Detectability of runs of homozygosity is influenced by analysis parameters and population-

specific demographic history

Avril M. Harder[1], Kenneth B. Kirksey[2], Samarth Mathur[3], Janna R. Willoughby[1*]

1. College of Forestry, Wildlife and Environment, Auburn University, Auburn, Alabama, USA

2. Walker College of Business, Appalachian State University, Boone, North Carolina, USA

3. Department of Evolution, Ecology, and Organismal Biology, The Ohio State University,

Columbus, Ohio, USA

*Author for correspondence

Abstract

Wild populations are increasingly threatened by human-mediated climate change and land use

changes. As populations decline, the probability of inbreeding increases, along with the potential

for negative effects on individual fitness. Detecting and characterizing runs of homozygosity

(ROHs) is a popular strategy for assessing the extent of individual inbreeding present in a

population and can also shed light on the genetic mechanisms contributing to inbreeding

depression. Here, we analyze simulated and empirical data sets to demonstrate the downstream

effects of program selection and long-term demographic history on ROH inference. We also

apply a sensitivity analysis to evaluate the effects of various parameter values on ROH-calling

results and demonstrate its utility for parameter value selection. We show that ROH inferences

can be biased when sequencing depth and the distribution of ROH length is not interpreted in

light of demographic history as well as program-specific tendencies. This is particularly

important for the management of endangered species, as underestimating inbreeding signals in

the genome can substantially undermining conservation initiatives. Based on our observations,

we suggest using a combination of ROH detection tools and ROH length-specific inferences to

generate robust population inferences regarding inbreeding history. We outline these

recommendations for ROH estimation at multiple levels of sequencing effort typical of

conservation genomics studies.

Key words: inbreeding, population genomics, PLINK, BCFtools

Introduction

Climate change and expanding human land use are increasingly partitioning wild populations

into smaller and smaller areas of available and suitable habitat, often leading to declining

populations sizes [1,2]. Decreases in population size can lead to increased inbreeding, which has

been reported to have negative fitness consequences for inbred individuals in many wild

populations [3–6]. When inbreeding depression is sufficiently severe, populations can be

threatened with extirpation; thus, assessing inbreeding extent is crucial for understanding and

mitigating risk in small populations of conservation concern. Prior to widespread application of

whole-genome sequencing strategies to non-model species, genetic estimates of inbreeding were

obtained using allozyme or microsatellite data or inferred from known pedigrees [7–10]. These

studies have been critically important to understanding the genetic dynamics of stable and

shrinking populations and have led to increasing recognition of inbreeding depression's

prevalence and ability to affect wild population persistence [4,11]. However, applying whole

genome sequencing strategies to identify runs of homozygosity (ROHs; genomic regions where

both inherited haplotypes are identical) opens up lines of inquiry previously not accessible via

pedigree- or microsatellite-based studies [12].

The identification and quantification of inbreeding depression, largely a result of

homozygosity of strongly deleterious mutations, is critical to conservation as it guides effective

management strategies for myriad species. For example, the introduction of genetically diverse

individuals into declining populations to reduce the proportion of deleterious alleles present in a

population is expected to reduce genetic load and, as a result, reduce extinction risks. Further,

coupling genomic data with fitness outcomes has shown that longer runs of homozygosity

(ROHs) are associated with higher mutational loads [13,14] and ROH analyses in historical and

extant populations have identified genetic trajectories leading to species declines and extinctions

[15,16]. Thus, the integration of genetic data via ROH estimation and existing demographic data

can provide important clues about population stressors that support management decision

making, ultimately reducing extinction risks for target species.

Estimating ROHs can provide crucial insights into populations' evolutionary histories,

but these histories can in turn affect which ROH-calling software and combination of parameter

values are most appropriate. For example, the settings best suited for inferring ROHs in a small,

long-isolated population experiencing high levels of inbreeding would not be suitable for

individuals sampled from a large, genetically diverse population because underlying sources of

error in these two scenarios are very different (*e.g.*, differences in ROH length distributions,

numbers of variable sites, expected minor allele frequencies; [17]). While some studies include

comparisons of results from multiple programs or parameter value combinations (*e.g.,* [18–21]),

many more studies rely on default settings and do not explore the effects of varying these

parameter values on their results. This has the potential to be misleading because it is always

impossible to know how close the resulting estimates approximate reality, and yet ROH are an

alluring metric for decision makers based on the phenomenon that they illustrate.

Here, we address the challenge of interpreting ROH population patterns using simulated

and empirical genomic sequencing data to compare ROH identification patterns among different

demographic histories. Here, we focus on whole-genome sequencing data, building previous

studies that have examined ROH inference for data sets with lower marker densities [22–24].

Specifically, we test a wide array of setting combinations for two programs commonly used in

population genomic studies—PLINK and BCFtools/RoH—and, for PLINK, apply a sensitivity

analysis to evaluate the effects of parameter values on ROH inference. Based on these results, we

outline a set of recommendations for ROH estimation for population genomics studies.

Methods

Part I: Simulated data

<u>Data generation and genotype calling</u>

We used SLiM v4.0.1 and modifications of scripts published by Stoffel et al [6] to simulate four

distinct demographic scenarios: (i) a long-term large population, (ii) a long-term small

population, (iii) a bottlenecked population, and (iv) a long-term declining population. Each

simulation was initiated as a population of 10,000 individuals, wherein each individual consisted

of a homologous pair of 30-Mb chromosomes [25,26]. Each population was simulated for 10,000

generations, followed by scenario-specific demographic changes (Fig. 1). The large population

scenario involved an additional 1,000 generations at a population size of 1,000 individuals,

whereas the small population scenario involved 1,000 additional generations at a population size

of 250 individuals. The bottlenecked population persisted with 1,000 individuals for 900

additional generations, followed by a 50-generation interval with 50 individuals (*i.e.*, the

bottleneck) and a final 50-generation interval with 250 individuals. The declining population

continued for 850 generations with a population size of 1,000, followed by three 50-generation

intervals of decreasing population sizes, from 500 to 250 and finally to 50 individuals.

Recombination rate (9.15×10^{-9} per site per generation) and base mutation rate (1.39×10^{-8} per

site per generation) were chosen based on estimates calculated for Tasmanian devils

(*Sarcophilus harrissii*)—the focal species for our empirical tests [27,28]. Each VCF file output

from SLiM was converted to FASTA sequence files using a custom script in R v4.0.3 and a

haploid ancestral sequence comprising a 30-Mbp segment randomly sampled from the

Tasmanian devil genome [29].

98   Using the known genotypes for all individuals, we generated a coordinate file for each

demographic scenario recording the start and end coordinates for all true ROHs $\geq$ 100 kb in

length (*i.e.*, $\geq$ 100 kb of sequential homozygous loci). We imposed this lower limit on ROH

length because ROHs less than 100 kb in length likely originated in a single common ancestor

approximately 500 generations ago (assuming a recombination rate of 1 cM/1 Mb; [30]), and

would not be expected to influence contemporary individual fitness as strongly as more recently

acquired autozygous segments [13]. This threshold is widely applied in population genetics

studies of non-model species [5,31–33], and we follow this convention for all downstream

analyses.

107   We randomly selected 50 individuals from the final generation of each population

simulation and generated FASTQ read files from each of the two FASTA files representing

homologous chromosomes using ART (version MountRainier-2016-06-05) [34]. We simulated

150-bp paired-end reads using the HiSeq 2500 error model to a depth of 50X per individual (*i.e.*,

25X per homologous chromosome). Each FASTQ file was quality-checked using FASTQC

v0.11.9 [35]. We aligned reads to the ancestral sequence using the BWA-MEM algorithm

implemented in BWA v2.0 and downsampled the resulting BAM files using SAMtools v1.17 to

simulate four additional levels of coverage per individual: 5X, 10X, 15X, and 30X [36,37].

Fig. 1. A) Diagrams describing population size changes for each of the simulated demographic
scenarios. Each diagram shows the final 1,000 generations of each scenario, which was preceded
in each by a burn-in of 10,000 generations with a population size of 10,000 diploid individuals.
Population sizes for each interval are noted for each scenario. B,D,F,H) True overall F_{ROH}
frequencies. C,E,G,I) Length bin-specific true F_{ROH} values; horizontal lines correspond to bin
median values.

For each sorted BAM file, we called genotypes using the 'HaplotypeCaller' algorithm in

Genomic Variant Call Format (GVCF) mode as implemented in GATK v4.1.9.0 [38]. For each

level of coverage, individual GVCF files were combined using 'CombineGVCFs' and genotyped

using 'GenotypeGVCFs'. We applied 'VariantFiltration' to these VCF files in GATK to flag

SNPs with low variant confidence (QualByDepth < 2), exhibiting strand bias (FisherStrand >

40), or with low mapping quality (RMSMappingQuality < 20). Finally, SNPs failing these filters

and indels were removed using 'SelectVariants.' Although population history and genetic dataset

characteristics can influence SNP filtering decisions, we chose to apply consistent filtering

across all demographic scenarios to instead focus on how changing software parameters affect

ROH inference. The effects of various SNP-filtering strategies on ROH inference have been

addressed elsewhere [24,39].

<u>ROH calling: hidden Markov model approach (BCFtools)</u>

We applied the same ROH calling approaches to all multisample VCF files produced from the

simulated data set using two of the programs most commonly applied to non-model species.

First, we tested an extension of the BCFtools software package, BCFtools/RoH v1.17 [40]. This

program uses a hidden Markov model to detect regions of autozygosity, requiring only a VCF

file for all samples, population allele frequency information, and an optional recombination map.

Because additional genetic information is not likely to be available for many wild populations,

we relied on allele frequencies calculated from each of our sample sets. The main decision faced

when running BCFtools/RoH is whether to estimate autozygous regions using called genotypes

or genotype likelihood values. We tested the effects of this decision on ROH estimation by either

including the *--GTs-only* setting to limit inference based on genotypes (hereafter, BCFtools

Genotypes) or omitting it and allowing genotype likelihood values to be considered (hereafter,

BCFtools Likelihoods) (Table 1). There are also two HMM options that can be set: *--hw-to-az*,

which sets the transition probability of Hardy-Weinberg to autozygous state, and *--az-to-hw*,

which sets the transition probability of autozygous to Hardy-Weinberg state. The default values

for these two probabilities are 6.7 x 10^{-8} and 5 x 10^{-9}, respectively. However, BCFtools/RoH can

be run in Viterbi training mode (*--viterbi-training* option) to estimate these two probabilities for

a specific data set. In total for BCFtools/RoH, we applied both methods, Genotypes and

Likelihoods, using default transition probabilities and population-specific probabilities calculated

using the *--viterbi-training* option and a custom R script (available in the GitHub repository for

this project) to each demographic scenario and coverage level.

Table 1. Parameter values applied during ROH calling for both simulated and empirical data. For PLINK, a total of 486 combinations
were tested. PLINK default values are underlined. ROH = run of homozygosity.
157

Program	Parameter	Abbreviation	Values	Description
BCFtools/RoH	--GTs-only	--	30	If set, uses genotypes only and ignores likelihood values
	--viterbi-training	--	1×10^{-10}	If set, runs BCFtools/RoH in Viterbi training mode to estimate HMM transition probabilities; value sets convergence threshold
	--hw-to-az	--	6.7×10^{-8}	Transition probability from Hardy-Weinberg to autozygous state
	--az-to-hw	--	5×10^{-9}	Transition probability from autozygous to Hardy-Weinberg state
PLINK	--homozyg-window-het	phwh	0, 1, 2	Number of heterozygous sites allowed within a window; default = 1 heterozygous site
	--homozyg-window-missing	phwm	2, 5, 50	Number of missing calls allowed in a window; default = 5 missing calls
	--homozyg-window-snp	phws	50, 100, 1000	Scanning window length in SNPs; default = 50 SNPs
	--homozyg-density	phzd	50	Minimum density in kb (*i.e.*, maximum inverse density (kb/variant); *e.g.*, to specify minimum 1 SNP per 50 kb, set to 50); default = 50 kb
	--homozyg-gap	phzg	500, 1000	Threshold distance in kb at which to split a ROH into two if two SNPs are too far apart; default = 1000 kb
	--homozyg-window-threshold	phwt	0.01, 0.05, 0.1	Proportion of overlapping windows that must be called homozygous to assign any SNP to a ROH; default = 0.05
	--homozyg-snp	phzs	10, 100, 1000	Minimum number of variants that must be included in a ROH of minimum length --homozyg-kb to report it; default = 100 SNPs
	--homozyg-kb	phzk	100	Required minimum length of sequence (in kb) spanned by number of homozygous sites specified by --homozyg-snp; default = 1000 kb

ROH calling: sliding window approach (PLINK)

We tested a large number of parameter value combinations in PLINK v1.90b6.26 [41,42].

Unlike BCFtools/RoH, PLINK employs a sliding window approach to ROH identification: for

each window placement, SNPs are examined for conformity to the PLINK parameter values

(*e.g.*, fewer than the number of heterozygous or missing calls allowed). It is then determined, for

each SNP, whether a sufficient proportion of windows overlapping that SNP are homozygous

and thus, whether the SNP is determined to be located within in a ROH. PLINK has multiple

parameters that can be set by the user, and we initially tested a total of 486 combinations of six

of these parameters for each level of coverage (see Table 1 for list of parameters, initial values,

and parameter descriptions). To explore the effects of these combinations on ROH inference, we

applied an iterative approach designed by Mathur et al. [43] to explore through visualization how

varying parameter settings influence overall F_{ROH} estimates across samples within each

population and coverage level. For each iteration, we performed four steps:

1. Run PLINK with all possible combinations of different parameters to be tested,

ultimately generating a matrix of parameter values (predictor variables) and inferred

F_{ROH} (response variable) for each sample.

2. Create a linear model for each combination of parameter values ($F_{RO\,H} = a + b_1 x_1 +$

$\cdots + b_n x_n + e$; where b_i = weight of parameter x_i), where the values of parameter x_i are

standardized to 1.

3. Extract standardized rank regression coefficients (SRC) from the linear regression

models using the *sensitivity* package in R and visualize sensitivity indices (SRC_i) to rank

weights of each parameter [44].

4. If $SRC_i \; \square \; 0$ with little individual variation, then set the parameter i to the default value. If

SRC_i is > 0 or < 0, then consider the effect described by SRC_i (*i.e.*, whether increasing

the value of the parameter increases or decreases F_{ROH} and how SRC_i varies with called

F_{ROH}) and either define a new set of parameter values to test or select a value from the

tested set.

We began the first iteration by reading the results from the initial 486 combinations of parameter

values for each demographic scenario into R v4.0.3 [29]. After examining the output figures, we

selected new sets of parameter values to test for each demographic scenario (S2 Table).

<u>Data summarization and statistical analyses</u>

Output files from BCFtools/RoH and the final PLINK runs were read into R for summarization

and statistical analyses. We also read in true ROH data (*i.e.*, start and end coordinates for known

ROHs $\geq$ 100 kb in length) and calculated true F_{ROH} values for each individual. We filtered all

called ROHs to retain ROHs $\geq$ 100 kb in length and calculated inferred F_{ROH} for each

demographic scenario, individual, coverage level, and method. To describe relationships

between true F_{ROH} and called F_{ROH} values, we constructed a linear model for each method and

coverage level with true F_{ROH} as the predictor variable and called F_{ROH} as the response variable.

For each model, we calculated the 95% confidence intervals (CIs) for the slope and *y*-intercept

parameters using the *confint* function in R. To determine whether true and called F_{ROH} values

differed for each model, we tested whether the model's *y*-intercept differed from zero and

whether the slope differed from one (*i.e.*, whether the 95% CIs included zero or one,

respectively). We also used the *y*-intercept and slope parameters to compare the results obtained

from each BCFtools methods when using default vs. Viterbi-trained HMM transition

probabilities within each demographic scenario and, for all three approaches, to determine

whether each method over- or underestimated true F_{ROH} at each coverage level and how the

degree of over- or underestimation changed with increasing true F_{ROH} values and across

demographic scenarios.

To further evaluate the accuracy of each ROH identification method, we also calculated

false negative (*i.e.*, failing to call a ROH present in an individual) and false positive (*i.e.*, calling

a ROH that was not present in an individual) rates for called ROHs. We began by identifying

overlap between true and called ROHs on a per-position basis by summing the number of bases

covered by both the true ROH and called ROH(s). From this information, we calculated (i) the

false negative rate: the total chromosomal length covered by true ROHs but not by called ROHs

divided by the total length of true ROHs; and (ii) the false positive rate: the total chromosomal

length covered by called ROHs but not by true ROHs divided by the total chromosomal length

not covered by true ROHs. For each demographic scenario, method, and level of coverage, we

calculated median false positive and negative rates and compared these medians and the 50%

quantiles between all method and coverage level combinations to provide insight into method-

specific differences in ROH calling errors.

We calculated F_{ROH} for ROHs in four different length bins to explore how ROH

identification methods may differ in their capabilities to accurately call ROHs of different sizes.

We defined length bins as:

(i)   Short ROHs: ≥ 100 kb and < 500 kb in length;

(ii)   Intermediate ROHs: ≥ 500 kb and < 1 Mb in length;

(iii) Long ROHs: ≥ 1 Mb and < 2 Mb in length; and

(iv)   Very long ROHs: ≥ 2 Mb in length.

We examined how F_{ROH} for each bin changed with increasing coverage and also how patterns of

over- and underestimation of F_{ROH} varied with increasing coverage by subtracting true F_{ROH}

from called F_{ROH} for each individual. For each method, level of coverage, and length bin, we

compared mean called F_{ROH} – true F_{ROH} and the 95% CI around these means (estimated using

the *quantiles* function in R), with CIs < 0 indicating underestimation of true F_{ROH} and CIs > 0

indicating overestimation. We further explored relationships between true and called ROHs by

examining how true and called ROHs overlap. We tabulated how many true ROHs each called

ROH overlaps (or contains) and vice-versa for each unique combination of demographic

scenario, ROH detection method, coverage level, and ROH length bin.

Part II: Empirical data

<u>Data curation and genotype calling</u>

To test the effects of program and parameter value selection on identifying ROHs from empirical

data, we analyzed publicly available whole genome sequencing data for a species of

conservation concern, the Tasmanian devil (BioProject PRJNA549794 in NCBI's Sequence

Read Archive; [45]). From the full data set, we selected the 15 individuals from this data set with

the highest number of reads. The accession numbers and relevant metadata for each set of

sequences are provided in S2 Table. Adapters and low-quality bases were trimmed from raw

sequences using Trim Galore v0.6.6 [46], and cleaned reads were mapped to the chromosome-

level mSarHar1.11 *S. harrissii* reference genome (NCBI GenBank accession

GCA_902635505.1) using BWA-MEM [36].

We used Qualimap v2.2.1 to determine mean coverage per individual from each sorted

BAM file [47]. These results were used to calculate the downsampling proportions required to

approximate 5X, 10X, 15X, and 30X coverage for each individual. Following downsampling,

BAM files were processed in the same manner as for the simulated data, with additional SNP

filtering criteria applied in VCFtools v0.1.17 [48], including filtering SNPs within 5 bp of indels

and requiring minor allele frequencies ≥ 0.05 and $< 20\%$ missing data across individuals.

<u>ROH calling and sensitivity analyses</u>

We called ROHs from the final multisample VCF files using the same approaches as for the

simulated data. We called ROHs in two ways, (i) using BCFtools/RoH (*i.e.*, relying on

genotypes or on genotype likelihood values and using default or Viterbi-trained HMM transition

probabilities) and (ii) testing 486 parameter combinations in PLINK at each level of coverage

and visualizing the effect of altering each parameter following the same sensitivity analysis

process described above. Parameter values for all iterations tested are provided in S1 Table with

additional details provided for the empirical data in the Supporting Information.

<u>Data summarization and statistical analysis</u>

Output files from BCFtools/RoH and the final PLINK runs for the empirical data were read into

R for summarization and statistical analyses. Following the approach we used for the simulated

data, we filtered all called ROHs to retain ROHs ≥ 100 kb in length and calculated inferred F_{ROH}

for each individual in each scenario, coverage level, and method. We also calculated F_{ROH} for

ROHs in four different length bins, where length bins were defined as described above for the

simulated data. To compare results across methods, coverage levels, and ROH lengths, we

calculated mean length-specific F_{ROH} values and compared the 95% CIs around these means

among methods and coverage levels.

Results

Part I: Simulated data

<u>Data collection and curation</u>

For each simulated population, all analyses were based on 50 randomly sampled individuals. For

the declining population, extremely low heterozygosity in two individuals precluded estimation

of HMM transition probabilities via Viterbi training, and all downstream analyses for this

demographic scenario were based on the remaining 48 individuals. Mean heterozygosity ($\square$ SD)

varied across scenarios: large population mean = 3.17×10^{-5} $\square$ 1.26×10^{-5}; bottlenecked

population mean = 3.05×10^{-5} $\square$ 5.81×10^{-6}; small population mean = 2.87×10^{-5} $\square$ 1.22×10^{-5};

declining population mean = 2.44×10^{-5} $\square$ 8.81×10^{-6}. Mean F_{ROH} generally increased with

decreasing mean heterozygosity: large population mean = 0.21; bottlenecked population mean =

0.48; small population mean = 0.62; declining population mean = 0.65 (Fig. 1B,D,F,H). For all

demographic scenarios, length-specific F_{ROH} mostly decreased with increasing ROH length with

some notable exceptions in the bottlenecked and declining population scenarios (Fig. 1C,E,G,I).

For these two final sets of samples, a few individuals had high proportions of their genomes

located in very long ROHs (Fig. 1G,I), consistent with very recent inbreeding events in each

population. Following downsampling and SNP filtering, the final mean coverage across all

demographic scenarios was 4.87, 9.77, 14.69, and 29.20, and 44.99 for the 5X, 10X, 15X, 30X,

and 50X VCF files, respectively.

<u>ROH calling</u>

Before comparing results from the three main methods—BCFtools Genotypes, BCFtools

Likelihoods, and PLINK—we examined the effect of using either default or Viterbi-trained

HMM transition probabilities on the results produced by each BCFtools method. Within each

method, we compared overall F_{ROH} values obtained when using default HMM transition

probabilities (default values: *--hw-to-az* = 6.7 x 10^{-8}, *--az-to-hw* = 5 x 10^{-9}) against those

calculated when using probabilities estimated via Viterbi training (mean values across all

groups: *--hw-to-az* = 9.2 x 10^{-7}, *--az-to-hw* = 3.8 x 10^{-6}; S3 Table). (In practice, mean

probabilities were calculated and applied for each demographic scenario and BCFtools method

combination, but note the small magnitude of variation in Viterbi-trained probabilities across all

groups relative to the difference between these probabilities and the program defaults.) The slope

and intercept parameters never significantly differed between the two sets of results across all

demographic scenarios and coverage levels, but for most comparisons, the slope and intercept

parameters of the model built using Viterbi-trained results were closer to 1 and 0, respectively,

than for the default values model (S4 Table; S1-4 Figs.). Using Viterbi-trained transition

probabilities slightly increased false positive rates but also decreased false negative rates (S5-8

Figs.). Because of these slight, if not significant, improvements in overall F_{ROH} estimation

accuracy and false negative rates, we chose to proceed with Viterbi-trained HMM transition

probabilities for both BCFtools methods and present only these results below (we also describe

one additional benefit of Viterbi-trained probabilities below).

We used our simulated data set and linear models to determine whether each approach tends to

over- or underestimate true F_{ROH}. All three methods underestimated F_{ROH}, with all model

intercepts across coverage levels negative and different from zero, except for PLINK when

calling ROHs in the large population at 5X coverage (*i.e.*, no other 95% CIs for intercepts

included zero; Fig. 2; S5 Table). Within each demographic scenario, the three methods produced

similar results with respect to overall F_{ROH} accuracy. For example, within all four scenarios,

intercepts did not differ among methods or coverage levels, with two exceptions at 5X coverage

when using PLINK: intercepts differed between this group and others within the large, declining,

and majority of coverage levels for small populations. Perhaps the most visualizing striking

points of comparison across demographic scenarios are the slope estimates specific to each.

Within each method, demographic scenarios were consistently ordered by slope estimates, with

the declining population having the largest slope estimate (for BCFtools, all slopes > 1 with no

95% CIs including one), followed by the small, bottlenecked, and large populations (for the large

population, all BCFtools slopes < 1, all PLINK slope 95% CIs include one; Fig. 2; S5 Table).

Thus, accuracy of overall F_{ROH} estimates calculated for the declining population varied the most

with respect to true F_{ROH} (*i.e.*, of all demographic scenarios, slope parameter estimates for this

population most greatly differed from one), followed by the large population. With respect to

true F_{ROH}, accuracy went up with increasing true F_{ROH} for the declining population and went

down with increasing true F_{ROH} for the large population.

Fig. 2. The relationship between true and inferred F_{ROH} values depends on inference method and population demographic scenario. Each regression line represents linear model results for a single level of coverage with the shaded areas representing 95% confidence intervals. Each point represents data for a single simulated individual. Dashed line is 1:1 line and *x*- and *y*-axes are consistent within each demographic scenario. Note the differing slopes across demographic scenarios (*e.g.*, among panels A-D) and differing overall accuracies across methods (*e.g.*, differing distances between regression lines and 1:1 line among panels D, H, and L).

We calculated false negative (*i.e.*, failing to call a ROH present in an individual) and

false positive (*i.e.*, calling a ROH that was not present in an individual) rates to further assess

each method's accuracy. With respect to false positive rates, PLINK performed poorly relative to

the other two methods, with false positive rates consistently higher than BCFtools Genotypes

and BCFtools Likelihoods (Fig. 3A). Across all demographic scenarios and coverage levels,

median false positive rates ranged from 6.14×10^{-7} to 0.031 for BCFtools Genotypes, from 4.54

$\times 10^{-8}$ to 0.007 for BCFtools Likelihoods, and from 0.008 to 0.112 for PLINK. For all three

methods, increasing coverage to 10X corresponded to decreasing false positive rates, but rates

did not continue to improve at coverages above 10X (*i.e.*, 50% quantiles around median values

overlap at 10-50X within each method and demographic scenario; results for all coverage levels

provided in S9-10 Figs.). Variation in false positive rates among samples at each coverage level

was smallest for BCFtools Likelihoods, followed by BCFtools Genotypes, with PLINK showing

the greatest variation across samples (summary statistics provided in S6 Table).

Fig. 3. PLINK outperforms BCFtools with respect to false negative rates, but underperforms
with respect to false positive rates. A) False positive (*i.e.*, incorrectly calling a base position as
being located in a ROH) and B) false negative (*i.e.*, failing to identify a base position as being
located in a ROH) rates across demographic scenarios and methods. Horizontal lines indicate
median values and shaded boxes are 50% quantiles. Note difference in scale of *y*-axis between
panels A and B. Both BCFtools approaches outperform PLINK with respect to false positive
rates but the reverse is true for false negative rates. Increasing coverage corresponds to
decreasing false positive rates and to increasing false negative rates. Values displayed for 5X and
50X coverages; data for all coverage levels presented in Fig. S9-10.

Patterns in false negative rates were generally in the opposite direction and magnitude to

those we observed with false positives: both BCFtools methods performed poorly relative to

PLINK, with BCFtools Genotypes producing slightly lower rates (median range across all

demographic scenarios and coverage levels = 0.36-0.74) than BCFtools Likelihoods (median

range = 0.49-0.78; Fig. 3B). PLINK exhibited lower false negative rates than the BCFtools

approaches (median range = 0.32-0.53) and less variation among samples at each coverage level.

364   All three methods produced false negative rates that increased slightly at 10X coverage relative

365   to 5X, with rates leveling off at 10X and higher coverage levels. Examples of false negative and

366   false positive scenarios can be seen in S11 Fig., which illustrates a full chromosome of true and

367   called ROHs for one exemplar individual from each demographic scenario.

368   We also examined how true and called values of F_{ROH} varied for ROHs of different

lengths. For the simulated data, all three methods almost always underestimated the proportion

of the genome located in short ROHs across demographic scenarios, with the 95% CI less than

zero for all tests other than PLINK at 5X coverage (Fig. 4; all levels of coverage presented in

S12-15 Figs.). For longer ROHs, results were variable and dependent on demographic scenario.

For example, all three methods appeared to approach accuracy for longer ROHs in the large

population, but this is likely due to the relative paucity of longer ROHs in this demographic

scenario (Fig. 1). For the bottlenecked and declining populations (wherein some individuals had

very long ROHs), both PLINK and BCFtools Genotypes tended to overestimate F_{ROH} for very

long ROHs. PLINK generally produced the highest overestimates of F_{ROH} and the most variation

across samples of the three approaches, followed by BCFtools Genotypes. Finally, one coverage-

related trend emerged across ROH length categories, methods, and demographic scenarios, with

F_{ROH} estimates calculated at 5X coverage often exceeding estimates calculated at higher

coverage levels. Across all length bins and demographic scenarios, individual estimates of

length-specific F_{ROH} calculated at 5X were greater than those calculated at 10X for 37%, 53%,

and 53% of BCFtools Likelihoods, BCFtools Genotypes, and PLINK estimates, respectively.

Fig. 4. Increasing true ROH length corresponds to increasing detection. Called F_{ROH} − True F_{ROH} by length bin and demographic scenario at 15X (results for all coverage levels presented in Fig. S12-15). For BCFtools Genotypes and PLINK, F_{ROH} for short ROHs is consistently underestimated whereas F_{ROH} for very long ROHs is overestimated when these ROHs are present. BCFtools Likelihoods does not overestimate ROHs in any length bin.

To further investigate how called ROHs correspond to true ROHs, we identified regions

of overlap between true and called ROHs within each individual and at each coverage level

using a unique identifier for each true and called ROH. We found no instances of true ROHs

being split into multiple called ROHs, but multiple true ROHs were often lumped together into a

single called ROH. This pattern held true for all three methods at most coverage levels (Fig. 5;

S16 Fig.). For BCFtools Genotypes and PLINK, increasing coverage did not appear to

ameliorate this problem (*i.e.*, the mean number of true ROHs lumped into a single called ROH

changed very little with increasing coverage; Fig. 5C,D). However, for BCFtools Likelihoods,

the number of true ROHs contained in a single called ROH decreased with increasing coverage,

reaching a 1:1 ratio at 30X. Across methods, the mean number of true ROHs combined into a

400   single called ROH increased with increasing ROH length, except for BCFtools Likelihoods at

401   coverage levels ≥ 30X (S16 Fig.). This lumping can be seen in Fig. 5B.

Fig. 5. All three methods tested combine multiple true ROHs into single called ROHs, with
increasing coverage only providing improvements for BCFtools Likelihoods. A) Diagram
illustrating this lumping issue. B) Examples of this issue at 5X and 50X in a single simulated
individual drawn from the small population demographic scenario. C) Number of true ROHs
combined into a single called ROH for ROHs of varying lengths when called by all three
methods at 5X and (D) at 50X in the small population (results for all coverage levels and
demographic scenarios provided in Fig. S16). Points correspond to mean values and vertical and
horizontal error lines indicate 95% confidence intervals. Dashed horizontal line corresponds to y
= 1 (a 1:1 relationship between numbers of true and called ROHs).

Part II: Empirical data

<u>Genotype and ROH calling results</u>

For the 15 sets of reads we downloaded from NCBI, the mean number of reads per sample was

9.75×10^8. Read mapping rates to the mSarHar1.11 *S. harrissii* reference genome were high,

with an average of 95.4% of reads mapped and properly paired. For the final sets of filtered

SNPs (n = 1,532,598), average depth across samples was 48.43 for the full coverage set (*i.e.*, not

downsampled) and 6.37, 11.84, 16.63, 30.75 for the 5X, 10X, 15X, and 30X downsampled sets,

respectively (S2 Table). Given the modest improvements in F_{ROH} accuracy conferred by using

Viterbi-trained HMM transition probabilities rather than the default values for the simulated

data, we calculated Viterbi-trained values for both BCFtools methods applied to the empirical

data.

423   Across methods, F_{ROH} estimated at 5X coverage was significantly higher than F_{ROH}

estimates at all higher levels of coverage (95% CIs did not overlap, Fig. 6A-C). At 5X coverage,

F_{ROH} estimates differed significantly among all three methods, with the lowest estimate

produced by PLINK and the highest by BCFtools Genotypes. For all higher levels of coverage,

F_{ROH} estimates produced by the two BCFtools methods did not differ from one another but

estimates from both methods differed from those produced by PLINK, with PLINK again

providing lower estimates.

430   When comparing how the three methods estimated length-specific F_{ROH} values, patterns

varied across ROH length categories. For short ROHs, PLINK produced much lower F_{ROH}

estimates then the two BCFtools methods across all coverage levels, with differences among the

three methods significant (*i.e.*, non-overlapping 95% CIs) at 5X coverage (Fig. 6D). For longer

ROHs, results of the three methods were generally concordant, with all methods producing

higher estimates of F_{ROH} at 5X than at higher coverages (Fig. 6D-G). However, PLINK

produced much higher estimates of F_{ROH} for very long ROHs at 5X than the other two methods,

with a mean value more than double those inferred by the BCFtools methods (Fig. 6G).

Fig. 6. When applied to the empirical data set, the three ROH calling methods differ greatly in
their inferences, particularly at 5X coverage. A-C) Overall F_{ROH} and (D-G) length bin F_{ROH}
results for each method and level of coverage, with means and 95% confidence intervals
indicated by points and vertical lines, respectively. Lighter background lines indicate results for
individual samples.

Discussion

In this study, we systematically analyzed the discrepancies that arise in ROH detection when

populations with varied demographic histories are analyzed and highlight effects on downstream

interpretations associated with these differences. This is important because the demographic

scenarios affected the accuracy, with declining populations showing the highest variance in F_{ROH}

estimation relative to true F_{ROH} values. Additionally, ROH detection accuracy varied with

sequencing coverage; improvements were observed up to a 10X coverage threshold beyond

which the benefits plateaued. Our findings emphasize the impact of choosing appropriate

computational parameters and underscore the challenges in standardizing ROH detection

methods values (*e.g.*, Saremi et al., 2019; Grossen et al., 2020; von Seth et al., 2021; Mueller et

al., 2022), highlighting the necessity for robust, systematic comparisons to enhance

reproducibility and reliability in genomic studies.

We demonstrated the exploratory utility of the sensitivity analysis process we followed to

select parameter values for our data (see the Supporting Information for corresponding

information for the empirical data). This process is important because disparate sequencing data

characteristics are likely to require different parameter values, meaning that it may not be

appropriate to use the values we used herein when analyzing other data. For example, studies

that use fewer SNPs (*e.g.*, populations that are less genetically diverse, studies with reduced

sequencing efforts) should test the effects of altering the minimum SNP density required on

ROH inference results. Sensitivity analysis provides a relatively quick and convenient way to

visualize how different parameter values affect F_{ROH} estimates for an entire data set and the

degree of variation in those effects across individuals.

Inference methodology biases F_{ROH} *values*

Between the two BCFtools methods when considering identified ROHs of all lengths, Genotypes

produced more accurate overall F_{ROH} estimates than Likelihoods, with F_{ROH} estimates from

Likelihoods also increasingly diverging from the true F_{ROH} value with increasing true F_{ROH} (Fig.

1A,B). For populations expected to have considerable variation in F_{ROH} among individuals (*e.g.*,

a population that has remained somewhat small for an extended period of time with evidence of

recent immigration), applying the BCFtools Likelihoods approach could result in increasingly

skewed values for the individuals with the highest levels of inbreeding. For example, using the

linear model parameters estimated for 15X coverage, an individual with a true F_{ROH} of 0.10

would be assigned an inferred F_{ROH} of 0.01 (difference = -0.09), whereas an individual with a

true F_{ROH} value of 0.40 would be assigned 0.23 (difference = -0.17). This could be particularly

problematic when dealing with species or populations of conservation concern because the

individuals with the highest true F_{ROH} also have the largest magnitude of error, meaning that

concerning signals of inbreeding could go undetected.

In contrast to the underestimations produced by the BCFtools/RoH methods, the sliding

window approach implemented in PLINK overestimated F_{ROH}. This was particularly evident at

5X coverage where F_{ROH} estimates differ more from their true values than any other method and

coverage level combination in our study (Fig. 1C). However, at coverages above 5X, PLINK

produced better estimates than either BCFtools approach (*i.e.*, in our linear models, intercepts for

PLINK at 10X-50X are closer to zero than for either BCFtools method and 95% CIs for these

parameter estimates do not overlap with any BCFtools intercept 95% CIs). In the context of

endangered species conservation, small overestimations of F_{ROH} may be more desirable than

underestimations because these are likely to be more conservative (*i.e.*, indicating more close

inbreeding than is present in reality) in many situations. Importantly though, as with BCFtools

Likelihoods, F_{ROH} estimates diverged from true F_{ROH} at increasing values of true F_{ROH}.

However, these values diverged at a much lower rate in the PLINK estimates compared to

BCFtools Likelihoods. Again using our simulated data as a model, an individual with a true

F_{ROH} value of 0.40 would be estimated to have an F_{ROH} of 0.46 (difference = 0.06) when

estimated at 10X-50X with PLINK.

For the two BCFtools methods, patterns of underestimation were consistent with these

approaches' high false negative rates and low false positive rates (Fig. 3). Conversely, PLINK

produced higher false positive rates and lower false negative rates than either BCFtools method,

consistent with overestimation of F_{ROH}. In terms of absolute difference between true and called

F_{ROH} values, PLINK outperformed BCFtools at 10X coverage and above, suggesting that PLINK

will often provide the most robust estimate of F_{ROH}. However, at lower coverages (5X-10X),

BCFtools Genotypes could be considered, given that this method produces F_{ROH} estimates closer

to true F_{ROH} than either PLINK or BCFtools Likelihoods. On the other hand, the underestimates

produced by this approach are likely related to the high false negative rates we observed

(especially relative to PLINK), and the appearance of convergence on true F_{ROH} may be due to

length-specific ROH calling rates by this program (see below) and therefore highly variable

across populations. It is important to note that while the trends we describe may be consistent

with some empirical results (*e.g.*, [5]), individual variation in genomic characteristics exerts

strong influence over F_{ROH} inference results as suggested by comparisons between our simulated

and empirical results.

Demographic history influences F_{ROH} *value accuracy*

The relationship between true and inferred F_{ROH} values varied with demographic scenario,

demonstrating the potential for uncertainty when incorporating these inbreeding values into

management action plans. For example, in the large population where true F_{ROH} values are small

(< 0.3), inferred F_{ROH} accuracy decreased with increasing values of true F_{ROH} for the BCFtools

methods, with both tending to underestimate F_{ROH} with increasing true F_{ROH} (Fig. 2A,E). The

opposite trend occurred for the declining population, where true F_{ROH} values exceed 0.9, with all

three methods approaching more accurate F_{ROH} estimates with increasing values of true F_{ROH}

(Fig. 2D,H,L). These contrasting patterns are likely due to different ROH length distributions in

these two demographic scenarios. In the declining population, ROHs in the individuals with the

highest overall F_{ROH} values tend to comprise very long ROHs (Fig. 1I). Consistent

overestimation of F_{ROH} for these very long ROHs by BCFtools Genotypes and PLINK directly

leads to higher inferred F_{ROH} estimates in these individuals. However, individuals with lower

overall F_{ROH} values have a greater proportion of their ROHs in shorter length categories, which

are underestimated by all three methods (Fig. 4D), leading to inferred overall F_{ROH} estimates that

diverge from true F_{ROH} as true F_{ROH} decreases (Fig. 2D,H,L). In the large population, although

some individuals have higher true F_{ROH} values, most ROHs are still short (Fig. 1C), and likely to

be underestimated by all three methods. Therefore, as true F_{ROH} increases, so do the number of

true short ROHs and the cumulative effects of underestimation across methods, leading to

overall F_{ROH} estimates that diverge from true F_{ROH} as true F_{ROH} increases.

Overall, we found that the accuracy of ROH detection significantly varies across

populations with different demographic histories. Our analysis revealed that declining

populations exhibited the highest variance in F_{ROH} estimation relative to true F_{ROH} values. In

contrast, large populations tended to show more stable estimates but with a general tendency for

underestimation of true F_{ROH}, particularly as true F_{ROH} values increased. These findings

highlight the need for an analysis approach that considers demographic history of the population.

For conservationists and researchers working with genomic data from populations known or

suspected to be in decline, it is vital to interpret ROH analysis outputs with the breadth of

demographic scenarios in mind. When demographic history is unknown and cannot be

estimated, the distribution of F_{ROH} present in the dataset could be compared to the distributions

simulated here (Fig. 1B,D,F,H), or to simulated distributions of F_{ROH} with more similar natural

history parameters. Although tedious, these approaches will minimize the effects of

misestimation of F_{ROH}, reducing the risk of exacerbating genetic diversity loses in populations

and species that are most at risk to extinction.

Coverage ≤ 10X strongly influences called ROH lengths

For the empirical data at 5X coverage, relative to higher coverage levels, all methods

consistently produced lower F_{ROH} estimates for short ROHs and higher F_{ROH} estimates for longer

ROHs (Fig. 7). The overcalling of intermediate to very long ROHs at 5X could be related to the

ROH-lumping issue noted in the simulated results, wherein multiple true ROHs are erroneously

called as a single ROH (Fig. 6). While we cannot confirm the accuracy of F_{ROH} inference for the

empirical data, comparisons between results generated at 5X and higher levels of coverage are

consistent with the simulated results, suggesting that these patterns are accurate (Fig. 2). For the

Tasmanian devil samples we analyzed, the results from 5X coverage suggest much more

frequent, recent inbreeding than the results from ≥ 10X coverage, painting a much more dire

demographic scenario than is presented when more coverage is obtained. Therefore, if one of the

goals of a whole-genome sequencing project is to assess recent or historical patterns of

inbreeding from ROH lengths, ~10X coverage appears to be a minimum requirement for

generating robust ROH inferences.

Patterns of under- or overestimation may vary with ROH length distributions

In our simulated and empirical data, we observed patterns indicating that underlying ROH length

distributions influence the patterns of F_{ROH} under- and overestimation. For example, even though

PLINK produced higher F_{ROH} estimates than both BCFtools methods for the simulated data and

PLINK and BCFtools Genotypes produced statistically indistinguishable estimates for the

empirical data (Fig. 2), length-specific F_{ROH} estimates suggest that differences in underlying true

ROH length distributions between the simulated and empirical data may be responsible for the

differences in relative F_{ROH} results we observed. For the simulated data, BCFtools Genotypes

increasingly overestimated F_{ROH} as ROH length increased (Fig. 5, Fig. S2), with increasing

numbers of true ROHs erroneously combined into single called ROHs (Fig. 6). Although we

cannot know the true ROH length distributions for the empirical data, long ROHs were called at

higher frequencies in the empirical data relative to the simulated data. The tendency of BCFtools

Genotypes to overestimate F_{ROH} for long ROHs combined with the presence of more called long

ROHs in our empirical data set may have minimized differences in overall F_{ROH} estimates

between BCFtools Genotypes and PLINK in the empirical results relative to the simulated

results (Fig. 2). Increased frequencies of long ROHs in the empirical data may have also led to

greater differences in F_{ROH} between 5X and 10X across all three methods for the empirical

results compared to the simulated results (Fig. 2). All three methods call significantly more

intermediate to very long ROHs from the empirical data at 5X than at 10X (Fig. 7B-D), and this

may be related to the increased false positive rates we noted at 5X in the simulated data. These

results again illustrate the effects of a population's or individual's actual ROH complement,

which is determined by typically unknown demographic and breeding patterns, on the relative

reliability and utility of ROH identification programs.

Particularly for endangered species with potentially complicated demographic histories,

interpreting ROH patterns in a population may be most accurate when multiple tools are used to

create an integrated picture. For example, comparing overall and length-specific F_{ROH} estimates

between BCFTools/RoH and PLINK can be used to understand the underlying length

distributions; an abundance of shorter ROHs would be indicated by higher overall F_{ROH}

estimates in PLINK compared to BCFtools Genotypes but similar length-specific F_{ROH} patterns,

whereas a ROH complement comprising many longer ROHs would be indicated by similar

overall F_{ROH} estimates between PLINK and BCFtools Genotypes but higher intermediate to very

long F_{ROH} estimates from BCFtools Genotypes related to PLINK. These accurate assessments of

past and ongoing inbreeding could then be used to inform management options, such as

translocations to ameliorate close inbreeding.

Conclusions

Inferring the presence and characteristics of ROHs can shed important light on population

demographic histories, detect inbreeding depression when combined with fitness information,

and even disentangle the mechanisms underlying or loci contributing to inbreeding depression.

However, given the variation in ROH-calling accuracy demonstrated here, we caution against

direct comparisons of F_{ROH} values generated from different data types or sources or using

different inference parameters. Data from disparate studies could be combined and re-analyzed

in a standardized fashion, although special attention should be paid to variation in reference

genome assembly quality for interspecific comparisons [49]. Regardless of the number of data

sets to be analyzed, we strongly recommend that studies relying on ROH inference (i) employ at

least two ROH-calling programs and interpret their results with each method's biases in mind

and (ii) compare multiple parameter value combinations via sensitivity analysis, taking care to

vary parameters of particular relevance to a data set.

Acknowledgments

We thank members of the Willoughby lab for constructive comments on earlier versions of this

manuscript as well as M Kardos and 2 anonymous reviewers.

Funding

This material is based on work supported by the National Science Foundation Postdoctoral

Research Fellowships in Biology Program under Grant No. 2010251 to AMH. This work was

also supported by the U.S. Department of Agriculture, National Institute of Food and

Agriculture, Hatch project 1025651 to JRW.

Data Availability

Code for all bioinformatic analyses available at https://github.com/avril-m-

harder/roh_inference_testing and https://github.com/kennethb22/roh_parameter_project_kk. All

FASTA files for simulated individuals and final VCF files (for simulated and empirical data and

all coverage levels) will be uploaded to a public repository upon manuscript acceptance.

Author Contributions

AMH and JRW conceived the study. AMH, KBK, and SM performed formal data analyses.

AMH wrote the manuscript with input and final approval from KBK, SM, and JRW.

References

1. Diffenbaugh NS, Field CB. Changes in ecologically critical terrestrial climate conditions.
Science. 2013;341: 486–492. doi:10.1126/science.1237123

2. Haddad NM, Brudvig LA, Clobert J, Davies KF, Gonzalez A, Holt RD, et al. Habitat
fragmentation and its lasting impact on Earth's ecosystems. Sci Adv. 2015;1: e1500052.
doi:10.1126/sciadv.1500052

3. Crnokrak P, Roff DA. Inbreeding depression in the wild. Heredity. 1999;83: 260–270.
doi:10.1038/sj.hdy.6885530

4. Keller L, Waller D. Inbreeding effects in wild populations. Trends in Ecology & Evolution.
2002;17: 230–241. doi:10.1016/S0169-5347(02)02489-8

5. Robinson JA, Räikkönen J, Vucetich LM, Vucetich JA, Peterson RO, Lohmueller KE, et al.
Genomic signatures of extensive inbreeding in Isle Royale wolves, a population on the
threshold of extinction. Sci Adv. 2019;5: eaau0757. doi:10.1126/sciadv.aau0757

6. Stoffel MA, Johnston SE, Pilkington JG, Pemberton JM. Genetic architecture and lifetime
dynamics of inbreeding depression in a wild mammal. Nat Commun. 2021;12: 2972.
doi:10.1038/s41467-021-23222-9

7. Gibbs HL, Grant PR. Inbreeding in Darwin's medium ground finches (*Geospiza fortis*).
Evolution. 1989;43: 1273–1284.

8. Liberg O, Andrén H, Pedersen H-C, Sand H, Sejberg D, Wabakken P, et al. Severe
inbreeding depression in a wild wolf (*Canis lupus*) population. Biol Lett. 2005;1: 17–20.
doi:10.1098/rsbl.2004.0266

9. Saccheri I, Kuussaari M, Kankare M, Vikman P, Fortelius W, Hanski I. Inbreeding and
extinction in a butterfly metapopulation. Nature. 1998;392: 491–494. doi:10.1038/33136

10. Slate J, Kruuk LEB, Marshall TC, Pemberton JM, Clutton-Brock TH. Inbreeding depression
influences lifetime breeding success in a wild population of red deer (*Cervus elaphus*). Proc
R Soc Lond B. 2000;267: 1657–1662. doi:10.1098/rspb.2000.1192

11. O'Grady JJ, Brook BW, Reed DH, Ballou JD, Tonkyn DW, Frankham R. Realistic levels of
inbreeding depression strongly affect extinction risk in wild populations. Biological
Conservation. 2006;133: 42–51. doi:10.1016/j.biocon.2006.05.016

12. Kardos M, Taylor HR, Ellegren H, Luikart G, Allendorf FW. Genomics advances the study
of inbreeding depression in the wild. Evol Appl. 2016;9: 1205–1218.
doi:10.1111/eva.12414

13. Stoffel MA, Johnston SE, Pilkington JG, Pemberton JM. Mutation load decreases with
haplotype age in wild Soay sheep. Evolution Letters. 2021;5: 187–195.
doi:10.1002/evl3.229

14. Szpiech ZA, Xu J, Pemberton TJ, Peng W, Zöllner S, Rosenberg NA, et al. Long runs of
homozygosity are enriched for deleterious variation. The American Journal of Human
Genetics. 2013;93: 90–102. doi:10.1016/j.ajhg.2013.05.003

15. Liu S, Westbury MV, Dussex N, Mitchell KJ, Sinding M-HS, Heintzman PD, et al. Ancient
and modern genomes unravel the evolutionary history of the rhinoceros family. Cell. 2021;
S0092867421008916. doi:10.1016/j.cell.2021.07.032

16. Palkopoulou E, Mallick S, Skoglund P, Enk J, Rohland N, Li H, et al. Complete genomes
reveal signatures of demographic and genetic declines in the woolly mammoth. Current
Biology. 2015;25: 1395–1400. doi:10.1016/j.cub.2015.04.007

17. Ceballos FC, Joshi PK, Clark DW, Ramsay M, Wilson JF. Runs of homozygosity: windows
into population history and trait architecture. Nat Rev Genet. 2018;19: 220–234.
doi:10.1038/nrg.2017.109

18. Saremi NF, Supple MA, Byrne A, Cahill JA, Coutinho LL, Dalén L, et al. Puma genomes
from North and South America provide insights into the genomic consequences of
inbreeding. Nat Commun. 2019;10: 4769. doi:10.1038/s41467-019-12741-1

19. Grossen C, Guillaume F, Keller LF, Croll D. Purging of highly deleterious mutations
through severe bottlenecks in Alpine ibex. Nat Commun. 2020;11: 1001.
doi:10.1038/s41467-020-14803-1

20. von Seth J, Dussex N, Díez-del-Molino D, van der Valk T, Kutschera VE, Kierczak M, et al.
Genomic insights into the conservation status of the world's last remaining Sumatran
rhinoceros populations. Nat Commun. 2021;12: 2393. doi:10.1038/s41467-021-22386-8

21. Mueller SA, Prost S, Anders O, Breitenmoser-Würsten C, Kleven O, Klinga P, et al.
Genome-wide diversity loss in reintroduced Eurasian lynx populations urges immediate
conservation management. Biological Conservation. 2022;266: 109442.
doi:10.1016/j.biocon.2021.109442

22. Ceballos FC, Hazelhurst S, Ramsay M. Assessing runs of homozygosity: a comparison of
SNP array and whole genome sequence low coverage data. BMC Genomics. 2018;19: 106.
doi:10.1186/s12864-018-4489-0

23. Duntsch L, Whibley A, Brekke P, Ewen JG, Santure AW. Genomic data of different
resolutions reveal consistent inbreeding estimates but contrasting homozygosity landscapes
for the threatened Aotearoa New Zealand hihi. Mol Ecol. 2021; mec.16068.
doi:10.1111/mec.16068

24. Meyermans R, Gorssen W, Buys N, Janssens S. How to study runs of homozygosity using
PLINK? A guide for analyzing medium density SNP data in livestock and pet species.
BMC Genomics. 2020;21: 94. doi:10.1186/s12864-020-6463-x

25. Haller BC, Messer PW. Evolutionary modeling in SLiM 3 for beginners. Hernandez R,
editor. Molecular Biology and Evolution. 2019;36: 1101–1109.
doi:10.1093/molbev/msy237

26. Haller BC, Messer PW. SLiM 3: Forward genetic simulations beyond the Wright-Fisher
model. Hernandez R, editor. Molecular Biology and Evolution. 2019;36: 632–637.
doi:10.1093/molbev/msy228

27. Feigin CY, Newton AH, Doronina L, Schmitz J, Hipsley CA, Mitchell KJ, et al. Genome of
the Tasmanian tiger provides insights into the evolution and demography of an extinct
marsupial carnivore. Nat Ecol Evol. 2018;2: 182–192. doi:10.1038/s41559-017-0417-y

28. Patton AH, Margres MJ, Stahlke AR, Hendricks S, Lewallen K, Hamede RK, et al.
Contemporary Demographic Reconstruction Methods Are Robust to Genome Assembly
Quality: A Case Study in Tasmanian Devils. Molecular Biology and Evolution. 2019;36:
2906–2921. doi:10.1093/molbev/msz191

29. R Core Team. R: a language and environment for statistical computing. Vienna, Austria: R
Foundation for Statistical Computing; 2020. Available: https://www.R-project.org/

30. Thompson EA. Identity by descent: Variation in meiosis, across genomes, and in
populations. Genetics. 2013;194: 301–326. doi:10.1534/genetics.112.148825

31. Hasselgren M, Dussex N, Seth J, Angerbjörn A, Olsen R, Dalén L, et al. Genomic and
fitness consequences of inbreeding in an endangered carnivore. Mol Ecol. 2021;30: 2790–
2799. doi:10.1111/mec.15943

32. Sánchez-Barreiro F, Gopalakrishnan S, Ramos-Madrigal J, Westbury MV, Manuel M,
Margaryan A, et al. Historical population declines prompted significant genomic erosion in
the northern and southern white rhinoceros (*Ceratotherium simum*). Mol Ecol. 2021;30:
6355–6369. doi:10.1111/mec.16043

33. Xie H-X, Liang X-X, Chen Z-Q, Li W-M, Mi C-R, Li M, et al. Ancient demographics
determine the effectiveness of genetic purging in endangered lizards. Harris K, editor.
Molecular Biology and Evolution. 2022;39: msab359. doi:10.1093/molbev/msab359

34. Huang W, Li L, Myers JR, Marth GT. ART: a next-generation sequencing read simulator.
Bioinformatics. 2012;28: 593–594. doi:10.1093/bioinformatics/btr708

35. Andrews S. FastQC: A quality control tool for high throughput sequence data. 2015.
Available: http://www.bioinformatics.babraham.ac.uk/projects/fastqc/

36. Li H. Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM.
arXiv:13033997 [q-bio]. 2013 [cited 8 May 2017]. Available:
http://arxiv.org/abs/1303.3997

37. Li H, Handsaker B, Wysoker A, Fennell T, Ruan J, Homer N, et al. The sequence
alignment/map format and SAMtools. Bioinformatics. 2009;25: 2078–2079.
doi:10.1093/bioinformatics/btp352

38. McKenna A, Hanna M, Banks E, Sivachenko A, Cibulskis K, Kernytsky A, et al. The
Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA
sequencing data. Genome Research. 2010;20: 1297–1303. doi:10.1101/gr.107524.110

39. Howrigan DP, Simonson MA, Keller MC. Detecting autozygosity through runs of
homozygosity: A comparison of three autozygosity detection algorithms. BMC Genomics.
2011;12: 460. doi:10.1186/1471-2164-12-460

40. Narasimhan V, Danecek P, Scally A, Xue Y, Tyler-Smith C, Durbin R. BCFtools/RoH: a
hidden Markov model approach for detecting autozygosity from next-generation
sequencing data. Bioinformatics. 2016;32: 1749–1751. doi:10.1093/bioinformatics/btw044

41. Chang CC, Chow CC, Tellier LC, Vattikuti S, Purcell SM, Lee JJ. Second-generation
PLINK: rising to the challenge of larger and richer datasets. GigaSci. 2015;4: 7.
doi:10.1186/s13742-015-0047-8

42. Purcell S, Neale B, Todd-Brown K, Thomas L, Ferreira MAR, Bender D, et al. PLINK: A
tool set for whole-genome association and population-based linkage analyses. The
American Journal of Human Genetics. 2007;81: 559–575. doi:10.1086/519795

43. Mathur S, Tomeček JM, Tarango-Arámbula LA, Perez RM, DeWoody JA. An evolutionary
perspective on genetic load in small, isolated populations as informed by whole genome
resequencing and forward-time simulations. Evolution. 2023;77: 690–704.
doi:10.1093/evolut/qpac061

44. Iooss B, Da Veiga S, Janon A, Pujol G. sensitivity: Global Sensitivity Analysis of Model
Outputs. 2021. Available: https://CRAN.R-project.org/package=sensitivity

45. Wright BR, Farquharson KA, McLennan EA, Belov K, Hogg CJ, Grueber CE. A
demonstration of conservation genomics for threatened species management. Mol Ecol
Resour. 2020;20: 1526–1541. doi:10.1111/1755-0998.13211

46. Krueger F. Trim Galore. Cambridge, U.K.: Babraham Institute; 2019.

47. Okonechnikov K, Conesa A, García-Alcalde F. Qualimap 2: advanced multi-sample quality
control for high-throughput sequencing data. Bioinformatics. 2016;32: 292–294.

48. Danecek P, Auton A, Abecasis G, Albers CA, Banks E, DePristo MA, et al. The variant call
format and VCFtools. Bioinformatics. 2011;27: 2156–2158.
doi:10.1093/bioinformatics/btr330

49. Brüniche-Olsen A, Kellner KF, Anderson CJ, DeWoody JA. Runs of homozygosity have
utility in mammalian conservation and evolutionary studies. Conserv Genet. 2018;19:
1295–1307. doi:10.1007/s10592-018-1099-y