Detecting ancient positive selection in humans using extended lineage sorting

Stéphane Peyrégne, Michael James Boyle, Michael Dannemann, and Kay Prüfer

Department of Evolutionary Genetics, Max Planck Institute for Evolutionary Anthropology, 04103 Leipzig, Germany

Natural selection that affected modern humans early in their evolution has likely shaped some of the traits that set present-day humans apart from their closest extinct and living relatives. The ability to detect ancient natural selection in the human genome could provide insights into the molecular basis for these human-specific traits. Here, we introduce a method for detecting ancient selective sweeps by scanning for extended genomic regions where our closest extinct relatives, Neandertals and Denisovans, fall outside of the present-day human variation. Regions that are unusually long indicate the presence of lineages that reached fixation in the human population faster than expected under neutral evolution. Using simulations, we show that the method is able to detect ancient events of positive selection and that it can differentiate those from background selection. Applying our method to the 1000 Genomes data set, we find evidence for ancient selective sweeps favoring regulatory changes and present a list of genomic regions that are predicted to underlie positively selected human specific traits.

[Supplemental material is available for this article.]

Modern humans differ from their closest extinct relatives, Neandertals, in several aspects, including skeletal and skull morphology (Weaver 2009), and may also differ in other traits that are not preserved in the archeological record (Varki et al. 2008; Laland et al. 2010). Natural selection may have played a role in fixing these traits on the modern human lineage. However, the selection events driving the fixation would have been restricted to a specific timeframe, extending from the split between archaic and modern humans ca. 650,000 yr ago to the split of modern human populations from each other around 100,000 yr ago (Prüfer et al. 2014). While methods exist that can be used to scan the genome for the remnants of past or ongoing positive selection (Nielsen et al. 2007; Pybus and Shapiro 2009), current methods have limited power to detect positive selection on the human lineage that acted during this older timeframe (see Sabeti et al. 2006 for a review on detection methods and their timeframes): an unusually high ratio of functional changes to nonfunctional changes, such as the d_N/d_S test, requires millions of years and often multiple events of selection to generate detectable signals (Kryazhimskiy and Plotkin 2008), while unusual patterns of genetic diversity between individuals and populations (e.g., extended homozygosity, Tajima's D, F_{ST}) are most powerful during the selective sweep or shortly after (Sabeti et al. 2006; Oleksyk et al. 2010). A favorable substitution is not expected to leave a mark on linked neutral variation beyond 250,000 yr in humans (Przeworski 2002, 2003).

The genome sequencing of archaic humans (Neandertals and Denisovans) to high coverage (Meyer et al. 2012; Prüfer et al. 2014) has spawned new methods to investigate the genetic basis of modern human traits that are not shared by the archaics (Pääbo 2014). One method, called 3P-CLR, models allele frequency changes before and after the split of two populations using the archaic genomes as an outgroup (Racimo 2016). 3P-CLR outperforms previous methods in the detection of older events of selection (up to 150,000 yr ago) (Fig. 2 from Racimo 2016) but has little power to detect events older than 200,000 yr ago in modern humans. A second method applied an approximate Bayesian computation on patterns of homozygosity and haplotype diversity around alleles that reach fixation (Racimo et al. 2014). Although this approach expands our ability to investigate older time frames, this signal of selection also fades over time, and events of positive selection older than 300 thousand years ago (kya) become undetectable.

Based on a method introduced by Green et al. (2010), Prüfer et al. (2014) presented a hidden Markov model that identifies regions in the genome where the Neandertal and Denisovan individuals fall outside of present-day human variation (i.e., the archaic lineages fall basal compared to all present-day humans) and applied the model to detect selective sweeps on the modern human lineage. Regions that are unusually long are candidates for ancient selective sweeps as variants are likely to have swept rapidly to fixation, dragging along with them large parts of the chromosomes that did not have time to be broken up by recombination. While this method is, in principle, expected to be able to detect events as old as the modern human split from Neandertals and Denisovans, this power was never formally tested, and it has several other shortcomings. First, the method was limited to modern human polymorphisms, ignoring the additional information given by fixed substitutions. Second, the method does not fit parameters to the data but requires these parameters to be estimated through coalescent simulations.

Here, we introduce a refined version of this method, called ELS (Extended Lineage Sorting), that models explicitly the longer regions produced under selection and includes the fixed differences between archaic and modern human genomes as an additional source of information. The ELS method also takes advantage of an expectation-maximization algorithm to estimate the model parameters from the data itself, making it free from assumptions regarding human demographic history.

To evaluate the power of the ELS method to detect ancient selective sweeps, we tested its performance under scenarios of background selection and neutrality. Finally, we present an updated list of candidate regions that likely underwent positive selection on the modern human lineage since the split from the common ancestor with Neandertals and Denisovans.

Results

Selection causes extended lineage sorting between closely related populations

The ancestors of modern humans split from the ancestors of Neandertals and Denisovans between 450,000 and 750,000 yr ago (Prüfer et al. 2014). Because the two newly formed descendant groups sampled the genetic variation from the ancestral population, a derived variant can be shared between some members of both groups, while other individuals show the ancestral variant. At these positions, some lineages from one group share a more recent common ancestor with some lineages in the other group than within the same group (Rosenberg 2002), a phenomenon called incomplete lineage sorting (Fig. 1A).

Eventually, a derived allele may reach fixation as part of a region that has not been unlinked by recombination. In these regions, all descendants will derive from one common ancestor, and any lineage from the other population will constitute an outgroup, i.e., all lineages are sorted. Because of recombination, the human genome is a mosaic of independent evolutionary histories, and the process of lineage sorting is expected to randomly affect regions, until, ultimately, all lineages will be sorted. In the case of modern humans, only a fraction of the regions in the genome are expected to show lineage sorting (Prüfer et al. 2014), and the genome can be partitioned into regions where an archaic lineage falls either within the variation of modern humans (internal region) or outside of the human variation (external region) (Fig. 1B).

While lineage sorting can occur under neutrality, selection on the modern human branch is expected to always lead to external regions as long as the selective sweep finished. In cases where the selective sweep is sufficiently strong, there will not be sufficient time for recombination to break the linkage with neighboring sites and a large region will reach fixation (extended lineage sorting) (Fig. 1C). In contrast, selection on standing variation may fail to generate such large regions, since recombination can act on the haplotype(s) with the prospective advantageous variant before selection sets in. We note that neither demography nor selection on the archaic lineage affect the lineage sorting within modern humans and thus the power to detect selective sweeps.

Expected incomplete lineage sorting among humans to archaics

We used coalescent simulations to determine the incidence and expected length of regions resulting from incomplete lineage
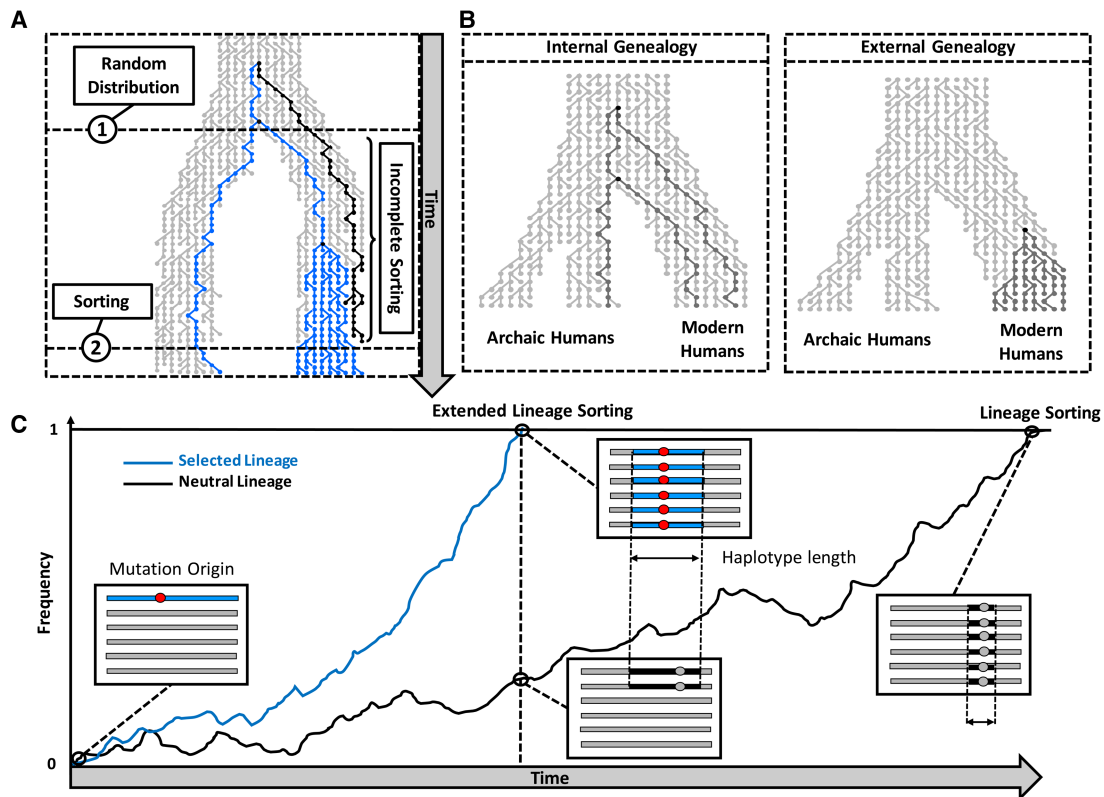
Figure 1. Illustration of the lineage sorting process. (*A*) Effects on the genealogy. The process starts with a random distribution of lineages when the ancestral population splits. The lineage in black is an outgroup to lineages in blue, so that the blue lineages show a closer relationship between populations than to the black lineage (incomplete lineage sorting). When the blue lineages in the top population reach fixation (through a selective sweep for instance), any lineage from the other populations will constitute an outgroup, thereby completing the sorting of lineages. (*B*) Two types of genealogies illustrating the possible relationships between an archaic lineage and modern human lineages. (*C*) Local effects in the genome at different time points. The curves represent the progression of lineage sorting for two independent regions, evolving under neutrality (black curve) and positive selection (blue curve), respectively. Longer fixation times are associated with more recombination so that neutrality produces smaller external regions.

sorting in modern humans. Using a model of human demographic history (Yang et al. 2014), we estimated the fraction of lineage sorting in modern humans in regard to Neandertals and Denisovans. In simulations with 370 African chromosomes, and assuming a uniform recombination rate, ~10% of the archaic genome is more divergent than the time to the most recent common ancestor of all sampled human variation. The length of the external regions is expected to be ~0.0016 centimorgan (cM) (95% CI: 0.001–0.0095 cM; e.g., 1–9.5 kb for a recombination rate of 1 cM/Mb) with the longest regions on the order of 0.02 cM. In contrast, internal regions are expected to be 0.012 cM long (95% CI: 0.0097–0.07 cM).

Minimum strength of selection to produce detectable sweep signals

We investigated the range of selection coefficients that could have led to the fixation of a lineage after the split with the archaic hominins but before the differentiation of genetically modern humans about 100–120 kya (Li and Durbin 2011) by simulating mutations occurring at different times and evolving with different selection coefficients. While the simulations show that completed selective sweeps could have occurred with selection coefficients as low as 0.0005 (Fig. 2A), the length distribution of haplotypes reaching fixation is indistinguishable from neutrality for selection coefficients under 0.001 (Fig. 2B,C). Under neutrality, the average length of external regions was 0.02 cM and remained below 0.03

cM for most simulations with a selection coefficient of 0.001. In contrast, external regions longer than 0.1 cM were observed for selection coefficients above 0.05. Therefore, detectable signals are expected to be biased toward strong events with a selection coefficient larger than 0.001.

## Hidden Markov model to detect extended lineage sorting

To detect regions of extended lineage sorting, we modeled the changes of local genealogies along the genome with a hidden Markov model. We distinguish two types of genealogies, internal or external, depending on whether the archaic lineage falls inside or outside of the human variation, respectively (Fig. 3A). The model includes a third state corresponding to extended lineage sorting, and external regions produced by this state are required to be longer, on average, than those produced by the external state. The three states are inferred from the state of the archaic allele (ancestral or derived) either at a polymorphic position in modern humans or at a position where modern humans carry a fixed derived variant. In the following, we describe the different statistical properties expected for each type of genealogy.

We first considered external regions. At modern human polymorphic sites, the archaic genome is expected to carry the ancestral variant since the derived variant would indicate incomplete lineage sorting. To account for sequencing errors or misassignment of the ancestral state, we allow a probability of 0.01 for carrying the derived allele (see Methods). At sites where the derived allele is fixed, the archaic genome will often carry the derived state, if the fixation event occurred before the split of the archaic from the modern human lineage, or occasionally, the ancestral state, if the fixation event is more recent and occurred after the split.

For internal regions, the archaic is expected to share the derived allele at modern human fixed derived sites but can carry the ancestral allele in our model to accommodate errors, albeit with low probability. In contrast, at sites that are polymorphic in modern humans, the probabilities of observing the ancestral or the derived allele in the archaic genome will depend on the age of the derived variant, with young variants being less likely to be shared compared to older variants. The frequency of the derived variant in the modern human population can be used as a proxy for its age, and the emission probabilities in our model take the modern human derived allele frequency into account (see Methods).

We modeled the transition probabilities between internal and external regions (related to the length of the regions) by exponential distributions. The extended lineage sorting state has the same chance of emitting derived alleles as the other external state but is required to have a larger average length. We used the Baum-Welch algorithm (Durbin et al. 1998), an expectation-
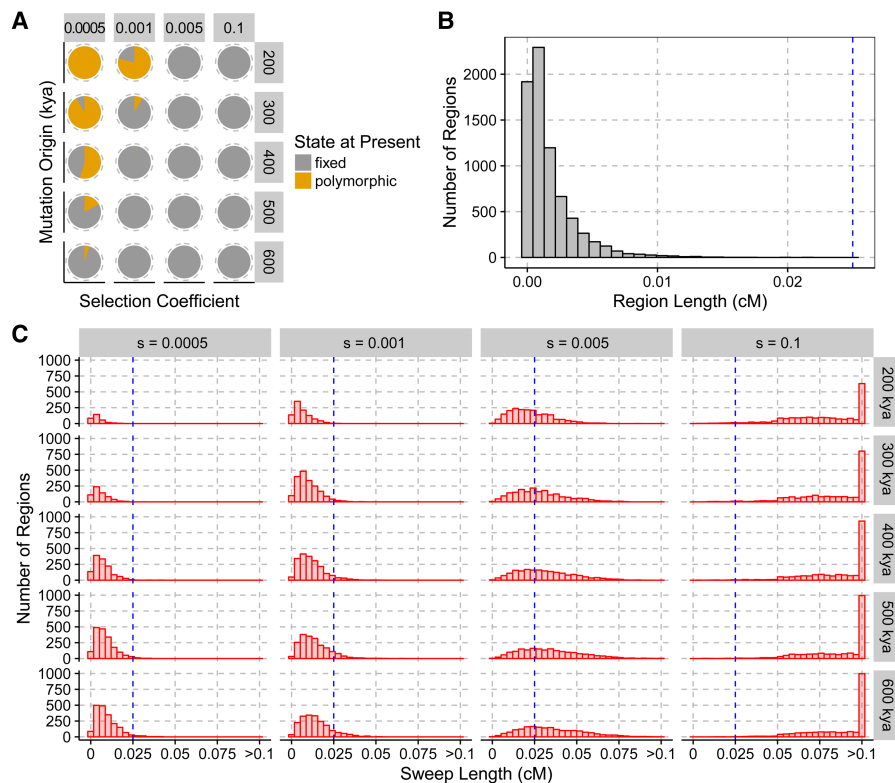
Figure 2. (A) Fraction of selected alleles reaching fixation (gray) or segregating (orange) at present, depending on the strength of selection (columns) and the age of the mutation (rows, in kya) in our simulations. Events for which the selected variant was lost are not shown. (B) Distribution of the genetic length of external regions simulated under neutrality. (C) Distributions of the genetic length of external regions depending on the strength of selection (columns) and age of mutations in kya (rows). The blue line corresponds to the upper limit for the length of external regions produced under neutrality from B.
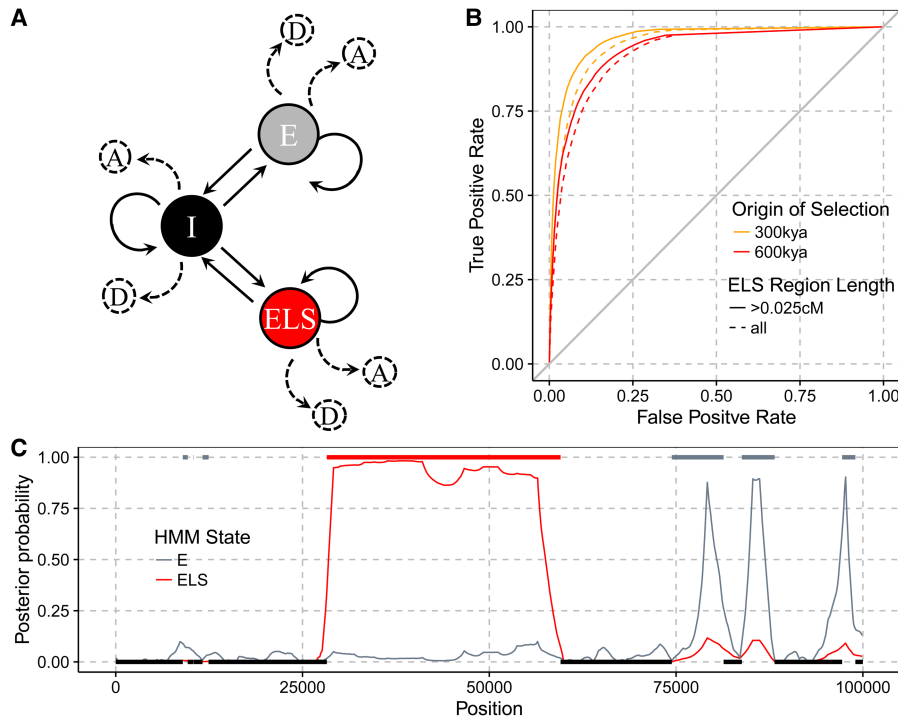
Figure 3. (A) Graphical representation of the Extended Lineage Sorting hidden Markov model. States are depicted by nodes and transitions by edges. Each state emits an archaic allele as either derived, D, or ancestral, A, depending on the type of site in the modern human population (fixed or segregating at a given frequency). States are labeled I for Internal, E for External, and ELS for Extended Lineage Sorting. (B) Receiver operator curves for varying cutoffs on the posterior probability of the ELS state and counting the number of sites in ELS regions that were correctly labeled. All bases labeled ELS outside of simulated ELS regions are considered false positives. Sites in ELS regions with a posterior probability below the cutoff are considered false negatives. (C) Example of the labeling of a simulated ELS region. Horizontal bars indicate true external (top) and internal (bottom) regions. The posterior probability is shown in red for ELS regions and in gray for E regions. The region overlapping position 50,000 (red bar) is caused by a simulated selective sweep.

maximization algorithm, to estimate the emission probabilities and estimate the transition probabilities with a likelihood maximization algorithm.

Accuracy of parameter estimates and inferred genealogies

We first investigated the performance of the parameter inference on simulated data under neutral evolution. We found that the estimated probabilities for encountering ancestral/derived alleles in external and internal regions fit the simulated parameters well (on average, less than ±0.08 from simulated under all tested conditions) (Supplemental Figs. S1, S2), while the estimated length of internal and external regions deviate more from the simulated lengths (around 15% overestimate of the mean length) (Supplemental Fig. S3). However, we found that the model exhibits better accuracy in labeling the correct genealogies with the estimated length parameters compared to the simulated true values (Supplemental Fig. S4). This difference seems to originate from the difficulty in accurately detecting very short external regions or internal regions with very few informative sites. We note that detecting selection is not affected by this problem since we are primarily interested in detecting long external regions. Including fixed differences improves the power to assign the correct genealogies compared to a version of the method without this additional source of information (Supplemental Fig. S4).

We do not expect ELS regions to be detected in our neutral simulations, and indeed we found that either the estimated proportion of ELS converged to zero or the maximum likelihood estimate for the length of ELS and external regions converge to the same value (49% and 51% of simulations, respectively). A likelihood ratio test comparing a model without the ELS state to the full model with the ELS state also showed no significant improvement with the additional state in almost all neutral simulations (only one likelihood ratio test out of 100 simulations showed a significant improvement after Bonferroni correction for multiple testing).

We then evaluated the accuracy of the ELS method to assign the correct genealogy to regions based on sequences obtained through coalescent simulations with selection (Fig. 3B,C). In these simulations, the underlying genealogy at each site along the sequences is known and can be compared to the estimates. To be conservative, we only focused on results with the smallest selection coefficient (s = 0.005) that produces regions long enough to be detectable. In Figure 3B, we show the accuracy for labeling the extended lineage sorting regions dependent on the posterior probability cutoff for the ELS state. The results demonstrate that the model has sufficient power to accurately label sites that experienced selection with a coefficient s ≥ 0.005 and an occurrence of the beneficial mutation as long as 600,000 yr ago.

We also used the simulations of positive selection events (s = 0.005) with two different times at which the beneficial mutation occurred, 300 and 600 kya, to test how often the beneficial simulated variant falls within a detected ELS region (Supplemental Table S1). To put this rate of true positives into perspective, we also counted how many ELS regions did not overlap the selected variant (false positives). A large fraction of selected mutations were detected (87%–92%). However, we also found a substantial fraction of false positive ELS regions (10%–11%). When restricting detected ELS regions to those that are longer than 0.025 cM, we find less than 0.1% false positives compared to 65%–68% true positives. Not all simulated regions with a selection coefficient of 0.005 produce ELS regions of this size, so that the rate of true positives for truly long regions is expected to be higher. For all following analyses, we used this minimal length cutoff of 0.025 cM.

Role of background selection

Background selection is defined as the constant removal of neutral alleles due to linked deleterious mutations (Charlesworth et al. 1993). In regions of the genome that undergo background selection, a fraction of the population will not contribute to subsequent generations, causing a reduced effective population size. As a consequence, remaining neutral alleles can reach fixation faster than

under neutrality, potentially producing unusually long external regions that could be mistaken as signals of positive selection. We investigated the effects of background selection by running forward simulations with parameters that mimic the strength and extent of background selection estimated for the human genome (Messer 2013). While background selection simulations did produce some long outlier regions that fall outside the distribution observed in neutral simulations, most regions are still smaller than regions simulated with positive selection at a conservative selection coefficient of 0.005 (Fig. 4A). Indeed, among the 1160 external regions detected in our simulations of background selection (s = 0.05) (Fig. 4A), only six were labeled as ELS and only three passed the minimal length filter of 0.025 cM.

Candidate regions of positive selection on the human lineage

To identify ancient events of positive selection on the human lineage, we applied the ELS method to African genomes from the 1000 Genomes Project (The 1000 Genomes Project Consortium 2012). We disregarded non-African populations since Neandertal introgression in these populations could mask selective sweeps and lead to false negatives. A model with ELS fits the data significantly better than a model without the ELS state for all chromosomes and for both tested recombination maps (P-value < 1 × 10^{-8}) (Supplemental Table S2).
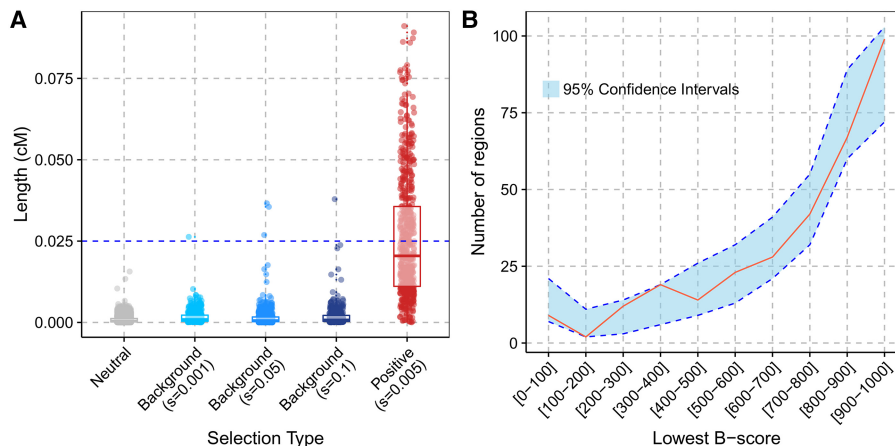
We identified 81 regions of human extended lineage sorting for which both recombination maps support a genetic length greater than 0.025 cM (average length: 0.05 cM). Depending on the recombination map, the longest overlap between the maps is 0.12 (African-American map) or 0.17 (deCODE map) cM long, which is three to four times longer than the longest regions produced under background selection in our simulations. An additional 233 regions are longer than 0.025 cM according to only one recombination map, with 71% of those additional regions showing support for the ELS state using both recombination maps. This suggests that the variation in the candidate set mostly stems from uncertain-

ty about recombination rates. We will refer to the set of 81 regions as the core set (Supplemental File S1) and the set including the 233 putatively selected regions found with just one recombination map as the extended set (314 regions) (Supplemental File S2).

For completeness, we also ran our model on the X Chromosome and identified 12 additional candidates (43 if we consider candidates found with at least one recombination map), applying a more stringent length cutoff of 0.035 cM to account for the stronger effects of random drift on this chromosome (Methods). Interestingly, we also found a significant increase of posterior probabilities for selection within previously reported regions under potential recurrent selective sweeps in apes (Mann-Whitney U one-sided test, P-value < 2.2 × 10^{-16}) (Supplemental Table S3; Dutheil et al. 2015; Nam et al. 2015).

The detected selection candidate regions on the autosomes do not show a decrease in B-scores (McVicker et al. 2009), a local measure of background selection strength, compared with random regions (Wilcoxon rank-sum test comparing the average B-scores with permuted regions, P-value = 0.565, or comparing the lowest B-scores in our regions to permuted regions, P-value = 0.504) (Fig. 4B). This suggests that candidate regions are not primarily generated by strong background selection.

We compared our candidate regions to the top candidates of eight previous scans for selection, including iHS, F_{ST}, XP-CLR, and HKA (Pybus et al. 2014; Cagan et al. 2016). Using the estimated time to the most recent common ancestor among Africans for each identified region/site, we found that our ELS scan identified significantly older events than other screens (Mann-Whitney U tests) (Fig. 5; Supplemental Table S4). We found 23 regions from the core set (detected by both recombination maps) overlapping with candidates from previous scans and 68 for the extended set (detected by at least one recombination map); neither overlap is more than expected at random (P-values are 0.06 and 0.595, respectively). In contrast, our candidate regions overlap candidate regions from 3P-CLR (Racimo 2016) and the ABC approach for detecting ancient selection (Racimo et al. 2014) more often than expected by chance (P-values < 0.05) (Supplemental Table S5).

Biological functions of the candidate regions

Since positive selection acts on advantageous phenotypes that are caused by changes to functional elements in the genome, we would expect that our candidate regions would overlap functional elements in the genome more often than expected.

We first tested this hypothesis by counting the overlap between sweep candidate regions and protein coding genes (Ensembl release 82) (Aken et al. 2016). We find no statistically significant overlap of ELS regions with protein coding genes compared to randomly placed regions of the same size (P-value = 0.671 and 0.124, for core and extended set, respectively) (Fig. 6). Previous work has identified 96 proteins that carry human fixed derived nonsynonymous changes compared to Neandertal and Denisova,
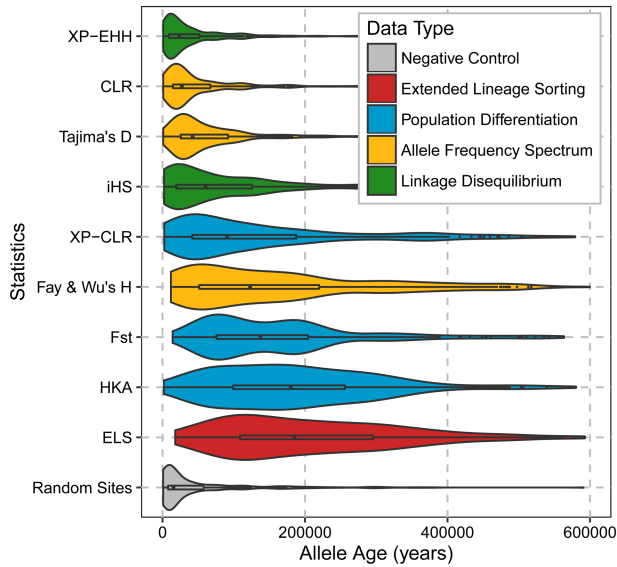

Figure 4. Effects of background selection. (A) Comparison of the length of ELS regions in simulations of different scenarios. For the distribution under background selection, the s parameter corresponds to the average selection coefficient from the gamma distribution (shape parameter of 0.2). We assumed that the deleterious mutations are recessive with dominance coefficient h = 0.1. The horizontal blue line corresponds to the length cutoff applied to the real data. (B) Distribution of B-scores in the candidate sweep regions (red curve) compared to sets of random regions with matching physical lengths (blue area with dotted blue lines indicating the 95% confidence intervals over 1000 random sets of regions). The lowest B-score (i.e., stronger background selection) was chosen when a region overlapped several B-score annotations.

Figure 5. Distributions of estimated ages of the modern human segregating derived variants with the highest frequency in putatively selected regions or the age of the derived variants at sites identified by various genome-wide scans. Our candidate regions are labeled as ELS, for Extended Lineage Sorting; other candidate regions are from Pybus et al. (2014) and Cagan et al. (2016). The color coding indicates the type of signal detected by each method. Ages were estimated by ARGweaver (Rasmussen et al. 2014). We only report events between 0 and 600 kya.

which constitute a particularly interesting subset of potentially functional changes to genes that may have been caused by selective sweeps (Prüfer et al. 2014). We found no overlap between these genes and the core set of sweep candidate regions that were identified by both recombination maps. However, when considering the extended set of sweep candidate regions, 11 regions overlapped such genes: *ADSL, BBIP1, ENTHD1, HERC5, KATNA1, KIF18A, NCOA6, PRDM10, SCAP, SLITRK1*, and *ZNHIT2*. This overlap is significantly larger than expected by chance (only two genes are expected on average; P-value $< 10^{-3}$). In all instances, the candidate regions contained at least one fixed amino acid change. Since fixed changes are part of the information used to infer external regions, it stands to reason that the presence of such a change may bias toward observing an overlap with candidate regions (72/81 core regions and 275/314 regions from the extended set contain fixed changes). However, we note that the overlap with fixed amino acid changes is also significantly larger than the overlap with other fixed changes (963 of 20,347 fixed changes fall within candidate regions from the extended set; binomial P-value $= 0.006$).

Phenotype may also be influenced by regulatory changes that affect gene expressions. Interestingly, we found a significant enrichment for regions overlapping enhancers and promoters (P-value < 0.001 and P-value $= 0.002$, respectively) (see Fig. 6) when considering the extended set of 314 candidate regions. However, this enrichment was not significant for the smaller core set of candidates.

To further investigate the biological function of our regions, we tested for Gene Ontology enrichment in genes within the extended set of regions (Prüfer et al. 2007). No category showed significant enrichment when comparing to randomly placed regions of identical sizes in the genome (see Supplemental Methods). We also assigned genes that overlap our extended data set to tissues in which they show the significantly highest expression (Anders

and Huber 2010) and found again no enrichment (Supplemental Table S6). In an attempt to include potential regulatory changes in the enrichment test, we assigned genes to candidate regions when a region fell upstream of or downstream from a gene (see Supplemental Methods). Although many candidate genes that were annotated in this way were most highly expressed in the brain or the heart (Odds ratio = 2.10 for both tissues), this enrichment is not significant when correcting for gene length and multiple testing (Family-wise error rate = 0.336 and 0.997, respectively) (Supplemental Table S7).

Additional work will be required to investigate the phenotypic consequences of changes in candidate regions for selection. To facilitate this work, we provide an annotated list of fixed or nearly fixed sites on the human lineage that fall within our candidate regions (Supplemental File S3).

Overlap with Neandertal introgression

Introgression from Neandertals and Denisovans into modern humans occurred approximately 37,000 to 86,000 yr ago (Sankararaman et al. 2012, 2016; Fu et al. 2014, 2015). For those advantageous derived variants that arose on the modern human lineage prior to introgression, we would expect that selection may have acted against the re-introduction of the ancestral variant through admixture. We tested whether this selection may have affected the distribution of Neandertal introgressed DNA around fixed changes in candidate sweep regions. Out of a total of 963 fixed derived variants in Africans overlapping the extended set of sweep regions, 240 (25%) show the ancestral allele in non-Africans and show evidence for re-introduction by admixture using a map of Neandertal introgression (Vernot and Akey 2014). This level of Neandertal ancestry is comparable to the genome-wide fraction of out-of-Africa ancestral alleles at African fixed derived sites (~26%; bootstrap P-value = 0.583). We also find no significant reduction in frequency of Neandertal ancestry around candidate substitutions in sweep regions, when comparing one
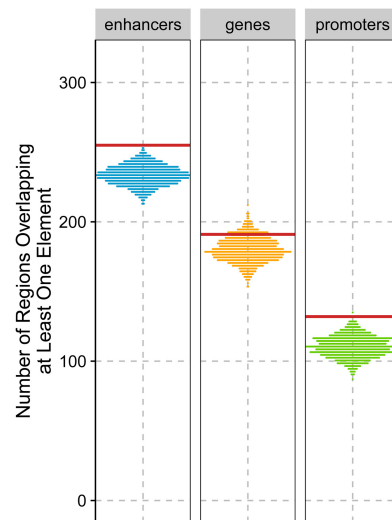
Figure 6. Enrichment for regulatory elements (enhancers, P-value $<$ 0.001; protein-coding genes, P-value = 0.124; and promoters, P-value = 0.002) in the extended set of 314 candidate sweep regions. The distributions were obtained by randomly placing candidate regions in the genome to obtain lists of regions with similar physical lengths. The red lines represent the value observed in the real extended set.

randomly sampled fixed African substitution per region against random regions matched for size and distance to genes (Supplemental Figs. S5, S6).

If selection against the re-introduction of an ancestral variant were very strong, selection may have depleted Neandertal ancestry in a large region surrounding the selected allele. Interestingly, we find some of our sweep candidate regions that fall within the longest deserts of both Neandertal and Denisova ancestry (Supplemental Table S8; Vernot et al. 2016). A significantly high number of the core set of regions fall in these deserts (5/81 regions, P-value = 0.024), while the extended set shows no significant enrichment (9/314 regions, P-value = 0.205).

Discussion

Many genetic changes set modern humans apart from Neandertals and Denisovans, but their functions remain elusive. Most of these changes probably resulted in either no change to the phenotype or to a selectively neutral change. However, in rare instances, selection may have favored changes modifying the appearance, behavior, and abilities of present-day humans. Unfortunately, current methods to identify selection have limited power to detect such old events of positive selection (Przeworski 2002, 2003; Sabeti et al. 2006).

Here, we introduce a hidden Markov model to detect ancient selective sweeps based on a signal of extended lineage sorting. Using simulations, we were able to show that the method can detect older events of selection as long as the selected variant was sufficiently advantageous. The power to detect older events is due to the fact that the method increases in power with the number of mutations that accumulated after the sweep finished. We also showed that background selection can cause false signals and have chosen a minimum length cutoff on candidate regions. While this cutoff reduces the number of false positives due to background selection, we note that this cutoff is expected to exclude bona fide events of positive selection, too.

We applied the ELS method to 185 African genomes, the Altai Neandertal genome, and the Denisovan genome and detected 81 candidate regions of selection when requiring a minimum genetic length supported by two independent recombination maps. The uncertainty in the recombination maps has a large effect on our results, as shown by the much larger number of 314 regions identified by either recombination map. Recombination rates over the genome are known to evolve rapidly (Lesecque et al. 2014), and of particular concern are recent changes in recombination rates that make some regions appear larger in genetic length than they were in the past. By comparing the current recombination rates in our regions to recombination rates in the ancestral population of both chimpanzee and humans (Munch et al. 2014), we identified some candidate regions that may have increased in recombination rates (Supplemental Table S9). However, it is currently impossible to date the change in recombination rates confidently, and these candidate sweeps may post-date the change.

A particular strength of our screen for selective sweeps is the ability to detect older events, as indicated by the estimated power to detect simulated events of positive selection of old age and moderate strength. This sets the ELS method apart from previous approaches that made use of archaic genomes, which were geared toward detecting younger events with an age of less than 300,000 yr ago (Racimo et al. 2014; Racimo 2016). Despite this difference, we found significant overlap between the ELS candidates and the candidates identified by these other approaches, while the overlap with other types of positive selection scans is smaller. Among our candidates, 55 are novel regions (234 if considering the extended set) that were not detected in any of the previous screens, including previous versions of the screen without fixed differences (Supplemental Fig. S7).

While we find no difference in the fraction of genes in selected regions compared to randomly placed regions, we detect an enrichment for enhancers and promoter regions. This result is in agreement with the hypothesis that regulatory changes may play an important role in human-specific phenotypes (King and Wilson 1975; Carroll 2003; Enard et al. 2014), maybe more so than amino acid changes (Hernandez et al. 2011; see also Enard et al. 2014; Racimo et al. 2014). Interestingly, several gene candidates falling within sweep regions play a role in the function and development of the brain. A particularly interesting observation is the potential selection on the genes encoding both the ligand, SLIT2, and its receptor, ROBO2, which reside on Chromosomes 4 and 3, respectively (see Supplemental File S3 for an annotated list of changes in those genes). Members of the Roundabout (*ROBO*) gene family play an important role in guiding developing axons in the nervous system through interactions with the ligands SLITs. SLITs proteins act as attractive or repulsive signals for axons expressing different ROBO receptors. *ROBO2* has been further associated with vocabulary growth (St Pourcain et al. 2014), autism (Suda et al. 2011), and dyslexia (Fisher and DeFries 2002) and is involved in the development of neural circuits related to vocal learning in birds (Wang et al. 2015). Interestingly, *ROBO2* is also in a long desert of both Denisovan and Neandertal ancestry in non-Africans.

We also identified interesting brain-related candidates on the X Chromosome, among them DCX, a protein controlling neuronal migration by regulating the organization and stability of microtubules (Gleeson et al. 1999). Mutations in this gene can have consequences for the expansion and folding of the cerebral cortex, leading to the "double cortex" syndrome in females and "smooth brain" syndrome in males (Gleeson et al. 1998).

We have presented a new approach to detect ancient selective sweeps based on a signal of extended lineage sorting. Applying this approach to modern human data revealed that selection may have acted primarily on regulatory changes. With population level sequencing of nonhuman species becoming more readily available, we anticipate that this approach will help to reveal the targets of ancient selection in other species.

Methods

Data

We used single nucleotide polymorphisms (SNPs) from 185 unrelated Luhya and Yoruba individuals from the 1000 Genomes Project phase I (The 1000 Genomes Project Consortium 2012) together with four ape reference genome assemblies (chimpanzee [panTro3] [Mikkelsen et al. 2005], bonobo [panPan1.1] [Prüfer et al. 2012], gorilla [gorGor3] [Scally et al. 2012], and orangutan [ponAbe2] [Locke et al. 2011]) to compile a list of polymorphic and fixed derived changes in Luhya and Yoruba. Neandertal and Denisova alleles at these positions were extracted from published VCFs (Danecek et al. 2011) using recommended filters (Prüfer et al. 2014; see Supplemental Material for further details). Sites where either Neandertal or Denisova carried a third allele were disregarded.

Genetic distances between these positions were calculated using the African-American (Hinch et al. 2011) and the deCODE (Kong et al. 2010) recombination maps (available in Build 37

from http://www.well.ox.ac.uk/~anjali/). Both maps were chosen since they estimate recombination rates from events that occurred within a few generations before the present. Recombination maps based on older events (i.e., LD-based map) can underestimate recombination rates in regions that underwent recent selective sweeps, potentially masking true signals.

Hidden Markov model

We would like to estimate for each informative position the probabilities for the three possible genealogies, external (E), internal (I), and extended lineage sorting (ELS), given the observed data. Formally and following the notation from Durbin et al. (1998), we calculate $P(\pi_i = k|x)$, where i denotes the position, $k \in \{E, I, ELS\}$ and x is the sequence of observations with the ith observation denoted x_i. With the genetic distance d between consecutive sites and l_k, the average genetic length of a region in state k, we specify the transition probabilities between identical states as $t_{k,k} = e^{-d/l_k}$. Transitions from I to the states ELS and E depend on an additional parameter p, the proportion of transitions from I to ELS, and their probability is given by $t_{I,ELS} = p(1 - e^{-d/l_I})$ and $t_{I,E} = (1 - p)(1 - e^{-d/l_I})$. Lastly, transitions from the two external states to internal have the probability $t_{j,I} = 1 - e^{-d/l_j}$, with $j \in \{E, ELS\}$. By construction, transitions between E and ELS genealogies are not allowed: it would not be possible to detect such transitions, as those two states have the same statistical properties.

The inference further requires the probability for observing an ancestral or derived allele in the archaic at a site i with a derived allele frequency $f_i > 0$ in modern humans (noted x_i) given that the true genealogy is $k \in \{I, E, ELS\}$: $e_k(x_i) = P(x_i| \pi_i = k)$. We assume that $\forall x : e_{ELS}(x) = e_E(x)$, i.e., that both external states give rise to ancestral and derived alleles in the archaic with equal probabilities given the same observation. Since external regions are not expected to give rise to derived sites when the derived allele is segregating in modern humans, the only sources for such an observation can be errors or independent coinciding identical mutations, and we define an error rate for external regions: $\varepsilon_E = e_E(x_i = derived, f_i < 1)$. Similarly fixed derived sites are expected to show the derived allele in the archaics if the local genealogy is internal, and we define an error rate for internal regions: $\varepsilon_I = e_I(x_i = derived, f_i = 1)$.

We compute the posterior probability $P(\pi_i = k \mid x)$ that an observation x_i came from state k given the observed sequence x as:

$$P(\pi_i = k \mid x) = \frac{P(x, \pi_i = k)}{P(x)}.$$

$P(x, \pi_i = k) = f_k(i)b_k(i)$, where $f_k(i) = P(x_1...x_i, \pi_i = k)$ and $b_k(i) = P(x_{i+1}...x_L| \pi_i = k)$ are the output of the Forward and Backward algorithms, respectively (Rabiner 1989; Durbin et al. 1998). $P(x)$ corresponds to the likelihood of the data given our model and was also calculated from the Forward algorithm.

Parameter estimate

We used the Baum-Welch algorithm to estimate all emission probabilities, with the exception of ε_E, the proportion of segregating sites derived in the archaic genome in external regions, due to limited accuracy in the estimates. We set this last parameter to a value of 0.01, a conservative upper limit on contamination and sequencing error in the two high-coverage archaic genomes. The Baum-Welch algorithm was run for a maximum of 40 iterations, and the convergence criteria were set to a log-likelihood maxima difference of less than 10^{-4}.

We estimated the remaining parameters (average lengths of regions and the proportion of transitions to the ELS state) using the derivative free optimization method COBYLA (Powell 1994)

as implemented in the NLopt library (Steven G. Johnson, The NLopt nonlinear-optimization package, http://ab-initio.mit.edu/nlopt) to maximize the log-likelihood values calculated by the Forward algorithm. Convergence was attained in a maximum of 1000 evaluations, and the log-likelihood maximization accuracy was set to 10^{-4}. To test for convergence to local maxima, we ran the algorithm twice with different starting points and used the parameters of the run with the highest likelihood to run the re-estimation algorithm a third time, starting with those parameters. All three runs gave similar results on all chromosomes.

Post–processing

The HMM was executed independently on all chromosomes for both Denisova and Neandertal and using the African-American and deCODE recombination maps. An external region was defined as a stretch of high posterior probabilities ($P \geq 0.7$) for the extended lineage sorting state that was uninterrupted by sites with a low probability ($P \leq 0.1$). The two cutoffs on the posterior probabilities were determined by simulating sequences with positive selection ($s = 0.005$, 500 kya; see below). Sites that were simulated external in both archaics were labeled as 1 and the remaining sites as 0. The HMM was then run on the simulations. By running a grid-search over possible cutoffs (step-sizes of 0.05 for the two parameters) and labeling the HMM output accordingly, we identified the set of chosen parameters by minimizing the root mean square error

$$\sqrt{\frac{\sum_i (t_i - o_i)^2}{n}},$$

with n the number of labeled sites, t_i the true label, and o_i the observed label.

Simulations

We simulated sequences using a model of recent human demography to test the performance of our HMM under different scenarios of neutral evolution, positive selection, or background selection. Each simulation consisted of one chimpanzee chromosome, one chromosome from each archaic hominin, and 370 human chromosomes, matching the 185 Luhya and Yoruba individuals used in our analysis. For all simulations in this study, a constant mutation rate of 1.45×10^{-8} bp^{-1}·generation^{-1}, a constant recombination rate of 1 cM.Mb^{-1}.generation^{-1}, and a generation time of 29 yr were assumed. We used estimates of population sizes from Yang et al. (2014) and population split estimates from Prüfer et al. (2014) as parameters for the simulated demography (Supplemental Information 1, 2).

Neutral simulations were generated with the coalescent simulator *scrm* (Staab et al. 2014) and give a good match to our observed data when plotting derived allele frequency in modern humans against the proportion of derived alleles in the outgroup (Supplemental Fig. S8). Simulations with positive selection were generated with the coalescent simulator *msms* (Ewing and Hermisson 2010), and background selection was explored using forward-in-time simulations generated by SLiM (Messer 2013). Further details on simulation parameters are given in the Supplemental Material.

Age comparison with other scans for selection

To compare our sweep screen with previous scans, we downloaded candidate regions from the 1000G positive selection database (Pybus et al. 2014). Only candidates with a P-value lower than 0.001 were considered. We added to this set of regions

the top reported regions from a HKA scan (Cagan et al. 2016). Allele age estimates were obtained from ARGweaver (Rasmussen et al. 2014).

F_{ST}, iHS, and XP-EHH are site-based statistics which localize sites that may have been selected (Malécot 1948; Wright 1951; Voight et al. 2006; Sabeti et al. 2007), whereas selective scans such as CLR, XP-CLR, Tajima's D, Fay and Wu's H, and HKA identify candidate regions (Hudson et al. 1987; Tajima 1989; Fay and Wu 2000; Kim and Stephan 2002; Chen et al. 2010). In order to compare the age of the selection events, we assumed that the selected variant in candidate regions was the site with the highest frequency. We note that this procedure will underestimate the age of events if the true selected site reached fixation, as is often expected for our method; the comparison is thus conservative.

Annotations

We annotated candidate regions using protein coding genes from Ensembl (release 82), promoters and enhancers mapped by GenoSTAN (Zacher et al. 2016), a measure of background selection (B-scores) (McVicker et al. 2009). Candidate regions were also overlapped with regions previously suggested to have experienced recurrent selective sweeps in apes on the X Chromosome (Dutheil et al. 2015; Nam et al. 2015), regions of Neandertal ancestry (Sankararaman et al. 2014; Vernot and Akey 2014), and long regions devoid of Neandertal and Denisova ancestry (Vernot et al. 2016).

To statistically test the overlap of our regions with these annotations, we randomly placed regions of similar physical sizes in the parts of the genome that passed our quality filters. Quality-filtered regions that were smaller than the longest gap present in our candidate ELS regions were regarded as sufficiently short to not prohibit the placement of regions.

Changes of recombination rates along the human lineage could limit our power to detect selected regions, and we used an ancestral recombination map of the human-chimpanzee ancestor to annotate top candidate regions (Supplemental Table S9; Munch et al. 2014).

Finally, we further characterized fixed or nearly fixed human-specific changes within the candidate regions using annotations of histone marks (enhancers, promoters), eQTLs, transcription factor binding sites, and conservation scores (Supplemental File S3).

Software availability

The software and input files used in this study have been made available through the website http://bioinf.eva.mpg.de/ELS/ and https://github.com/StephanePeyregne/ELS/. A version of the source code is also available as Supplemental Code in the online version of this article.

Acknowledgments

References

The 1000 Genomes Project Consortium. 2012. An integrated map of genetic variation from 1,092 human genomes. *Nature* **491:** 56–65.

Aken BL, Ayling S, Barrell D, Clarke L, Curwen V, Fairley S, Fernandez Banet J, Billis K, García Girón C, Hourlier T, et al. 2016. The Ensembl gene annotation system. *Database (Oxford)* doi: 10.1093/database/baw093.

Anders S, Huber W. 2010. DESeq: differential expression analysis for sequence count data. *Genome Biol* **11:** R106.

Cagan A, Theunert C, Laayouni H, Santpere G, Pybus M, Casals F, Prüfer K, Navarro A, Marques- Bonet T, Bertranpetit J, et al. 2016. Natural selection in the great apes. *Mol Biol Evol* **33:** 3268–3283.

Carroll SB. 2003. Genetics and the making of *Homo sapiens*. *Nature* **422:** 849–857.

Charlesworth B, Morgan MT, Charlesworth D. 1993. The effect of deleterious mutations on neutral molecular variation. *Genetics* **134:** 1289–1303.

Chen H, Patterson N, Reich D. 2010. Population differentiation as a test for selective sweeps. *Genome Res* **20:** 393–402.

Danecek P, Auton A, Abecasis G, Albers CA, Banks E, DePristo MA, Handsaker RE, Lunter G, Marth GT, Sherry ST, et al. 2011. The variant call format and VCFtools. *Bioinformatics* **27:** 2156–2158.

Durbin R, Eddy S, Krogh A, Mitchison G. 1998. *Biological sequence analysis: probabilistic models of proteins and nucleic acids*. Cambridge University Press, Cambridge.

Dutheil JY, Munch K, Nam K, Mailund T, Schierup MH. 2015. Strong selective sweeps on the X chromosome in the human-chimpanzee ancestor explain its low divergence. *PLoS Genet* **11:** e1005451.

Enard D, Messer PW, Petrov DA. 2014. Genome-wide signals of positive selection in human evolution. *Genome Res* **24:** 885–895.

Ewing G, Hermisson J. 2010. MSMS: a coalescent simulation program including recombination, demographic structure and selection at a single locus. *Bioinformatics* **26:** 2064–2065.

Fay JC, Wu CI. 2000. Hitchhiking under positive Darwinian selection. *Genetics* **155:** 1405–1413.

Fisher SE, DeFries JC. 2002. Developmental dyslexia: genetic dissection of a complex cognitive trait. *Nat Rev Neurosci* **3:** 767–780.

Fu Q, Li H, Moorjani P, Jay F, Slepchenko SM, Bondarev AA, Johnson PL, Aximu-Petri A, Prüfer K, de Filippo C, et al. 2014. Genome sequence of a 45,000-year-old modern human from western Siberia. *Nature* **514:** 445–449.

Fu Q, Hajdinjak M, Moldovan OT, Constantin S, Mallick S, Skoglund P, Patterson N, Rohland N, Lazaridis I, Nickel B, et al. 2015. An early modern human from Romania with a recent Neanderthal ancestor. *Nature* **524:** 216–219.

Gleeson JG, Allen KM, Fox JW, Lamperti ED, Berkovic S, Scheffer I, Cooper EC, Dobyns WB, Minnerath SR, Ross ME, et al. 1998. Doublecortin, a brain-specific gene mutated in human X-linked lissencephaly and double cortex syndrome, encodes a putative signaling protein. *Cell* **92:** 63–72.

Gleeson JG, Lin PT, Flanagan LA, Walsh CA. 1999. Doublecortin is a microtubule-associated protein and is expressed widely by migrating neurons. *Neuron* **23:** 257–271.

Green RE, Krause J, Briggs AW, Maricic T, Stenzel U, Kircher M, Patterson N, Li H, Zhai W, Fritz MH, et al. 2010. A draft sequence of the Neandertal genome. *Science* **328:** 710–722.

Hernandez RD, Kelley JL, Elyashiv E, Melton SC, Auton A, McVean G, 1000 Genomes Project, Sella G, Przeworski M. 2011. Classic selective sweeps were rare in recent human evolution. *Science* **331:** 920–924.

Hinch AG, Tandon A, Patterson N, Song Y, Rohland N, Palmer CD, Chen GK, Wang K, Buxbaum SG, Akylbekova EL, et al. 2011. The landscape of recombination in African Americans. *Nature* **476:** 170–175.

Hudson RR, Kreitman M, Aguadé M. 1987. A test of neutral molecular evolution based on nucleotide data. *Genetics* **116:** 153–159.

Kim Y, Stephan W. 2002. Detecting a local signature of genetic hitchhiking along a recombining chromosome. *Genetics* **160:** 765–777.

King MC, Wilson AC. 1975. Evolution at two levels in humans and chimpanzees. *Science* **188:** 107–116.

Kong A, Thorleifsson G, Gudbjartsson DF, Masson G, Sigurdsson A, Jonasdottir A, Walters GB, Jonasdottir A, Gylfason A, Kristinsson KT, et al. 2010. Fine-scale recombination rate differences between sexes, populations and individuals. *Nature* **467:** 1099–1103.

Kryazhimskiy S, Plotkin JB. 2008. The population genetics of dN/dS. *PLoS Genet* **4:** e1000304.

Laland KN, Odling-Smee J, Myles S. 2010. How culture shaped the human genome: bringing genetics and the human sciences together. *Nat Rev Genet* **11:** 137–148.

Lesecque Y, Glémin S, Lartillot N, Mouchiroud D, Duret L. 2014. The Red Queen model of recombination hotspots evolution in the light of archaic and modern human genomes. *PLoS Genet* **10:** e1004790.

Li H, Durbin R. 2011. Inference of human population history from individual whole-genome sequences. *Nature* **475:** 493–496.

Locke DP, Hillier LW, Warren WC, Worley KC, Nazareth LV, Muzny DM, Yang SP, Wang ZY, Chinwalla AT, Minx P, et al. 2011. Comparative and demographic analysis of orang-utan genomes. *Nature* **469:** 529–533.

Malécot G. 1948. *Les mathématiques de l'hérédité.* Masson & Cie, Paris, France.

McVicker G, Gordon D, Davis C, Green P. 2009. Widespread genomic signatures of natural selection in hominid evolution. *PLoS Genet* **5:** e1000471.

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

Meyer M, Kircher M, Gansauge MT, Li H, Racimo F, Mallick S, Schraiber JG, Jay F, Prüfer K, de Filippo C, et al. 2012. A high-coverage genome sequence from an archaic Denisovan individual. *Science* **338:** 222–226.

Mikkelsen TS, Hillier LW, Eichler EE, Zody MC, Jaffe DB, Yang SP, Enard W, Hellmann I, Lindblad-Toh K, Altheide TK, et al. 2005. Initial sequence of the chimpanzee genome and comparison with the human genome. *Nature* **437:** 69–87.

Munch K, Mailund T, Dutheil JY, Schierup MH. 2014. A fine-scale recombination map of the human-chimpanzee ancestor reveals faster change in humans than in chimpanzees and a strong impact of GC-biased gene conversion. *Genome Res* **24:** 467–474.

Nam K, Munch K, Hobolth A, Dutheil JY, Veeramah KR, Woerner AE, Hammer MF, Great Ape Genome Diversity Project, Mailund T, Schierup MH. 2015. Extreme selective sweeps independently targeted the X chromosomes of the great apes. *Proc Natl Acad Sci* **112:** 6413–6418.

Nielsen R, Hellmann I, Hubisz M, Bustamante C, Clark AG. 2007. Recent and ongoing selection in the human genome. *Nat Rev Genet* **8:** 857–868.

Oleksyk TK, Smith MW, O'Brien SJ. 2010. Genome-wide scans for footprints of natural selection. *Philos Trans R Soc Lond B Biol Sci* **365:** 185–205.

Pääbo S. 2014. The human condition—a molecular approach. *Cell* **157:** 216–226.

Powell MJD. 1994. A direct search optimization method that models the objective and constraint functions by linear interpolation. *Adv Optim Numer Anal* **275:** 51–67.

Prüfer K, Muetzel B, Do HH, Weiss G, Khaitovich P, Rahm E, Pääbo S, Lachmann M, Enard W. 2007. FUNC: a package for detecting significant associations between gene sets and ontological annotations. *BMC Bioinformatics* **8:** 41.

Prüfer K, Munch K, Hellmann I, Akagi K, Miller JR, Walenz B, Koren S, Sutton G, Kodira C, Winer R, et al. 2012. The bonobo genome compared with the chimpanzee and human genomes. *Nature* **486:** 527–531.

Prüfer K, Racimo F, Patterson N, Jay F, Sankararaman S, Sawyer S, Heinze A, Renaud G, Sudmant PH, de Filippo C, et al. 2014. The complete genome sequence of a Neanderthal from the Altai Mountains. *Nature* **505:** 43–49.

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

Przeworski M. 2003. Estimating the time since the fixation of a beneficial allele. *Genetics* **164:** 1667–1676.

Pybus OG, Shapiro B. 2009. Natural selection and adaptation of molecular sequences. In *The phylogenetic handbook* (ed. Lemey P, et al.), pp. 415–417. Cambridge University Press, New York.

Pybus M, Dall'Olio GM, Luisi P, Uzkudun M, Carreño-Torres A, Pavlidis P, Laayouni H, Bertranpetit J, Engelken J. 2014. 1000 genomes selection browser 1.0: a genome browser dedicated to signatures of natural selection in modern humans. *Nucleic Acids Res* **42:** D903–D909.

Rabiner LR. 1989. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc IEEE* **77:** 257–286.

Racimo F. 2016. Testing for ancient selection using cross-population allele frequency differentiation. *Genetics* **202:** 733–750.

Racimo F, Kuhlwilm M, Slatkin M. 2014. A test for ancient selective sweeps and an application to candidate sites in modern humans. *Mol Biol Evol* **31:** 3344–3358.

Rasmussen MD, Hubisz MJ, Gronau I, Siepel A. 2014. Genome-wide inference of ancestral recombination graphs. *PLoS Genet* **10:** e1004342.

Rosenberg NA. 2002. The probability of topological concordance of gene trees and species trees. *Theor Popul Biol* **61:** 225–247.

Sabeti PC, Schaffner SF, Fry B, Lohmueller J, Varilly P, Shamovsky O, Palma A, Mikkelsen TS, Altshuler D, Lander ES. 2006. Positive natural selection in the human lineage. *Science* **312:** 1614–1620.

Sabeti PC, Varilly P, Fry B, Lohmueller J, Hostetter E, Cotsapas C, Xie X, Byrne EH, McCarroll SA, Gaudet R, et al. 2007. Genome-wide detection and characterization of positive selection in human populations. *Nature* **449:** 913–918.

Sankararaman S, Patterson N, Li H, Pääbo S, Reich D. 2012. The date of interbreeding between Neandertals and modern humans. *PLoS Genet* **8:** e1002947.

Sankararaman S, Mallick S, Dannemann M, Prüfer K, Kelso J, Pääbo S, Patterson N, Reich D. 2014. The genomic landscape of Neanderthal ancestry in present-day humans. *Nature* **507:** 354–357.

Sankararaman S, Mallick S, Patterson N, Reich D. 2016. The combined landscape of Denisovan and Neanderthal ancestry in present-day humans. *Curr Biol* **26:** 1241–1247.

Scally A, Dutheil JY, Hillier LW, Jordan GE, Goodhead I, Herrero J, Hobolth A, Lappalainen T, Mailund T, Marques-Bonet T, et al. 2012. Insights into hominid evolution from the gorilla genome sequence. *Nature* **483:** 169–175.

St Pourcain B, Cents RA, Whitehouse AJ, Haworth CM, Davis OS, O'Reilly PF, Roulstone S, Wren Y, Ang QW, Velders FP, et al. 2014. Common variation near ROBO2 is associated with expressive vocabulary in infancy. *Nat Commun* **5:** 4831.

Staab PR, Zhu S, Metzler D, Lunter G. 2014. Scrm: efficiently simulating long sequences using the approximated coalescent with recombination. *Bioinformatics* **31:** 1680–1682.

Suda S, Iwata K, Shimmura C, Kameno Y, Anitha A, Thanseem I, Nakamura K, Matsuzaki H, Tsuchiya KJ, Sugihara G, et al. 2011. Decreased expression of axon-guidance receptors in the anterior cingulate cortex in autism. *Mol Autism* **2:** 14.

Tajima F. 1989. Statistical method for testing the neutral mutation hypothesis by DNA polymorphism. *Genetics* **123:** 585–595.

Varki A, Geschwind DH, Eichler EE. 2008. Explaining human uniqueness: genome interactions with environment, behaviour and culture. *Nat Rev Genet* **9:** 749–763.

Vernot B, Akey JM. 2014. Resurrecting surviving Neandertal lineages from modern human genomes. *Science* **343:** 1017–1021.

Vernot B, Tucci S, Kelso J, Schraiber JG, Wolf AB, Gittelman RM, Dannemann M, Grote S, McCoy RC, Norton H, et al. 2016. Excavating Neandertal and Denisovan DNA from the genomes of Melanesian individuals. *Science* **352:** 235–239.

Voight BF, Kudaravalli S, Wen X, Pritchard JK. 2006. A map of recent positive selection in the human genome. *PLoS Biol* **4:** 446–458.

Wang R, Chen CC, Hara E, Rivas MV, Roulhac PL, Howard JT, Chakraborty M, Audet JN, Jarvis ED. 2015. Convergent differential regulation of SLIT-ROBO axon guidance genes in the brains of vocal learners. *J Comp Neurol* **523:** 892–906.

Weaver TD. 2009. The meaning of Neandertal skeletal morphology. *Proc Natl Acad Sci* **106:** 16028–16033.

Wright S. 1951. The genetical structure of populations. *Ann Eugen* **15:** 322–354.

Yang MA, Harris K, Slatkin M. 2014. The projection of a test genome onto a reference population and applications to humans and archaic hominins. *Genetics* **198:** 1655–1670.

Zacher B, Michel M, Schwalb B, Cramer P, Tresch A, Gagneur J. 2016. Accurate promoter and enhancer identification in 127 ENCODE and roadmap epigenomics cell types and tissues by GenoSTAN. *PLoS One* **12:** e0169249.