P. & Siegmund, H. R. Regional Genetic
735 Structuring and Evolutionary History of the Impala *Aepyceros melampus*.
736 *Journal of Heredity* **97**, 119–132 (2006).

737 29. Lorenzen, E. D., Heller, R. & Siegmund, H. R. Comparative phylogeo-
738 graphy of African savannah ungulates. *Molecular Ecology* **21**, 3656–3670
739 (2012).

740 30. Faith, J. T., Tryon, C. A., Peppe, D. J., Beverly, E. J. & Blegen,
741 N. Biogeographic and Evolutionary Implications of an Extinct Late
742 Pleistocene Impala from the Lake Victoria Basin, Kenya. *Journal of*
743 *Mammalian Evolution* **21**, 213–222 (2013).

744 31. Baumdicker, F. *et al.* Efficient ancestry and mutation simulation with
745 msprime 1.0. *Genetics* **220** (2021).

746 32. Kelleher, J., Thornton, K. R., Ashander, J. & Ralph, P. L. Efficient
747 pedigree recording for fast population genetics simulation. *PLOS Com-
748 putational Biology* **14**, e1006581 (2018).

749 33. Koster, J. & Rahmann, S. Snakemake—A Scalable Bioinformatics Work-
750 flow Engine. *Bioinformatics* **28**, 2520–2522 (Aug. 2012).

751 34. Ruppert, D. *Efficient Estimations from a Slowly Convergent Robbins-
752 Monro Process* tech. rep. (Cornell University, 1988).

753 35. Polyak, B. T. & Juditsky, A. B. Acceleration of Stochastic Approximation
754 by Averaging. *SIAM Journal on Control and Optimization* **30**, 838–855
755 (1992).

756 36. Jain, P. *et al.* A Markov Chain Theory Approach to Characterizing the
757 Minimax Optimality of Stochastic Gradient Descent (for Least Squares)
758 (2018).

759 37. Neu, G. & Rosasco, L. *Iterate Averaging as Regularization for Stochastic
760 Gradient Descent* in *Conference On Learning Theory* (2018), 3222–3242.

761 38. Neal, R. M. & Hinton, G. E. *A View of the EM Algorithm that Justifies
762 Incremental, Sparse, and Other Variants* in *Learning in Graphical Models*
763 355–368 (Springer Netherlands, 1998).

764 39. Sato, M.-a. & Ishii, S. On-line EM Algorithm for the Normalized Gaus-
765 sian Network. *Neural Computation* **12**, 407–432 (2000).

766 40. Cappé, O. & Moulines, E. On-line Expectation-Maximization Algorithm
767 for Latent Data Models. *Journal of the Royal Statistical Society: Series
768 B (Statistical Methodology)* **71**, 593–613 (2009).

769 41. Liang, P. & Klein, D. *Online EM for Unsupervised Models* in *Proceed-
770 ings of Human Language Technologies: The 2009 Annual Conference
771 of the North American Chapter of the Association for Computational
772 Linguistics* (2009), 611–619.

773 42. Chen, J., Zhu, J., Teh, Y. W. & Zhang, T. *Stochastic Expectation Maxi-
774 mization with Variance Reduction* in *Proceedings of the 32nd International
775 Conference on Neural Information Processing Systems* **31** (2018), 7967–
776 7977.

777 43. Mas-Sandoval, A. *et al.* Fast and Accurate Estimation of Multidimen-
778 sional Site Frequency Spectra from Low-Coverage High-Throughput
779 Sequencing Data. *GigaScience* **11** (2022).

780 44. Lu, M., Zhao, J., Luo, Q. & Wang, B. *Accelerating Minor Allele Fre-
781 quency Computation with Graphics Processors* in *Proceedings of the 1st
782 International Workshop on Big Data, Streams and Heterogeneous Source
783 Mining Algorithms, Systems, Programming Models and Applications -
784 BigMine '12* (ACM Press, 2012).

785 45. Momigliano, P., Florin, A.-B. & Merilä, J. Biases in Demographic Model-
786 ling Affect Our Understanding of Recent Divergence. *Molecular Biology
787 and Evolution* (ed Kim, Y.) (Feb. 2021).

788 46. Excoffier, L. & Foll, M. fastsimcoal: a Continuous-Time Coalescent
789 Simulator of Genomic Diversity Under Arbitrarily Complex Evolutionary
790 Scenarios. *Bioinformatics* **27**, 1332–1334 (Mar. 2011).

791 47. Jouganous, J., Long, W., Ragsdale, A. P. & Gravel, S. Inferring the Joint
792 Demographic History of Multiple Populations: Beyond the Diffusion
793 Approximation. *Genetics* **206**, 1549–1567 (July 2017).

794 48. Kamm, J. A., Terhorst, J. & Song, Y. S. Efficient Computation of the
795 Joint Sample Frequency Spectra for Multiple Populations. *Journal of
796 Computational and Graphical Statistics* **26**, 182–194 (Jan. 2017).

797 49. Kim, B. Y., Huber, C. D. & Lohmueller, K. E. Inference of the Distribution
798 of Selection Coefficients for New Nonsynonymous Mutations Using
799 Large Samples. *Genetics* **206**, 345–361 (May 2017).

800 50. Huang, X. *et al.* Inferring Genome-Wide Correlations of Mutation Fitness
801 Effects between Populations. *Molecular Biology and Evolution* **38** (ed
802 Nielsen, R.) 4588–4602 (May 2021).

803 51. Martin, S. H. & Amos, W. Signatures of Introgression across the Allele
804 Frequency Spectrum. *Molecular Biology and Evolution* **38** (ed Harris,
805 K.) 716–726 (Sept. 2020).