

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

An ancestral recombination graph of human, Neanderthal, and Denisovan genomes

Permalink

<https://escholarship.org/uc/item/57r8j8rz>

Journal

Science Advances, 7(29)

ISSN

2375-2548

Authors

Schaefer, Nathan K

Shapiro, Beth

Green, Richard E

Publication Date

2021-07-16

DOI

10.1126/sciadv.abc0776

Peer reviewed

GENETICS

An ancestral recombination graph of human, Neanderthal, and Denisovan genomes

Nathan K. Schaefer^{1,2,3†}, Beth Shapiro^{1,2,3}, Richard E. Green^{3,4*}

Many humans carry genes from Neanderthals, a legacy of past admixture. Existing methods detect this archaic hominin ancestry within human genomes using patterns of linkage disequilibrium or direct comparison to Neanderthal genomes. Each of these methods is limited in sensitivity and scalability. We describe a new ancestral recombination graph inference algorithm that scales to large genome-wide datasets and demonstrate its accuracy on real and simulated data. We then generate a genome-wide ancestral recombination graph including human and archaic hominin genomes. From this, we generate a map within human genomes of archaic ancestry and of genomic regions not shared with archaic hominins either by admixture or incomplete lineage sorting. We find that only 1.5 to 7% of the modern human genome is uniquely human. We also find evidence of multiple bursts of adaptive changes specific to modern humans within the past 600,000 years involving genes related to brain development and function.

INTRODUCTION

Much of the current genetic variation within humans predates the split, estimated at 520 to 630 thousand years (ka) ago (1), between the populations that would become modern humans and Neanderthals. The shared genetic variation present in our common ancestral population is still largely present among humans today and was present in Neanderthals up until the time of their extinction. This phenomenon, which is known as incomplete lineage sorting (ILS), means that any particular human will share many alleles with a Neanderthal that are not shared with some other humans. Therefore, humans often share genetic variation with Neanderthals not by admixture but rather by shared inheritance from a population ancestral to us both. Because of this, any effort to map ancestry from archaic hominins in human genomes must disentangle admixture from ILS. Furthermore, a technique able to identify both admixture and ILS could produce a catalog of uniquely human genomic regions that is free of both and thereby shed light on the evolutionary processes that have been important in our origin as a unique species.

Ancestral recombination graph (ARG) inference (2) is a powerful starting point for such an analysis. An ARG can be conceptualized as a series of trees, mapped to individual sites, over phased haplotypes (chromosomes) in a panel of genomes. Ancestral recombination events, or sites at which chromosome segments with different histories were joined together by historical recombination, form boundaries between trees. Each ancestral recombination event manifests as a clade of haplotypes, all of which descend from the first ancestral haplotype to have it, moving from one position in the tree upstream of the event to a new position in the downstream tree (3). ARGs are complete descriptions of phylogenomic datasets and present for recombining genomes what single trees present for nonrecombining genomes, i.e., a complete description of their genetic relationships.

¹Howard Hughes Medical Institute, University of California, Santa Cruz, Santa Cruz, CA 95064, USA. ²Department of Ecology and Evolutionary Biology, University of California, Santa Cruz, Santa Cruz, CA 95064, USA. ³Genomics Institute, University of California, Santa Cruz, Santa Cruz, CA 95064, USA. ⁴Department of Biomolecular Engineering, University of California, Santa Cruz, Santa Cruz, CA 95064, USA.

*Corresponding author. Email: ed@soe.ucsc.edu

†Present address: The Eli and Edythe Broad Center of Regeneration Medicine and Stem Cell Research, University of California, San Francisco, San Francisco, CA 94143, USA.

Copyright © 2021
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
NonCommercial
License 4.0 (CC BY-NC).

As prior techniques for ancestry mapping can be thought of as summaries of the ARG, higher resolution ancestry maps could be produced if the ARG were known. In addition, the ARG can be used to estimate the time to most recent common ancestor (TMRCA) between admixed and admixer haplotypes, providing additional information about historical admixture between humans and archaic hominins.

Given the utility of an ARG, it is expected that several methods have been devised for estimating ARGs from genetic data. These published approaches all have different strengths and limitations. BEAGLE (3), ArgWeaver (4), and Rent+ (5) were designed for small datasets and require substantial time and/or memory to be used with large sequencing panels. Margarita (6) randomly samples histories at ancestral recombination event boundaries and does not seek to produce parsimonious recombination histories (6). ArgWeaver (4), which is widely considered the gold standard in ARG inference, requires prior knowledge of demographic model parameters. Relate (7) is a relatively new method that scales well to large datasets and produces trees without polytomies and with branch lengths but, in doing so, necessarily samples some relationships that are not directly inferred from the data, as do several other methods (4, 5). The most computationally efficient approach, tsinfer (8), also scales to large datasets but assumes that frequency of an allele is correlated with its age. Since this assumption is violated at loci undergoing either admixture or selection, tsinfer is not well suited for ARG inference using genetic data from Neanderthals, Denisovans, and modern humans.

Here, we present a heuristic, parsimony-guided ARG inference algorithm called SARGE (Speedy Ancestral Recombination Graph Estimator) and use it to build a genome-wide ARG of both modern human and archaic hominin genomes. SARGE can run on thousands of phased genomes, makes no prior assumptions other than parsimony, heuristically estimates branch lengths, and avoids making inferences about unobserved relationships by leaving polytomies in output trees. We validate SARGE using simulated data and demonstrate that it has high specificity compared to existing methods in reconstructing the topology of trees, making it suitable for identifying archaic admixture segments. To achieve this high specificity, SARGE avoids describing some relationships in output trees, resulting in lower sensitivity than existing methods.

We run SARGE on a panel of 279 modern human genomes, two high-coverage Neanderthal genomes, and one high-coverage

Denisovan genome. Using the resulting ARG, we map Neanderthal and Denisovan ancestry, ILS, and the absence of both across modern human genomes. We find evidence of at least one wave of Neanderthal admixture into the ancestors of all non-Africans. We also identify several long and deeply divergent Neanderthal haplotype blocks that are specific to some human populations. We find support for the hypothesis that Denisovan-like ancestry is the result of multiple introgression events from different source populations (9, 10). We also detect an excess of Neanderthal and Denisovan haplotype blocks unique to South Asian genomes. Last, we pinpoint human-specific changes likely to have been affected by selection since the split with archaic hominins, many of which are involved in brain development.

RESULTS

ARG algorithm

To build an ARG containing both modern human and archaic hominin genomes without the use of a demographic model or the

need to infer ancestral haplotypes, we developed a parsimony-based ARG inference technique, SARGE. SARGE uses both shared derived alleles and inferred, shared ancestral recombination events to articulate trees (Fig. 1A and Supplementary Methods). SARGE uses the four-gamete test (11) to determine regions of recombination and the affected haplotypes. The crux of SARGE is a fast algorithm for choosing the branch movement(s) capable of explaining the highest number of discordant clades across a genomic segment that fails the four-gamete test. Once the branch movements, i.e., inferred ancestral recombinations, are determined, further definition of clades is possible. Thus, the trees are articulated by both shared alleles and shared ancestral recombination events (figs. S1 and S2 and Supplementary Methods). SARGE infers branch lengths via a heuristic method, compensating for mutation rate variation across the genome by comparing the number of mutations on each branch to the divergence to an outgroup genome in a fixed width region around each site (fig. S7 and Supplementary Methods).

In the interest of parsimony, our method attempts to infer a set of ancestral recombination events that each explains as many four-gamete

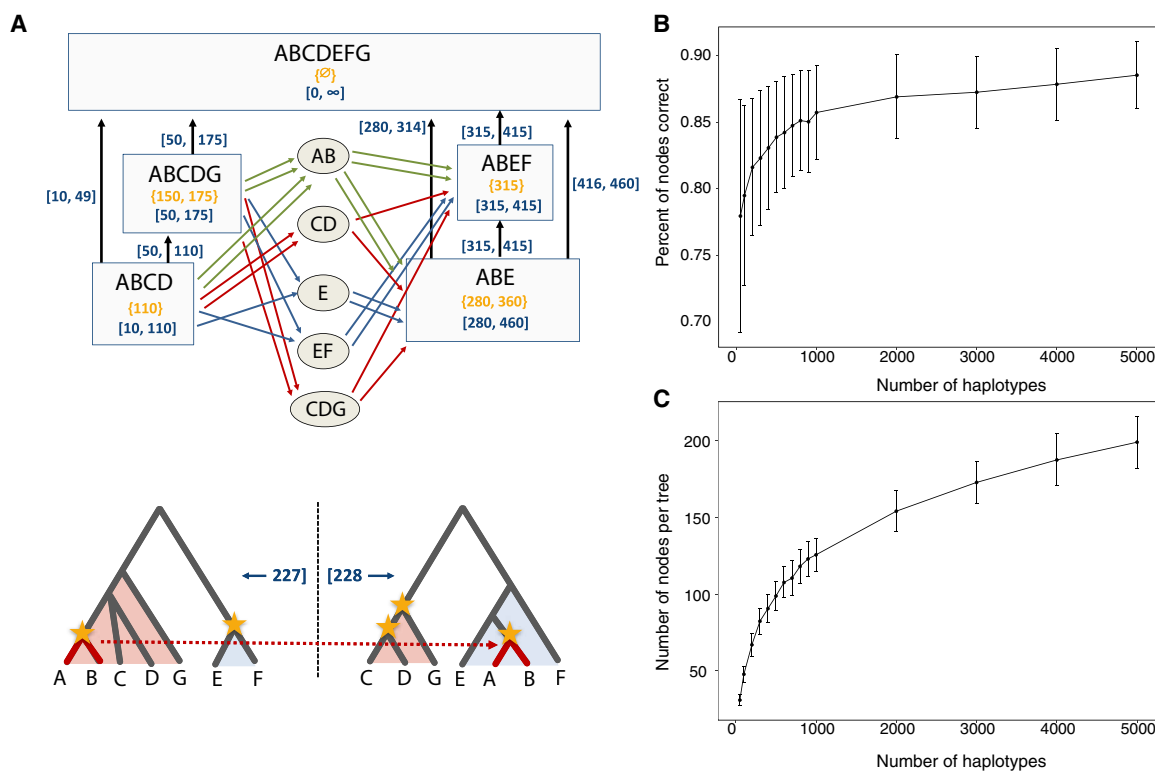


Fig. 1. Data structure and performance of SARGE on simulated data. (A) Schematic of data structure. Top: Rectangles are “tree nodes” representing clades in trees. Each clade has member haplotypes (shown with letters A to G) and a start and end coordinate (blue numbers in brackets) determined by coordinates of single-nucleotide polymorphism (SNP) sites tagging the clade (yellow numbers in braces), along with a propagation distance parameter (100 in this example). Parent/child edges (vertical arrows) also have start and end coordinates determined by the nodes. Ovals are candidates for clades sharing an ancestral recombination event that can explain four-gamete test failures; colored edges indicate potential paths between tree nodes through candidate nodes that could explain four-gamete test failures (colors indicate types of paths). The candidate node with the most edges (here, AB) is eventually chosen as the most parsimonious branch movement, allowing for the inference of new nodes. The two trees at the bottom show the “solved” ancestral recombination event with the branch movement marked in red and all clades inferred without SNP data marked with yellow stars (haplotypes A and B share an ancestral recombination event; their ancestry is shared with haplotypes C, D, and G upstream of the recombination event and haplotype E downstream of it). The coordinates of the recombination event (blue numbers in brackets) are taken to be midway between the highest-coordinate upstream site (left side) and the lowest-coordinate downstream site (right side) involved in recombination. For a more detailed overview of the data structure, see figs. S3 to S5. (B) Accuracy of SARGE on simulated data (defined as percent of all clades correct according to the true ARG in the simulation), with increasing numbers of human-like haplotypes from an unstructured population. Error bars are one SD across five replicates. (C) Number of nodes per tree with increasing number of haplotypes in simulated data.

test failures as possible. Because the four-gamete test is known to underestimate the true number of ancestral recombination events (12, 13), SARGE will systematically underestimate the true number of ancestral recombination events in a dataset by design. Because of this, SARGE is not well suited to certain tasks, such as the creation of fine-grained recombination maps. We have attempted to mitigate cases where a clade in the ARG should be broken by an unobserved ancestral recombination event, however, by introducing a propagation distance parameter that limits the genomic distance over which each observed clade is allowed to persist (Fig. 1A and Supplementary Methods).

SARGE is scalable to large datasets and achieves higher specificity than many other methods at the cost of lower sensitivity, by leaving uncertainty (polytomies) in the output data. Using simulated data, we find that SARGE runs quickly (figs. S8D and S10), requires little memory, and has 78.93% specificity [95% confidence interval (CI), 78.09 to 78.95%] on average across a range of simulated datasets that include between 50 and 500 haplotypes (see Supplementary Methods). SARGE is at least as specific as alternative techniques (fig. S9, A and C). Conversely, SARGE's sensitivity (25.36%; 95% CI, 25.32 to 25.40%) is lower than that of other methods (fig. S9, B and D), as SARGE leaves an increasingly large number of polytomies in output trees as the number of input haplotypes increases (Fig. 1, B and C). As expected, SARGE (as well as similar techniques) performs best when the mutation to recombination rate ratio is high, as this makes clades easier to detect (figs. S11 and S12) and suffers slightly in accuracy with increasing amounts of population structure (fig. S13).

We also find that the sensitivity of SARGE can be increased by increasing the propagation distance parameter (see Supplementary Text and fig. S15), that missing clades are likely to be small clades that are likely to be close to the leaves of trees (see Supplementary Text and fig. S16), and that incorrectly inferred clades tend to be within a few kilobases of sites at which those clades exist in truth (see Supplementary Text and fig. S17). We also find, using simulated data, that the SARGE's branch lengths do not appear to be systematically biased upward or downward (see Supplementary Text and fig. S18).

We ran SARGE on 279 phased human genomes from the Simons Genome Diversity Project (SGDP) (14), together with two high-coverage Neanderthal genomes (1, 15) and one high-coverage Denisovan genome (16). In our analyses, we relied on modern human population labels defined by the SGDP for many analyses, but we split sub-Saharan Africans into one population containing only the most deeply diverged lineages (Biaka, Mbuti, and Khomani-San), which we call "Africa-MBK," and the remaining genomes ("Africa"). Using these data, we find that the completeness of trees in the ARG (the extent to which all possible clades are present rather than in polytomies) is positively correlated with the local mutation rate to recombination rate ratio (fig. S20A; Spearman's $\rho = 0.40$; $P < 2.2 \times 10^{-16}$) and that the number of inferred ancestral recombination events per genomic window agrees with a previously published population recombination map (17) (fig. S20B; Spearman's $\rho = 0.46$; $P < 2.2 \times 10^{-16}$), as expected. Estimates of the mean TMRCA of pairs of haplotypes, taken across all trees, were also concordant with prior knowledge (Fig. 2A).

Using these data, we found SARGE's inferences of ancestral recombination events to be accurate. Because SARGE articulates tree clades using either shared allelic variation or shared inferred ancestral

recombination, it is possible to test the concordance of trees made from each source. On average, 13.2% of clades in the ARG are known only from inference of shared ancestral recombination events and not by the presence of a shared, derived allele. We created a similarity score between every pair of phased human genome haplotypes in our dataset based on how often the haplotypes share ancestral recombination events. This score recapitulates relationships among humans known from single-nucleotide polymorphism (SNP) data alone (Fig. 2, B and C; Pearson's r^2 with scores from SNP data = 0.989; $P < 2.2 \times 10^{-16}$). We note that genomes with the poorest correlation between SNP-based and recombination-based similarity scores to other genomes are those most likely to contain phasing errors (table S1).

Archaic hominin admixture

We used our ARG to find regions of each phased human genome that derived from admixture with archaic hominins (see Supplementary Methods and fig. S24). If humans and the archaic hominins in our panel were in populations that had sorted their lineages, then this exercise would be simple with a complete and correct ARG. However, since human genome regions are often within a clade that includes hominin haplotypes due to ILS, finding admixed segments requires analysis beyond simply finding clades that unite some human and archaic hominin haplotypes.

We started by selecting clades from ARG trees that united some modern humans with archaic hominins to the exclusion of some other modern humans. We then assigned each human genome haplotype in each such clade as putative Neanderthal, Denisovan, or ambiguous ancestry, depending on whether the clade contained Neanderthal, Denisovan, or both types of haplotypes. We then performed several filtering steps to remove these clades likely to result from ILS. First, we removed any clades that included more than 10% of the Africa-MBK haplotypes from the most basal human lineages, which are unlikely to be admixed. We then discarded clades that persisted for a short distance along the chromosome (which likely represent older haplotypes broken down over time by recombination) or in which the TMRCA between modern humans and archaic hominins was high (see Supplementary Methods and fig. S24). This ascertainment strategy was designed to identify haplotype blocks that we could confidently identify as archaic-introgressed and therefore likely underestimated the true extent of admixture across the genome. Because our method relies on both the haplotype block length and the TMRCA between admixed and introgressor haplotypes to identify admixed segments, we were able to identify some haplotypes that resemble archaic admixture in modern humans but that have relatively high sequence divergence to published archaic genomes (manifesting as high TMRCA between archaic and modern genomes within these segments).

Using the resulting maps, we calculated genome-wide percent admixture estimates across populations and compared them to \hat{f} estimates based on the population-wide D statistic (18, 19) using basal Africa-MBK lineages as an outgroup. Since our alleles were polarized relative to the chimpanzee genome and only sites with derived alleles present in hominins were considered, our calculations were of the form $D(\text{Africa-MBK, Test, Introgressor, Chimpanzee}) / D(\text{Africa-MBK, Introgressor}_1, \text{Introgressor}_2, \text{Chimpanzee})$, where Introgressor_1 and Introgressor_2 were randomly chosen subsets of half of the introgressor (Neanderthal or Denisovan) haplotypes and the derived allele frequency in chimpanzee was set to 0 at all sites in

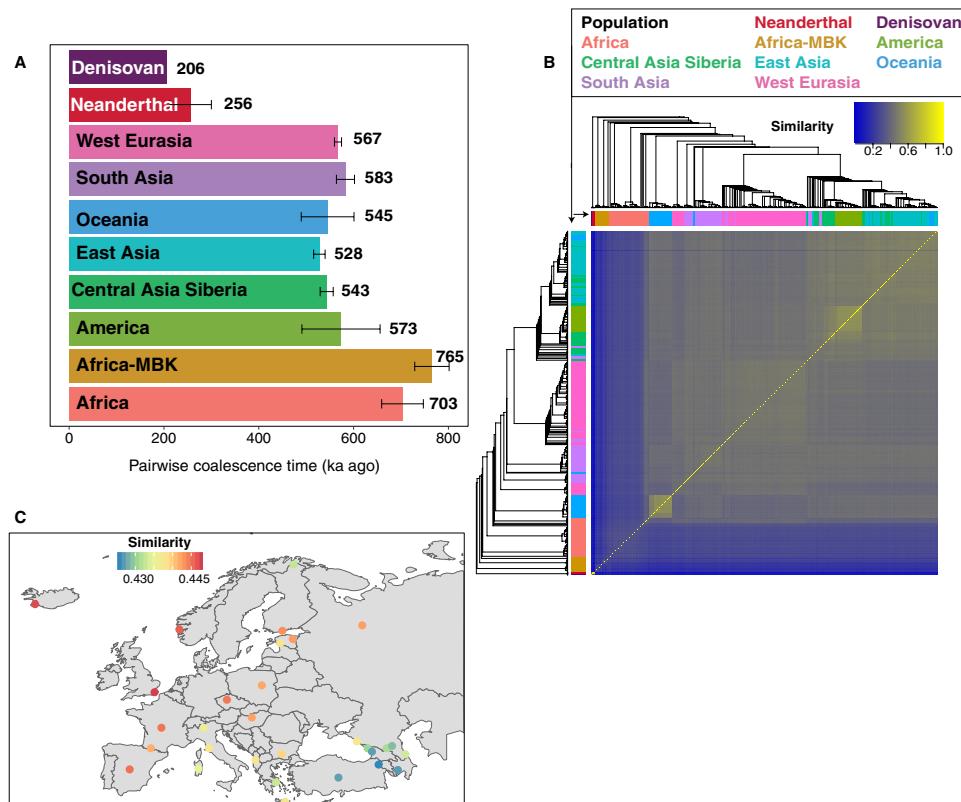


Fig. 2. Performance of SARGE on SGDP and archaic hominin dataset. (A) Pairwise coalescence times for randomly sampled sets of up to 10 pairs of phased genome haplotypes per population (every possible pair was considered for archaic hominins, since fewer genomes were available). Values are calibrated using a 13-Ma human-chimp divergence time (see Supplementary Methods) and averaged across every variable site in the dataset, error bars show one SD, and branch shortening values for archaic samples were incorporated into calculations using mean values reported in (7). The lower value for humans comes from removing archaic-admixed clades from trees. (B) UPGMA (unweighted pair group method with arithmetic mean) trees computed using nucleotide diversity from SNP data (top and left) against similarity matrix from shared recombination events inferred by SARGE. Light yellow boxes (similar groups) are Native Americans and Papuans. (C) Average similarity between Orcadian haplotypes in the SGDP panel and other European haplotypes calculated on the basis of the number of shared ancestral recombination events. The best matches are in England, Iceland, and Norway, as expected.

our dataset. ARG-based estimates are similar to, but lower than, D statistic-based estimates in all non-African genomes, which we expected because of our aggressive filtering strategy for eliminating ILS (see Supplementary Methods). We detected slightly more admixture in sub-Saharan Africans (excluding Africa-MBK) than using the D statistic (Fig. 3A), even when considering the lower end of 95% confidence weighted block jackknife intervals (table S2). We note that a recent study that used an outgroup-free method to detect Neanderthal ancestry blocks in human genomes also found a higher average amount of Neanderthal ancestry in African genomes than has been previously reported (20). As another quality check, we compared our maps of Neanderthal ancestry to those published in prior studies (20–23). We found that maps produced by SARGE are about as concordant with the published maps as the published maps are with each other (fig. S25).

We next looked for population-specific differences in archaic hominin ancestry in modern humans. Lengths of archaic haplotype segments and the TMRCA to admix across those segments are both affected by the time of admixture and the divergence between the true admixers and available archaic hominin genomes. We therefore computed the mean of these two values for each ancestry type per phased genome haplotype and compared them across individuals from different populations to look for evidence of distinctive,

population-specific admixture events. This analysis revealed distinctive population-specific patterns for Neanderthal and Denisovan ancestry, and many pairwise comparisons of these values between populations are significant (table S3). Segments of ambiguous ancestry produce a pattern resembling a mixture of Neanderthal and Denisovan ancestry, as expected (Fig. 3, B to D, and figs. S26, S28, S37, and S38). We caution, however, that our approach can artificially shorten haplotype block lengths (see Supplementary Methods and figs. S21 and S22), especially for populations such as Papuans and Australians that were absent from the 1000 Genomes Project panel (24) that was used for phasing (14). Nonetheless, Neanderthal haplotype block lengths in Oceania are not significantly shorter than in other populations (Fig. 3B), and incorrect phasing in archaic genomes does not appear to negatively affect results of admixture scans using simulated data (see Supplementary Text).

As expected, the ARG classifies a smaller fraction (0.096 to 0.46%) of sub-Saharan African genomes (excluding Mozabite and Saharawi individuals) as resulting from Neanderthal admixture compared to non-African genomes (0.73 to 1.3%). The haplotype segments of African genomes that are grouped together in clades with Neanderthal haplotypes are distinctive from the haplotype segments found in the genomes of people with non-African ancestry (Fig. 3B and fig. S26A). Namely, the African haplotypes are more dissimilar to the Neanderthal

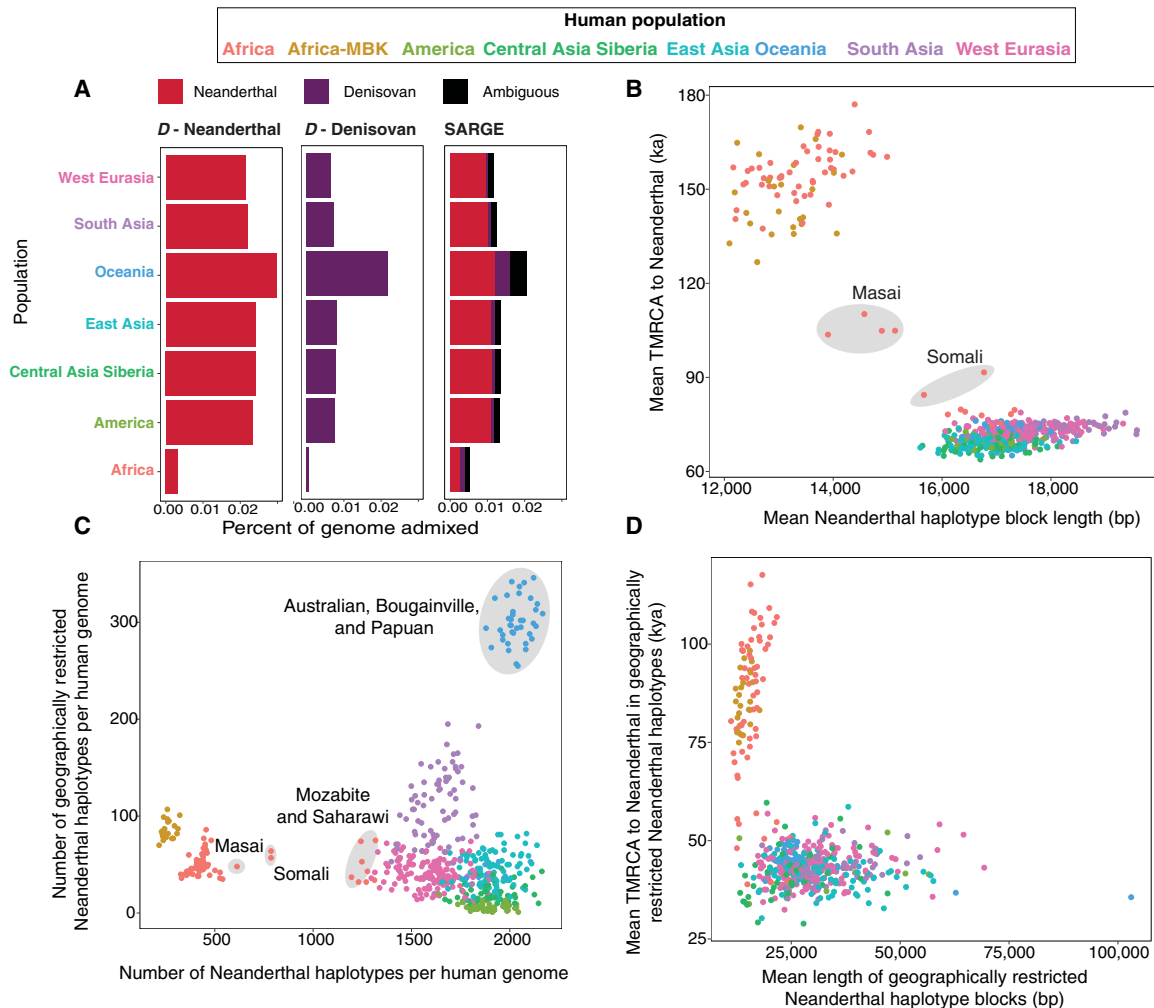


Fig. 3. Neanderthal ancestry inferred in modern humans. (A) Genome-wide percent Neanderthal, Denisovan, and ambiguous (either Neanderthal or Denisovan) across SGDP populations, using the ARG and an estimator based on the *D* statistic. *D* statistic calculations considered only one archaic population at a time as introgressor and thus do not detect ambiguous ancestry and also might count some Denisovan ancestry as Neanderthal and vice versa. (B) For individual phased human genome haplotypes (points), mean TMRCAs with Neanderthal in Neanderthal haplotype blocks (y axis) and mean Neanderthal haplotype block length (x axis) for all Neanderthal-introgressed haplotypes. TMRCAs calculations assume a total of 6.5-Ma human-chimpanzee divergence and branch shortening values from (1), with a mutation rate of 1×10^{-9} per site per year. bp, base pair. (C) Overall number Neanderthal haplotype blocks versus geographically restricted (unique to a 3000-km radius) Neanderthal haplotype blocks. (D) Same as (B), but limited to geographically restricted (unique to a 3000-km radius) Neanderthal haplotype blocks. Only haplotypes with more than 10 geographically restricted segments are shown.

haplotypes with which they are grouped and tend to be shorter. These observations are qualitatively consistent with the model where in genetic drift may group Neanderthal and African haplotypes, independent of a specific admixture event. It is also possible that these haplotypes are the result of true introgression events from unknown archaic hominins distantly related to the Neanderthal/Denisovan lineage (25). Another recent study using an inferred ARG also found mysterious, divergent haplotypes within sub-Saharan Africans that resembled unknown archaic introgression (7).

Unexpectedly, however, two of the SGDP African populations, Masai and Somali, are intermediate between non-African and African genomes when measuring lengths of archaic haplotype segments and TMRCAs to admixers within them (Fig. 3B). These Neanderthal haplotype blocks may have originated in ancient European migrants to eastern Africa (26) and spread beyond eastern Africa through

gene flow, which is known to have affected even the basal Africa-MBK lineages (27).

To test this hypothesis, we recomputed the mean length and TMRCAs of admixer genomes within archaic-introgressed haplotype segments across all individuals, using only geographically restricted segments. We defined these as any archaic haplotype segments found only in genomes that were sampled within a 3000-km radius of each other (using geodesic distance between sampling coordinates). This analysis showed Masai and Somali genomes to have a comparable number of geographically restricted Neanderthal haplotypes to most other African genomes (Fig. 3C), with similar haplotype block lengths and TMRCAs to admixers within these geographically restricted haplotypes (Fig. 3D). This observation is concordant with the idea that the unusual Neanderthal haplotypes in these populations originated in Eurasian migrants.

Outside of Africa, our Neanderthal introgression maps largely agree with prior studies. We detect a mean TMRCA to Neanderthal of about 74 ka ago across all Neanderthal haplotype blocks in non-African populations, using published corrections for branch shortening in the archaic genomes (1). This value is consistent with published estimates of Neanderthal admixture times and the phylogenetic distance between the Vindija33.19 Neanderthal and the introgressing Neanderthal calculated by a recent study (28). The mean TMRCA between genomic segments detected as Neanderthal admixture segments and the Neanderthal itself is consistent within about 5000 years for all populations outside of Africa (Fig. 3B). We see slightly more Neanderthal ancestry in Central Asia, East Asia, and the Americas than in Europe, South Asia, and Southwest Asia (Fig. 3A). We also find more geographically restricted Neanderthal haplotype blocks in South Asia than elsewhere in mainland Eurasia, and the fewest geographically restricted Neanderthal haplotype blocks in the Americas (Fig. 3C and fig. S35).

Humans in Central and East Asia are known to have elevated Neanderthal ancestry compared to other populations (22). However, there is debate over whether this elevated Neanderthal ancestry is due to smaller past population size relative to other groups and the resulting stronger effect of genetic drift (22) or to additional pulses of Neanderthal admixture specific to these populations (9, 29). Although we detect more Neanderthal ancestry in Central and East Asians than in West Eurasians, we detect a similar number of geographically restricted haplotype blocks (unique to a 3000-km radius) in both groups (Fig. 3C). Further, Neanderthal haplotype blocks are shorter on average and therefore older in Central and East Asians than in West Eurasians (Fig. 3B). This implies that the excess Neanderthal ancestry in Central and East Asians mostly comprises broadly shared haplotype blocks from introgression common to all non-Africans, consistent with the drift scenario. Another recent study (20) suggested that excess Neanderthal ancestry in Central and East Asians could be mostly due to underestimating Neanderthal ancestry in West Eurasians as a result of using sub-Saharan Africans, who share some Neanderthal ancestry with West Eurasians, as a model unadmixed outgroup. We reject this explanation because removing Neanderthal haplotypes from West Eurasians would likely increase the number of geographically restricted Neanderthal haplotypes in Asians, contrary to our observation. Furthermore, South Asians appear to have a comparable amount of Neanderthal ancestry to West Eurasians (Fig. 3, B and C), despite sharing few Neanderthal haplotype blocks with Africans relative to West Eurasians (fig. S37A). Last, because our strategy for ascertaining Neanderthal haplotype used only Mbuti, Biaka, and Khomani-San genomes as an outgroup and allowed up to 10% of these genomes to have Neanderthal ancestry, we do not believe our results were significantly biased by our choice of outgroup. This is further demonstrated by the fact that we detect a non-negligible amount of Neanderthal ancestry in all African groups (table S2).

Aside from these broadly shared haplotype blocks, we also observe geographically restricted Neanderthal haplotype blocks in each non-African population in our panel. These population-specific haplotype blocks tend to be longer than the shared haplotype blocks and to have an older TMRCA to the Neanderthal genome than the broadly shared haplotype blocks (Fig. 3D, compared to all blocks shown in Fig. 3B). These observations suggest that the population-specific haplotype blocks may be the result of more recent population-specific Neanderthal admixture, as has recently been suggested (26–28, 30).

We next investigated population-specific patterns within Denisovan ancestry segments and found that these segments probably originate from admixture with multiple, divergent individuals that were distantly related to the Denisovan genome. This implies that the Denisovan genome is not a good model for the actual population that admixed with humans with “Denisovan” ancestry. Prior studies have suggested that Denisovan-like haplotype blocks in humans have two or three distinct sources with different levels of divergence to the Denisovan genome, with the best-matching haplotype blocks in East Asia (9, 10). We uncover the same signal: Geographically restricted Denisovan haplotype blocks have the lowest TMRCA to the Denisovan genome in East Asian genomes (mean TMRCA to Denisovan of 125 ka ago) (figs. S26 and S27).

Unexpectedly, we detected many Neanderthal and Denisovan-like haplotype blocks that are unique to South Asia (Fig. 3C and figs. S26C, S45, S35, and S36) and many Neanderthal haplotype blocks that are unique to Oceania (Fig. 3C and figs. S44 and S35). These geographically restricted Neanderthal haplotype blocks are no more divergent to the Neanderthal genome than those specific to other populations (Fig. 3D), complicating any interpretation of these regions.

Genomic regions free of admixture and ILS

Our ARG strategy allows us to bin the human genome into regions containing archaic admixture in at least some humans, regions of ILS, and regions free of both archaic admixture and ILS in all humans (hereafter archaic “deserts”). We find that approximately 7% of the human autosomal genome is human-unique and free of both admixture and ILS. Roughly 50% of the human genome contains regions where one or more humans has archaic ancestry obtained through admixture. If deserts are further restricted to regions that contain a high-frequency, human-specific derived allele, i.e., a substitution that can be assigned to the human lineage (hereafter “human-specific regions”), then these comprise only 1.5% of the assayed genome (Fig. 4A). Despite comprising very little of the genome, however, human-specific regions are significantly enriched for genes, exons, and regulatory element binding sites, while deserts are enriched for both genes and regulatory element binding sites (table S4). In line with previous studies (21, 31), we find admixed regions to be depleted of genes. Regions of ILS are enriched for overlap with genes but significantly depleted of exons (table S4).

To obtain an expectation of the extent of these different types of genomic regions, we ran a series of coalescent simulations with different amounts of archaic hominin admixture, as well as with no admixture (see Supplementary Methods and Supplementary Text). Our observation in the real data, which only 7% of the autosomal genome is free of both archaic admixture and ILS, is consistent with the results of a simulation that resulted in both higher archaic ancestry estimates per individual and more variability in individual archaic ancestry estimates than we observed in real data (Fig. 4A, Supplementary Text, and fig. S47). Two nonmutually exclusive explanations for this difference are (i) the existence of more, geographically limited, archaic hominin admixture events than the two we modeled (see Supplementary Text) and (ii) purifying selection acting against deleterious alleles found within archaic admixed segments (32). Under the model including selection, epistasis may be important, as the strength of selection against archaic alleles would vary depending on which archaic alleles are present in each human population.

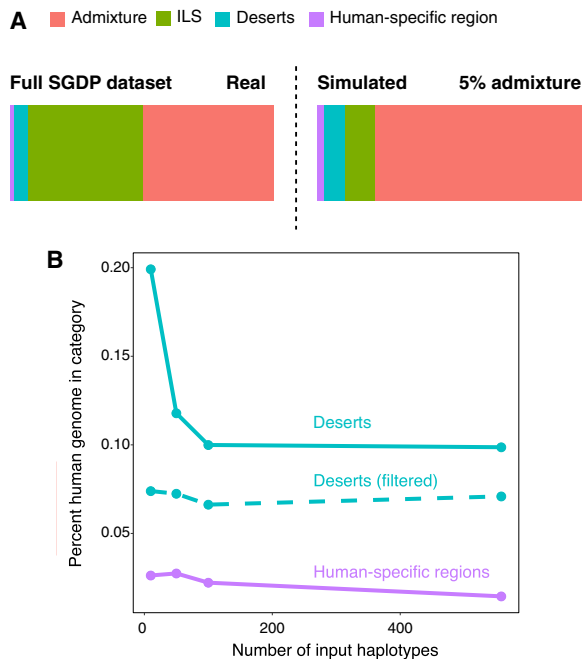


Fig. 4. Extent of deserts of archaic hominin ancestry. (A) Left: In the SGDP dataset, fractions of the genome where any individual has archaic ancestry (admixture), where any individual is incompletely sorted with archaic hominin lineages but where there is no archaic admixture (ILS), where there is no evidence of either admixture or ILS with archaic hominins (deserts), and where there is a fixed derived allele private to and shared by all humans (human-specific regions). Right: The same values from a simulated dataset with single pulses of Neanderthal and Denisovan admixture, both with an admixture proportion (2%) that produced reasonable amounts of archaic ancestry per individual genome (fig. S47C). (B) For random subsamples of the SGDP dataset, along with the full dataset, fractions of the genome comprising deserts, deserts filtered for candidate archaic alleles using another dataset, and human-specific regions are shown. The points on the far right (full dataset) correspond to the desert and ILS bars in the left panel of (A).

The power to detect deserts, i.e., regions in which no human carries a haplotype shared with an archaic hominin by ILS or admixture, can be expected to be affected by the number of human genomes available for analysis. To be certain that we have found the true extent of archaic deserts, we inferred ARGs over random subsamples of the human panel, computing the extent of deserts and human-specific regions for each (Fig. 4B). We were able to recover the full extent of deserts using a subsample of 100 haplotypes, less than half the size of the full panel, suggesting that the panel is sufficiently large.

Given the limited sensitivity of our method compared to alternatives, we also compared the ability of SARGE, Relate, and tsinfer to recover true desert regions in demographic simulations with different numbers of haplotypes (see Supplementary Methods). We found that all methods, including SARGE, tended to overestimate the extent of desert regions, with this tendency increased in the larger dataset (figs. S48 and S49), further highlighting the unexpectedly small fraction of the human genome found to be in these deserts in the actual human genome data.

Timing of human-specific mutations

Given a clade of interest, mutations shared by all members of the clade must have arisen between that clade's TMRCA and its parent

clade's TMRCA. Using this logic, and using the local divergence to the chimpanzee genome in the surrounding region of each tree to calibrate the local mutation rate, assuming human-chimpanzee divergence 13 Ma ago (see Supplementary Text) (18), we estimated ages of all human-specific mutations within deserts. For a schematic of how branch lengths were estimated, see fig. S7. Because the order of mutations along any given branch is unknowable, we took the midpoint of each branch, in years, to be the approximate age of each mutation. Combining these dates with a catalog of high-frequency, human-specific mutations and other annotation data (see Supplementary Methods) allowed us to construct a picture of human-specific evolutionary changes through time. We note that our date estimates rely on the assumption of a constant local mutation rate around each locus across all lineages since human/chimpanzee divergence. Some studies have found evidence for mutation rate heterogeneity in different hominoid lineages, with mutations accumulating more slowly in humans than in other apes, e.g., (33). In this case, some of our TMRCA estimates could be biased downward, as more time would have been required to accumulate the mutations tagging recently formed clades than those that occurred branches closer to the root.

We first examined whether there were one or multiple bursts of human-specific adaptive changes since divergence with Neanderthals and Denisovans. We compiled the ages of all fixed or nearly fixed human-specific derived mutations within archaic hominin deserts that either were annotated as nonsynonymous substitutions (34) or fell within annotated regulatory element binding sites. The age distribution of these mutations is centered around 600 ka ago (Fig. 5A).

We then compared the ages of mutations affecting pairs of genes that interact, according to the STRING database (35), to see whether any clustered around specific time points (see Supplementary Methods). We find two distinct bursts of these mutations, one concentrated around 600 ka ago and another around 200 ka ago (Fig. 5B). We note that, because many of our human-specific genes are likely functionally important and purifying selection can decrease genetic diversity, some of the time estimates for these mutations may be biased downward.

Estimating how and when the modern human lineage arose remains controversial. Dating the oldest population split within modern humans using genetic data has suggested times as recent as 200 to 100 ka ago (36, 37). Archeological evidence paints a more complex and older story; however, a recent study reported human remains with many modern features but archaic cranial morphology dated to about 315 ka ago (38), suggesting that not all human-specific traits arose at the same time. Other studies have found that accumulation of derived morphological features in humans occurred in approximately three periods, whose boundaries correspond roughly to the timing of mutational bursts we found (39).

Functional consequences of human-specific mutations

Comparison of the human and extinct hominin genomes could reveal instances of positive selection that are undetectable via allele frequency or haplotype-based analyses within modern humans or through comparative genomics between humans and other primate genomes (40). The ARG framework is an attractive approach as it pinpoints truly human-specific genomic regions unaffected by either admixture or ILS.

We performed a gene ontology (41) enrichment analysis on the human-specific haplotype regions that accounted for the lengths of

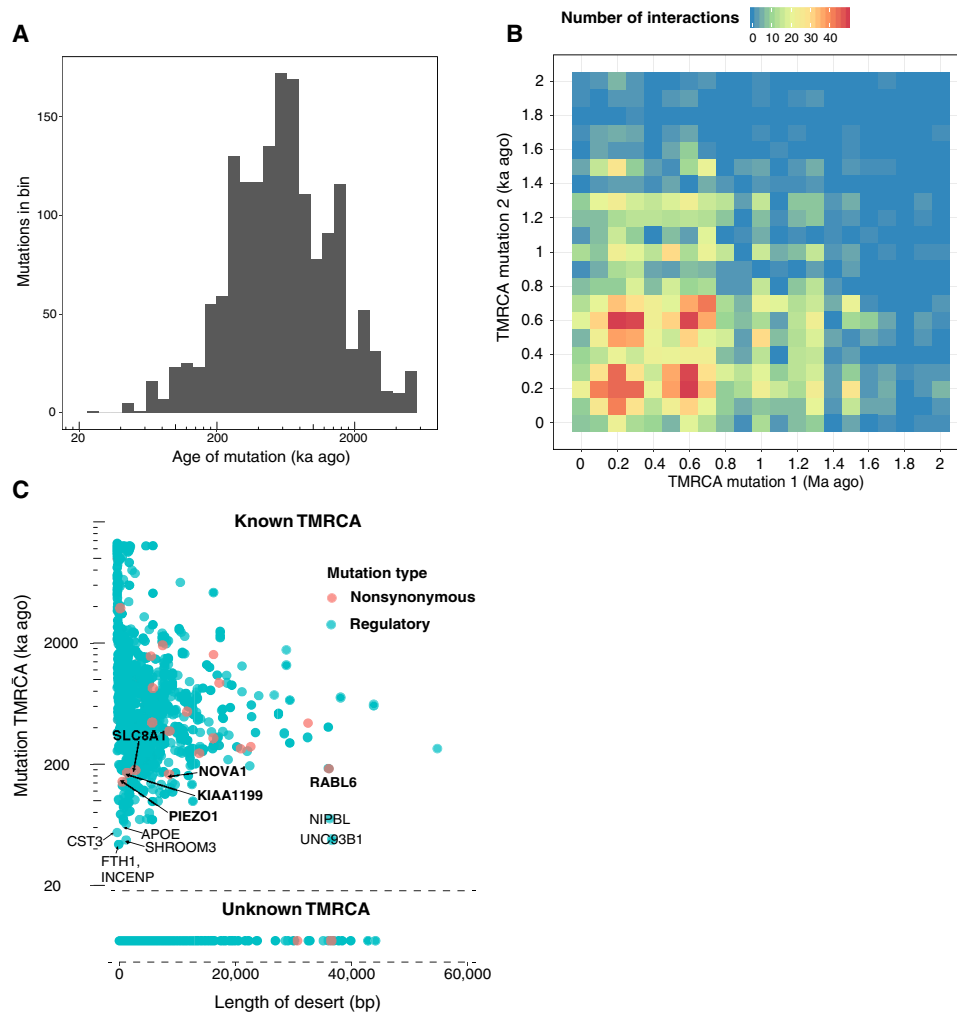


Fig. 5. Approximate ages and consequences of fixed human-specific derived mutations. (A) Ages of candidate human-specific functional mutations (nonsynonymous substitutions and mutations within regulatory element binding sites) inferred by SARGE within desert regions (free of ILS and admixture between archaic hominins and modern humans). (B) For interacting pairs of genes in the STRING database (35) for which candidate human-specific functional mutations affect both genes, the ages of the two mutations are shown. (C) For each candidate human-specific functional mutation, the length of the surrounding desert region (x axis) and inferred mutation age (y axis) are shown. Mutations for which SARGE did not infer a date (mutations within CpG sites or for which the ancestral allele was unknown) are shown in the bottom panel. Mutations were scored on the basis of length of desert and age; genes with regulatory mutations in the top 99.5th percentile of this distribution, or nonsynonymous mutations in the top 95th percentile, are shown, with nonsynonymous changes in bold.

candidate genes (see Supplementary Methods) and found these regions to be heavily enriched for genes related to neuron growth, synapse assembly and function, and cell adhesion (table S5). We note that this may occur because of positive selection in the ancestors of all humans, strong purifying selection that reduces the TMRC within humans so that it excludes extinct hominins, or a combination of both. Further, this test does not suggest specific functional consequences of specific mutations.

We ordered human-specific derived mutations within desert regions found in our ARG by the strength of evidence that they were targeted by selection. To this end, we first sought to limit analysis to potentially functionally relevant mutations, defined as mutations that either caused a nonsynonymous substitution relative to archaic hominin genomes or fell within annotated binding sites for regulatory elements known to affect specific genes (see Supplementary Methods). We developed a simple score for each mutation based on

its inferred age, where available, and the length of the surrounding desert region (Fig. 5C). The rationale for this approach is that older human-specific substitutions should be in shorter haplotypes as they would have undergone more generations of recombination. Mutations in haplotype regions that run counter to this expectation are a priori more likely to have been affected by positive selection. This approach is similar to the recently described extended lineage sorting (ELS) scan (42), which prioritized long genomic intervals where modern human and archaic hominin lineages are completely sorted. Unlike the ELS method, however, our method only considers alleles that are fixed in modern humans. Our model also does not use a hidden Markov model to smooth transitions between sorted and unsorted haplotypes.

Several patterns emerge when considering genes with high-scoring human-specific mutations, and we highlight some of these key findings. Gene ontology terms related to mRNA splicing, processing,

and export are enriched in genes with high-scoring mutations (table S7). Of these, we find a regulatory mutation affecting one—LUC7L3—that is somewhat tissue specific (tau, a measure of tissue specificity scaled from 0 to 1 = 0.713), most highly expressed in cerebellar tissue, and annotated to be involved in splice site selection. Its paralog LUC7L and the gene KHDC4, both of which also have high-scoring (top 50th percentile) regulatory mutations, are involved in the same process. In addition, the gene NOVA1, which harbors a nonsynonymous mutation in the top 95th percentile of our score distribution (Fig. 5C), is a neuronal splicing factor that regulates splicing of genes involved in synapse formation within the brain (43).

Other types of genes, largely related to brain function and development, appear to be affected by high-scoring mutations. Many genes localized to the centrosome and mitotic spindle are involved in maintaining the polarity of dividing neuroblasts, and some mutations affecting these genes are thought to be critical for the development of the human neocortex (44). We find the term “asymmetric neuroblast division” to be enriched in high-scoring genes (table S7). Among individual genes, we find a high-scoring nonsynonymous mutation affecting the centrosomal protein RABL6, which is highly expressed in cerebellar tissue and overexpressed in cancer (45) and a high-scoring regulatory mutation affecting INCENP, a protein crucial for cytokinesis that localizes to the mitotic spindle and centromere (46). Axon pathfinding is another process suggested to have been targeted by human-specific changes; the gene PIEZO1 is involved in this process (47) and harbors a high-scoring nonsynonymous mutation. In addition, the protocadherin PCDHGB7, which contains a nonsynonymous substitution within a long desert region but which we could not date, is a member of a gene family that generates neuronal cell surface identity and is thought to help guide growing neurites (48). In addition to these, we find a number of other mutations potentially affecting genes involved in histone acetylation, neural cell migration, and the clearing of toxic substances from the brain (see Supplementary Text).

DISCUSSION

We implemented a new ARG inference approach, SARGE, and used it to build the first genome-wide ARG of both human and archaic hominin genomes. Analysis of the topology of these ARG trees confirms prior findings about archaic hominin admixture but with important new biological insights. For one, we find that a low fraction, 1.5 to 7%, of the human genome is uniquely human, with the remainder comprising lineages shared with archaic hominins from either ILS or admixture. This small human-specific fraction of the genome is enriched for genes related to neural development and function. We also find evidence for multiple waves of human-specific mutations that occurred through time, suggesting that the modern human phenotype may have developed in stages.

In addition to Neanderthal admixture into the ancestors of all modern non-African populations, we find evidence for other, population-specific episodes of admixture throughout Eurasia. The TMRCA to these population-specific Neanderthal haplotype blocks is deeper than the TMRCA to the Neanderthal haplotype blocks shared by all non-African populations. This deeper TMRCA suggests that Neanderthals contributing population-specific ancestry were less closely related to published (Altai and Vindija) Neanderthal genomes than were the Neanderthals that contributed the broadly shared Neanderthal haplotype blocks. We also find that Neanderthal

ancestry is present to a smaller extent in some African genomes due to back-migration, consistent with other recent reports (20).

We note that our estimated TMRCA to Neanderthal within Neanderthal-introgressed segments in all non-African populations is recent, ~74 ka ago, and implies therefore that little genetic drift separates admixed humans from sequenced Neanderthals in these segments. This recent TMRCA suggests that the majority of Neanderthal ancestry in modern humans originated from Neanderthal gene flow into the ancestors of all non-Africans before populations diversified. It also suggests that at least one of the Neanderthal genomes used here is closely related to the Neanderthal(s) involved in this admixture event. The slightly elevated Neanderthal ancestry that others have described in Central and East Asian populations also appears to have originated in this first pulse, as Central and East Asian Neanderthal haplotypes are mostly shared with other, geographically distant populations. This observation favors the hypothesis that the increased Neanderthal ancestry in these populations relative to others is due to weaker selection against alleles that may be mildly deleterious (32), made possible because of smaller historical population sizes in this part of Eurasia, rather than to additional admixture events (22). Our evidence of many small-scale, population-specific admixture events, however, together with a simulation study that found a single-pulse admixture model followed by drift unable to explain the discrepancies in admixture proportions in European and Asian genomes (49), hints at a complex history of admixture throughout Eurasia not fully captured by either of these two hypotheses.

Last, the genomes of some Oceanian and other populations harbor genes from a population most closely related to the archaic Denisovan genome. The available Denisovan genome is less genetically similar to the admixing genome than the available Neanderthal genomes are to the admixing Neanderthals. While we are hopeful that future work may uncover a DNA-bearing fossil better representing the population involved in the Denisovan admixture, our approach allows identification of admixed regions that can be used to better describe the genome of the archaic hominin group involved in the admixture event. Larger panels of Denisovan admixed genomes may one day provide a nearly complete Denisovan genome scavenged in parts from the genomes of admixed human individuals.

The ARG also allows for prioritizing the selective importance of mutations specific to, and shared by, all modern humans by considering the TMRCA of those mutations together with the lengths of their surrounding human-unique regions. Many of these selected human-specific mutations appear to affect genes involved in neural development and function, as well as RNA splicing. Using new tools for genome editing and brain organoid models for neural function, these mutations are obvious and important targets for experimental studies to determine what was selected in our human ancestors after divergence from our most closely related, extinct relatives.

MATERIALS AND METHODS

Data processing

We downloaded data from the SGDP panel (14), along with two Neanderthal (1, 15) genomes and one Denisovan (16) genome. The Simons data were downloaded in prephased form from https://sharehost.hms.harvard.edu/genetics/reich_lab/sgdp/phased_data/; phasing was performed using SHAPEIT2 (50). We note that the hosts of the data state that the genotypes they provide at sites lacking a

homologous chimpanzee are unreliable; we discarded all these sites from analysis.

Existing variant call sets for the ancient samples were either created using a genotype caller that did not account for ancient DNA damage (15, 16) or were subjected to a mapability filter that discarded many sites in the genome (1). Because our method is sensitive to genotype errors and seeks to make inferences at every possible site in the genome, we chose to recall variants in these three genomes using the ancient DNA-aware genotype caller *snpAD* version 0.3.0 (51). For all *snpAD* runs, we required a minimum base quality of 25 and treated different types of libraries separately, separating uracil-DNA glycosylase (UDG)-treated and non-UDG-treated libraries in the case of the Vindija Neanderthal and separating single-stranded and double-stranded library data for the Altai Neanderthal and Denisovan.

Although the SGDP data were already phased, phasing posed a challenge for the Neanderthal and Denisovan data, for which there is no reference panel and for which DNA is fragmented into short segments. Fortunately, the comparatively low nucleotide diversity in these archaic hominins results in the presence of long runs of homozygosity, which are phased by definition. As an unbiased first step, we performed read-backed phasing using *WhatsHap* version 0.16 (52) (with default parameters, plus `-ignore-read-groups`). Before filtering SNPs for quality and coverage, this phased 722,828 of 11,746,838 heterozygous sites (6.2%) in the Altai Neanderthal, 346,992 of 48,083,469 heterozygous sites (0.7%) in the Vindija Neanderthal, and 514,575 of 33,951,346 heterozygous sites (1.5%) in the Denisovan. Many of the remaining, unconfidently phased heterozygous sites were excluded by other, later filtering steps; however, in our final, filtered dataset, we were left with only 1,677,774 of 49,876,210 total SNPs (3.4%) for which at least one archaic hominin individual was heterozygous and not phased by read-backed phasing.

Following read-backed phasing, we merged archaic hominin VCF files [using *bcftools merge* from *bcftools* version 1.8 (53)] and then phased the merged files using *Eagle2.4* (54), with the 1000 Genomes Project data (24) as a reference panel. We used *Eagle2*'s default parameters but specified that it should not impute missing data (`--noImpMissing`) and that it should output alleles that it could not phase (`--outputUnphased`). After this, we randomly assigned both alleles at every unphased heterozygous site to one or the other haplotype. Although this decision, along with the use of a modern human reference panel, undoubtedly introduced haplotype switch errors, we deemed this preferable to excluding sites that were not confidently phased (which would require us to exclude data from all of the SGDP individuals at the same sites). To mitigate problems arising from this decision, we avoided performing any haplotype-specific analyses on the archaic hominin genomes. When creating maps of archaic hominin ancestry in modern humans, for example, we track only whether a modern human haplotype is in a clade with one or more archaic hominin haplotypes at each site but not which specific archaic hominin haplotype is in the clade. After running the ARG, we computed the discordance between similarity scores per genome haplotype computed using SNP data and computed using shared ancestral recombination events; this discordance should be largely driven by phasing switch errors, which can cause the inference of erroneous ancestral recombination events. We found that this discordance ranged from about 8× higher (in Denisovan) to 2 to 3× higher (in Neanderthals) in archaic hominins than in the most discordant human genomes (table S1).

We merged the phased archaic hominin files with the SGDP data, using *bcftools merge* with the `-missing-to-ref` option, and then used *bcftools norm* to remove duplicate alleles (`-d`). To avoid misidentifying all SGDP samples as homozygous reference at sites that were originally excluded from the SGDP dataset, we limited the variant call set for each chromosome to the sites between the first and last site in the SGDP data on that chromosome. To mitigate the same problem, we also removed any site for which all nonreference alleles in our SGDP data were private to archaic hominins but for which nonreference alleles were present in modern humans within the 1000 Genomes dataset (24). We then discarded all sites for which any individual had a missing genotype or genotype quality below 25 or for which any archaic sample fell within the upper or lower tail of its genome-wide coverage distribution (extracted from the VCF file). The allowed coverage ranges (determined by eye) were 23 to 70× for the Altai Neanderthal, 10 to 43× for the Denisovan, and 10 to 47× for the Vindija33.19 Neanderthal.

Last, we polarized our variant call set into ancestral and derived alleles, using the chimpanzee reference genome *panTro4* (55) [mapped to *hg19* by the UCSC Genome Browser team (56) and downloaded in AXT format] as an ancestral sequence, discarding any variant that was an indel, had more than two alleles, or lacked a known chimpanzee homolog. We chose *panTro4* as an ancestral sequence rather than a composite ancestral sequence as some other studies have done [e.g., (57)] because it allowed us to more easily estimate branch lengths, at the cost of discarding some sites. In addition, because our approach assumes the infinite sites model of mutation, we excluded all CpG dinucleotide sites from analysis, as methylated cytosines in CpG dinucleotides are highly mutable and are thus more likely than other nucleotides to undergo repeated mutations (58).

ARG inference

We developed an ARG inference program called *SARGE* (available at <https://github.com/nkschaefer/sarge>), which is optimized for speed and low memory usage, in addition to making minimal model assumptions. *SARGE* assumes parsimony and the infinite sites model and uses the four-gamete test (11) as a central insight. *SARGE* avoids using statistical techniques to smooth branch lengths or infer clades, opting instead to describe only that which can be inferred directly from the input data. The result is a set of trees that contain polytomies and have relatively low-resolution branch lengths.

Our algorithm centers on the observation that a single tree cannot contain two clades that share members unless one is a superset of the other. We assume that every shared derived allele in our dataset defines a clade. It has been shown that, under this assumption, pairs of sites for which the inferred clades share members but for which neither is a superset of the other, mark ancestral recombination events, or breakpoints between different trees. This is referred to as the “four-haplotype test” or “four-gamete test” (3, 11). One could use this technique to map ancestral recombination events, which mark boundaries between trees, articulate trees using the sites within these boundaries. In practice, however, this can only produce minimally articulated trees. In the case of organisms with low nucleotide diversity, this is because there will not often be enough polymorphic sites between ancestral recombination breakpoints to observe many of the possible clades per tree. In the case of organisms with high nucleotide diversity, however, it will be possible to detect far more ancestral recombination events, thus making

the size of “bins” between ancestral recombination breakpoints smaller and leading to the same problem.

Our algorithm therefore seeks to infer all relevant information about each ancestral recombination event. An ancestral recombination event can be conceptualized as a branch movement (3) and so each consists of a set of haplotypes moving from one clade in an upstream tree into a new clade in a downstream tree. Given two clades that share members, but for which neither is a superset of the other (henceforth described as a failure of the four-haplotype test), and assuming that this four-haplotype test failure describes only one ancestral recombination event, there are then three possible branch movements that can explain it (fig. S1). We refer to the clade in the upstream tree from which a subclade moved as α , the clade in the downstream tree into which a subclade moved as β , and the subclade that moved positions as γ . Four-haplotype test failures are possible between the following sets of clades (with the clade in the upstream tree listed first and the clade in the downstream tree listed second): α/α , α/β , and β/β . In the case of an upward branch movement, all four-haplotype test failures are α/α , and all four-haplotype test failures are of the type β/β in the case of downward branch movements. The members of the moving clade γ can then be inferred once the type of four-haplotype test failure is known. Denoting the members of the upstream clade as U and the members of the downstream clade as D , γ contains $U \setminus D$ in the α/α case, $U \cap D$ in the α/β case, or $D \setminus U$ in the β/β case.

If two trees that differ by an ancestral recombination event are known a priori, then we devised an algorithm to describe the ancestral recombination event by choosing the candidate branch movement(s) (γ clades) that explain the most four-haplotype test failures (see Supplementary Methods and fig. S2). We then created a data structure to store ARG nodes over the genomic ranges they span, using a maximum, preset “propagation distance” parameter that determines the maximum range of a node defined by a single SNP. We then developed an algorithm to continually create and adjust ARG nodes as input SNP data are read in order across the genome (see Supplementary Methods and fig. S3). When four-haplotype test failures are encountered, candidate γ clades are added to the data structure, and when all sites that can inform a given ancestral recombination event have been encountered, the most popular (parsimonious) γ clade is chosen and used to infer new ARG nodes (figs. S4 and S5). Our use of propagation distance parameter serves as a bound on computational complexity and memory requirements (see Supplementary Methods and fig. S6) and limiting the impact of incorrectly inferred clades on the rest of the data structure. SARGE produces a tree at every variable site in the input data, with branch lengths estimated heuristically from the mutations encountered on each branch, normalized relative to the branch length to an outgroup genome (see Supplementary Methods and fig. S7).

Quality control (QC) simulations

For the sake of assessing our and other ARG inference programs, we simulated sampling an increasing number of haplotypes from a single panmictic population with no history of growth or bottlenecks (QC simulations). Our QC simulations were performed using msprime (59). We chose a recombination rate of 1 centimorgan per megabase (cM/Mb) and a mutation rate of 1×10^{-9} per year with a 25-year generation time, giving a per-generation mutation rate of 2.5×10^{-8} . In addition, we chose a heterozygosity value of 10.1 per 10,000 bases, comparable to the rate in modern sub-Saharan Africans (15).

We simulated 1 Mb of sequence per run, running five replicates each of simulations with 50, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000, 2000, 3000, 4000, and 5000 haplotypes. The complete command used was `mspms X 1 -t 1010.0 -r 404.0 1000000 --precision 6 -T`, where X is the number of haplotypes. Whenever there were duplicate base positions in a simulated dataset, we ignored the allele data at all but the first occurrence of each position. We also repeated these simulations with a lower mutation rate of 1×10^{-8} per site per generation, to see how this affected SARGE’s ability to correctly infer clades. The mean sensitivity and specificity given in the Results section are computed on both of these sets of simulations, with numbers of haplotypes ranging from 50 to 500.

We then ran SARGE on each dataset with a propagation distance of 25,000 bases, along with tsinfer (8) (converting its output to a sequence of trees linked to specific variable sites), Relate (7) with the mutation rate set to that known from the simulation and with the haploid effective population size set to two times the simulation’s effective population size, and Rent+ (5) with the `-t` option to infer branch lengths. For each inferred tree, we loaded the tree output by msprime for the same variable site and defined the inferred ARG’s specificity as the fraction of all clades in all inferred trees that existed as clades in the msprime tree at the same sites. We also computed the sensitivity, defined as the fraction of all clades in true trees that existed as clades in the inferred ARGs. For the two methods, Relate and Rent+, that produced fully articulated trees (without polytomies), these two values were the same. Results of these QC simulations are shown in figs. S8 to S10.

We also repeated a similar set of simulations with the goal of assessing the effect of the mutation rate, the recombination rate, and population structure on the accuracy of SARGE, compared to Rent+, Relate, and tsinfer. To this end, we used the same base parameters (with a 2.5×10^{-8} per site per generation mutation rate) and chose 100 haplotypes for each simulation.

For testing the effect of heterozygosity, simulations were of the form `mspms 100 1 -t X -r 202.0 1000000 --precision 6 -T` with $X = 505, 1010, 2020, 5050, \text{ and } 10,100$. We set the effective population size parameter in Relate to two times the effective population size in the simulation (the `-t` parameter divided by 1,000,000 bases) and ran SARGE with a 25-kb propagation distance on each. Results of these simulations are shown in fig. S11.

For testing the effect of the recombination rate, simulations were of the form `mspms 100 1 -t 1010.0 -r X 1000000 --precision 6 -T` with $X = 202, 404, 808, 2020, \text{ and } 4040$. We set the effective population size parameter in Relate to 20,200 and ran SARGE with a 25-kb propagation distance on each. Results of these simulations are shown in fig. S12.

For testing population structure, in addition to the panmictic simulation, we ran one simulation with two populations that split 100 ka ago, another in which these both split into two additional populations 50 ka ago, and a third in which all four populations split into two more populations 25 ka ago, for a total of eight subpopulations. We sampled an equal number of haplotypes from each population for 100 total, except in the case of the eight-population simulation, in which we sampled either 10 or 15 haplotypes from each population, for a total of 100 haplotypes. The command for the two-population simulation was `mspms 100 1 -t 1010.0 -I 2 50 50 -r 404.0 1000000 -ej 0.0990099009901 2 1 --precision 6 -T`, the command for the four-population simulation was `mspms 100 1 -t 1010.0 -I 4 25 25 25 25 -r 404.0 1000000 -ej 0.049504950495 3 1 -ej`

0.049504950495 4 2 -ej 0.0990099009901 2 1 --precision 6 -T, and the command for the eight-population simulation was `mspms 100 1 -t 1010.0 -I 8 10 15 10 15 10 15 10 15 -r 404.0 1000000 -ej 0.0247524752475 5 1 -ej 0.0247524752475 6 3 -ej 0.0247524752475 7 2 -ej 0.0247524752475 8 4 -ej 0.049504950495 3 1 -ej 0.049504950495 4 2 -ej 0.0990099009901 2 1 --precision 6 -T`. We ran Relate with an effective population size parameter of 20,200 and SARGE with a 25-kb propagation distance parameter on each. Results of the simulations are shown in fig. S13.

We did another QC simulation that assessed the effect of a single population bottleneck on the sensitivity and specificity of all programs considered in this study. We chose the same parameters as in other QC simulations but simulated the effect of reducing N_e to $\frac{1}{4}$ its original value 80 ka ago. The command for this simulation was `mspms 100 1 -t 1010.0 -r 4040.0 1000000 --precision 6 -G 17.50197 -eN 0.0792079207921 1 -T`. Results of these simulations are shown in fig. S14.

We also used some of our QC simulations for other assessments. We used our initial QC simulation, including five replicates, with 100 sampled haplotypes to assess the effect of the propagation distance parameter of the accuracy of SARGE compared to tsinfer and Relate (fig. S15 and “Testing the propagation distance parameter” section in Supplementary Text). To show the effects of sample size on the properties of incorrectly inferred clades, we also used one of our QC simulations with 50 sampled haplotypes and another with 500 sampled haplotypes to explore the numbers of clades missing, correctly inferred, and incorrectly inferred in SARGE output trees by clade size (fig. S16) and to find the shortest distance along the simulated chromosome that each incorrectly inferred clade exists from a true clade with the same members (fig. S17 and “Properties of missing and incorrectly inferred clades” section in Supplementary Text).

Demographic simulations

Several of our other analyses and tests required simulated data more similar to our real data (a global panel of human genomes, plus one eastern Neanderthal, one western Neanderthal, and one Denisovan genome). We therefore created a second type of simulated dataset using a popular human demographic model.

We used `scrm` (60) for these simulations, providing branch shortening values for ancient samples. We combined a popular, three-population demographic model for modern humans (61) with populations meant to approximate the Altai (15) and Vindija (1) Neanderthals and the Altai Denisovan (16). We assumed a 1 cM/Mb recombination rate and a 0.5×10^{-9} per year mutation rate, along with a 25-year generation time, giving a per-generation mutation rate of 1.25×10^{-8} . This mutation rate is different from that in the original description of the demographic model but concordant with more recent studies that measured divergence between modern and archaic hominins (1). In addition to the demographic model parameters listed in (61), we modeled a Neanderthal/human split time of 575 ka ago (15), an Altai/Vindija split time of 137.5 ka ago, a Neanderthal/Denisovan split of 420 ka ago, and modeled the heterozygosity in all Neanderthals as 1.6×10^{-4} and in Denisovans as 1.8×10^{-5} (1). We chose 100 ka ago as the divergence time between the Vindija and introgressing Neanderthal, no split between the introgressing and sampled Denisovan populations, and we modeled human/Neanderthal admixture as a single pulse 50 ka ago, in the population ancestral to both Europeans and Asians, and human/Denisovan admixture as a single pulse 20 ka ago, in the Asian population.

We assigned 57 ka ago of missing evolution to the Vindija haplotypes, 123 ka ago of missing evolution to the Altai haplotypes, and 80 ka ago of missing evolution to the Denisovan haplotypes (1). Our simulated chromosome was 25 Mb long, and we sampled 2 haplotypes from each Neanderthal and Denisovan (but not the introgressing Neanderthal), as well as 150 haplotypes from each modern human population (African, Asian, and European) population, for a total of 450 modern human haplotypes. For the sake of simplicity ascertaining archaic introgressed haplotypes, we modeled no archaic introgression into Africans and set the gene flow parameters between Africans and other populations to zero.

We then ran five different simulations, varying the human/archaic admixture proportion in each (the admixture proportion was the same for both the Neanderthal and Denisovan admixture events). The admixture proportions we used were 1, 2, 5, and 10%. The full command for each run was `scrm 456 1 -t 17253.7128713 -r 13802.970297 25000000 -T -I 7 150 150 150 0 0 0 0 -eI q 0 0 0 0 2 0 -eI 0.0891112545728 0 0 0 2 0 0 0 -eI 0.0579585395596 0 0 0 0 0 2 -n 1 1.68 -n 2 3.74 -n 3 7.29 -n 4 0.231834158238 -n 5 0.231834158238 -n 6 0.231834158238 -n 7 0.0260813428018 -eg 0 2 116.010723 -eg 0 3 160.246047 -m 2 3 2.797460 -m 3 2 2.797460 -ej 0.028985 3 2 -en 0.028985 2 0.287184 -em 0.028985 1 2 7.293140 -em 0.028985 2 1 7.293140 -es 0.0144896348899 3 [1-x] -ej 0.0144896348899 8 7 -es 0.03622408722472 [1-x] -ej 0.036224087224795 -ej 0.0724481744495 6 5 -ej 0.197963 2 1 -en 0.303501 1 1 -ej 0.099616239868 5 4 -ej 0.304282332688 7 4 -ej 0.416577003084 4 1`, where [1-x] means one minus the admixture proportion in a given run. We again discarded all but the first instance of every unique base position in the output file, and we converted the “true” trees into SARGE format for running analyses.

We used the simulation with what we deemed was the most realistic admixture proportion (0.02) for other QC assessments of SARGE, including checking accuracy of inferred branch lengths (since SARGE requires knowledge of the TMRCA of all humans to convert branch lengths into years; see “Testing accuracy of inferred branch lengths” section in Supplementary Text), and testing the ability of SARGE, Relate, and tsinfer to recover the true extent of deserts of archaic hominin admixture and ILS across the simulated chromosome (see “Comparing sizes of deserts inferred using different methods” section in Supplementary Methods).

We also ran one smaller version of the demographic simulation, with a 0.05 admixture proportion and 20 haplotypes sampled per human population, for 60 total modern human haplotypes. We used this simulation along with the bigger simulation to assess haplotype block lengths in SARGE-inferred versus true ARGs, to assess the effects of sample size on inferred haplotype block length accuracy (fig. S21). We also used this simulation alongside the larger one to assess accuracy of inferred desert sizes (see “Comparing sizes of deserts inferred using different methods” section in Supplementary Methods and figs. S48 and S49).

Admixture scans

The central challenge of creating admixture maps is to disentangle ILS from admixture. Both processes create local trees in the genome that group candidate admixed haplotypes with admixer haplotypes. Clades resulting from ILS are older than those resulting from admixture; however, they should therefore persist for shorter stretches along the genome and have older TMRCA. Clades resulting from ILS are separable from those resulting from admixture using these

two metrics, although the low resolution of branch lengths in the inferred ARG makes this problem more difficult than when true branch lengths are known (fig. S19). Because of this, we established several filtering steps to distinguish ILS from admixed clades in our dataset.

To map Neanderthal and Denisovan ancestry, we first scanned through ARG output for all clades that grouped some modern human haplotypes with one or more admixer haplotypes (Neanderthal and/or Denisovan) to the exclusion of some other modern human haplotypes. Since SARGE produces many polytomies, this carries the risk of observing a parent of one or more true admixed clades, but not the true admixed clade. This would manifest as a clade containing many modern humans, in addition to one or more archaic hominins, and would falsely be interpreted as a very high-frequency archaic-introgressed haplotype. To mitigate this problem, we defined the Mbuti, Biaka, and Khomani-San genomes as an outgroup population (Africa-MBK) and discarded any clade that contained more than 10% of the outgroup members. We set this 10% threshold, rather than excluding all clades containing outgroup members, because of later filtering steps also designed to eliminate ILS clades. Next, we discarded any clades that persisted for less than 5 kb along the chromosome, as well as one extremely long haplotype that fell within a centromere, as annotated by the UCSC Genome Browser (56). For each clade that passed our selection criteria, we visited each nonarchaic hominin member and determined whether that member had candidate Neanderthal, Denisovan, or undetermined ancestry by assessing whether it was closer (by tree topology, ignoring branch lengths) to a Neanderthal or Denisovan haplotype, or equidistant to both.

We then computed the mean TMRCA between each human member and the candidate archaic introgressor across each haplotype. SARGE reports each branch length as a percent of the total height of the tree both above and below that branch (fig. S7), which in our dataset includes all fixed differences between the genomes in our dataset and the reference chimpanzee genome. The units in which these TMRCA are expressed are therefore a percent of two times the total divergence time between humans and chimpanzees. These values were then corrected for branch shortening, according to the mean branch shortening or “missing evolution” values published in a prior study (1), converted to branch units by dividing each number by 26 Ma. To correct for branch shortening, we added the amount of missing evolution reported for a given archaic genome (in the same units as ARG branch lengths) to the archaic genome’s branch when computing TMRCA. That is, if the TMRCA between a modern and archaic haplotype is x and the branch shortening value for the archaic haplotype is y , then the corrected TMRCA between them is $(2x + y)/2$.

Our ascertainment strategy allowed some candidate admixed haplotypes to intersect. If a human haplotype was inferred to have Neanderthal ancestry at one site and Denisovan ancestry at the next variable site, for instance, then it is likely that the two haplotypes actually originated from the same source. Because of this possibility, we merged all these haplotypes together, averaging the TMRCA to admixer in each, weighting by the relative lengths of the two haplotypes being merged. We then repeated this process until each haplotype was unable to merge with any others. Because each such haplotype resulted from the combination of haplotypes with different types of ancestry, all merged haplotypes were labeled as ambiguous, rather than Neanderthal or Denisovan, origin.

At this stage, the set of haplotypes likely still contained haplotypes resulting from ILS rather than admixture. To help mitigate this, we assigned each a score designed to be high in cases of admixture and low in cases of ILS. We chose a date that predates the human/archaic hominin separation, 700 ka ago, and computed the P value of each haplotype having originated at that time or earlier according to its length and according to its TMRCA to the candidate admixer. For this, we assumed both TMRCA to admixer and haplotype lengths to be exponentially distributed, and we assumed neutral evolution, a standard recombination rate of 1 cM/Mb, a 25-year generation time, and 13 Ma ago human-chimpanzee divergence. We assigned each haplotype a probability of resulting from ILS, based on its TMRCA: If the TMRCA (as a percent of the height of the tree) is y , then $p_{\text{TMRCA}} = 1 - e^{-\left(\frac{2666}{700,000}\right)(y)}$. We then assigned each a probability of resulting from ILS, based on its length: If x is the length of a haplotype, in bases, then $p_{\text{len}} = e^{-\frac{700,000}{1.68 \times 10^{25}}x}$. We then computed a score, using a pseudo-count of 0.001 to avoid taking the logarithm of 0: $S_{\text{adm}} = -(\log_{10}(p_{\text{TMRCA}} + 0.001) + \log_{10}(p_{\text{len}} + 0.001))$.

Because our prior filtering strategies already removed many ILS haplotypes, we sought to find the smallest threshold for S_{adm} that gave reasonable results. We therefore tested S_{adm} cutoffs at every fifth percentile of the distribution across all haplotypes and computed the coefficient of variation (SD divided by mean) for TMRCA to admixer in each class of archaic-introgressed haplotypes (Neanderthal, Denisovan, and ambiguous) within each human population defined within the SGDP (14). Because we expect members of these populations to have somewhat shared histories, we expected the coefficient of variation to the admixer to decrease and level out as an appropriate cutoff was reached, reflecting the removal of highly variable segments with old TMRCA (ILS). We found such a cutoff at the 30th percentile of the score distribution and excluded haplotypes with scores below this threshold (fig. S24). Because our resulting archaic introgression maps still underestimated archaic ancestry per genome relative to the D statistic (Fig. 3A) and because our real dataset overestimated the extent of ILS relative to every demographic simulation we ran and underestimated admixture relative to the size of inferred deserts (Fig. 4), we believe that this cutoff was reasonable, if not overly conservative.

We also quantified uncertainty in our percent admixture estimates using the weighted block jackknife technique (18) with 10-Mb blocks. We generated windows using BEDTools, omitted each from our admixture maps in turn, and recalculated the percent of each type of archaic ancestry in each SGDP genome from our admixture maps. We then used this distribution of archaic ancestry estimates for each individual to obtain a 95% CI. Mean percent archaic ancestry, along with the minimum lower and maximum upper bound of the 95% CI for each type of ancestry across all genomes in each population are given in table S2.

Because our method relies on phased data, improper phasing could pose problems for this by breaking haplotypes where they should not be broken. Improper phasing is most likely to be a problem in the archaic genomes, for which reads were short and there is no phasing reference panel, so many sites were phased randomly, after phasing the others using a (likely inadequate) modern human reference panel (see “Data processing” section above). We worked around this issue by only conducting admixture scans on modern human genomes, those more likely to be properly phased, and by tracking only whether each candidate admixed clade contained Neanderthal or Denisovan haplotypes, without regard to which

specific Neanderthal or Denisovan haplotypes they contained. That is, a clade containing one modern human genome haplotype and one Neanderthal haplotype would still be considered the same clade (potentially from the same admixture event) if at the next site the Neanderthal haplotype contained within was replaced by another. To ensure that this did not negatively affect our results, we repeated some admixture scans on simulated data after intentionally introducing phasing errors into the archaic genomes; we did not see evidence that this negatively affected the results (see Supplementary Text).

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/7/29/eabc0776/DC1>

[View/request a protocol for this paper from Bio-protocol.](#)

REFERENCES AND NOTES

- K. Prüfer, C. de Filippo, S. Grote, F. Mafessoni, P. Korlević, M. Hajdinjak, B. Vernot, L. Skov, P. Hsieh, S. Peyrégne, D. Reher, C. Hofse, S. Nagel, T. Maricic, Q. Fu, C. Theunert, R. Rogers, P. Skoglund, M. Chintalapati, M. Dannemann, B. J. Nelson, F. M. Key, P. Rudan, Ž. Kučan, I. Gušić, L. V. Golovanova, V. B. Doronichev, N. Patterson, D. Reich, E. E. Eichler, M. Slatkin, M. H. Schierup, A. M. Andrés, J. Kelso, M. Meyer, S. Pääbo, A high-coverage Neanderthal genome from Vindija Cave in Croatia. *Science* **358**, 655–658 (2017).
- R. C. Griffiths, P. Marjoram, Ancestral inference from samples of DNA sequences with recombination. *J. Comput. Biol.* **3**, 479–502 (1996).
- Y. S. Song, J. Hein, Constructing minimal ancestral recombination graphs. *J. Comput. Biol.* **12**, 147–169 (2005).
- M. D. Rasmussen, M. J. Hubisz, I. Gronau, A. Siepel, Genome-wide inference of ancestral recombination graphs. *PLoS Genet.* **10**, e1004342 (2014).
- S. Mirzaei, Y. Wu, RENT+: An improved method for inferring local genealogical trees from haplotypes with recombination. *Bioinformatics* **33**, 1021–1030 (2017).
- M. J. Minichiello, R. Durbin, Mapping trait loci by use of inferred ancestral recombination graphs. *Am. J. Hum. Genet.* **79**, 910–922 (2006).
- L. Speidel, M. Forest, S. Shi, S. R. Myers, A method for genome-wide genealogy estimation for thousands of samples. *Nat. Genet.* **51**, 1321–1329 (2019).
- J. Kelleher, Y. Wong, A. W. Wohns, C. Fadil, P. K. Albers, G. McVean, Inferring whole-genome histories in large population datasets. *Nat. Genet.* **51**, 1330–1338 (2019).
- S. R. Browning, B. L. Browning, Y. Zhou, S. Tucci, J. M. Akey, Analysis of human sequence data reveals two pulses of archaic denisovan admixture. *Cell* **173**, 53–61.e9 (2018).
- G. S. Jacobs, G. Hudjashov, L. Saag, P. Kusuma, C. C. Darusallam, D. J. Lawson, M. Mondal, L. Pagani, F.-X. Ricaut, M. Stoneking, M. Metspalu, H. Sudoyo, J. S. Lansing, M. P. Cox, Multiple deeply divergent denisovan ancestries in Papuans. *Cell* **177**, 1010–1021.e32 (2019).
- R. R. Hudson, N. L. Kaplan, Statistical properties of the number of recombination events in the history of a sample of DNA sequences. *Genetics* **111**, 147–164 (1985).
- S. R. Myers, R. C. Griffiths, Bounds on the minimum number of recombination events in a sample history. *Genetics* **163**, 375–394 (2003).
- V. Bafna, V. Bansal, The number of recombination events in a sample history: Conflict graph and lower bounds. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **1**, 78–90 (2004).
- S. Mallick, H. Li, M. Lipson, I. Mathieson, M. Gymrek, F. Racimo, M. Zhao, N. Chennagiri, S. Nordenfelt, A. Tandon, P. Skoglund, I. Lazaridis, S. Sankararaman, Q. Fu, N. Rohland, G. Renaud, Y. Erlich, T. Willems, C. Gallo, J. P. Spence, Y. S. Song, G. Poletti, F. Balloux, G. van Driem, P. de Knijff, I. G. Romero, A. R. Jha, D. M. Behar, C. M. Bravi, C. Capelli, T. Hervig, A. Moreno-Estrada, O. L. Posukh, E. Balanovska, O. Balanovsky, S. Karachanak-Yankova, H. Sahakyan, D. Toncheva, L. Yepiskoposyan, C. Tyler-Smith, Y. Xue, M. S. Abdullah, A. Ruiz-Linares, C. M. Beall, A. di Rienzo, C. Jeong, E. B. Starikovskaya, E. Metspalu, J. Parik, R. Villems, B. M. Henn, U. Hodoglugil, R. Mahley, A. Sajantila, G. Stamatoyannopoulos, J. T. S. Wee, R. Khushainova, E. Khusnutdinova, S. Litvinov, G. Ayodo, D. Comas, M. F. Hammer, T. Kivisild, W. Klitz, C. A. Winkler, D. Labuda, M. Bamshad, L. B. Jorde, S. A. Tishkoff, W. S. Watkins, M. Metspalu, S. Dryomov, R. Sukernik, L. Singh, K. Thangaraj, S. Pääbo, J. Kelso, N. Patterson, D. Reich, The Simons Genome Diversity Project: 300 genomes from 142 diverse populations. *Nature* **538**, 201–206 (2016).
- K. Prüfer, F. Racimo, N. Patterson, F. Jay, S. Sankararaman, S. Sawyer, A. Heinze, G. Renaud, P. H. Sudmant, C. de Filippo, H. Li, S. Mallick, M. Dannemann, Q. Fu, M. Kircher, M. Kuhlwiilm, M. Lachmann, M. Meyer, M. Ongyerth, M. Siebauer, C. Theunert, A. Tandon, P. Moorjani, J. Pickrell, J. C. Mullikin, S. H. Vohr, R. E. Green, I. Hellmann, P. L. F. Johnson, H. Blanche, H. Cann, J. O. Kitzman, J. Shendure, E. E. Eichler, E. S. Lein, T. E. Bakken, L. V. Golovanova, V. B. Doronichev, M. V. Shunkov, A. P. Dereviianko, B. Viola, M. Slatkin, D. Reich, J. Kelso, S. Pääbo, The complete genome sequence of a Neanderthal from the Altai Mountains. *Nature* **505**, 43–49 (2014).
- M. Meyer, M. Kircher, M. T. Gansauge, H. Li, F. Racimo, S. Mallick, J. G. Schraiber, F. Jay, K. Prüfer, C. de Filippo, P. H. Sudmant, C. Alkan, Q. Fu, R. Do, N. Rohland, A. Tandon, M. Siebauer, R. E. Green, K. Bryc, A. W. Briggs, U. Stenzel, J. Dabney, J. Shendure, J. Kitzman, M. F. Hammer, M. V. Shunkov, A. P. Dereviianko, N. Patterson, A. M. Andres, E. E. Eichler, M. Slatkin, D. Reich, J. Kelso, S. Paabo, A high-coverage genome sequence from an archaic Denisovan individual. *Science* **338**, 222–226 (2012).
- International HapMap Consortium, K. A. Frazer, D. G. Ballinger, D. R. Cox, D. A. Hinds, L. L. Stuve, R. A. Gibbs, J. W. Belmont, A. Boudreau, P. Hardenbol, S. M. Leal, S. Pasternak, D. A. Wheeler, T. D. Willis, F. Yu, H. Yang, C. Zeng, Y. Gao, H. Hu, W. Hu, C. Li, W. Lin, S. Liu, H. Pan, X. Tang, J. Wang, W. Wang, J. Yu, B. Zhang, Q. Zhang, H. Zhao, H. Zhao, J. Zhou, S. B. Gabriel, R. Barry, B. Blumenstiel, A. Camargo, M. Defelice, M. Faggart, M. Goyette, S. Gupta, J. Moore, H. Nguyen, R. C. Onofrio, M. Parkin, J. Roy, E. Stahl, E. Winchester, L. Ziaugra, D. Altshuler, Y. Shen, Z. Yao, W. Huang, X. Chu, Y. He, L. Jin, Y. Liu, Y. Shen, W. Sun, H. Wang, Y. Wang, Y. Wang, X. Xiong, L. Xu, M. M. Y. Wayne, S. K. W. Tsui, H. Xue, J. T.-F. Wong, L. M. Galver, J.-B. Fan, K. Gunderson, S. S. Murray, A. R. Oliphant, M. S. Chee, A. Montpetit, F. Chagnon, V. Ferretti, M. Leboeuf, J.-F. Olivier, M. S. Phillips, S. Roumy, C. Sallée, A. Verner, T. J. Hudson, P.-Y. Kwok, D. Cai, D. C. Koboldt, R. D. Miller, L. Pawlikowska, P. Taillon-Miller, M. Xiao, L.-C. Tsui, W. Mak, Y. Q. Song, P. K. H. Tam, Y. Nakamura, T. Kawaguchi, T. Kitamoto, M. Morizono, A. Nagashima, Y. Ohnishi, A. Sekine, T. Tanaka, T. Tsunoda, P. Deloukas, C. P. Bird, M. Delgado, E. T. Dermitzakis, R. Gwilliam, S. Hunt, J. Morrison, D. Powell, B. E. Stranger, P. Whittaker, D. R. Bentley, M. J. Daly, P. I. W. de Bakker, J. Barrett, Y. R. Chretien, J. Maller, S. M. Carroll, N. Patterson, I. Pe'er, A. Price, S. Purcell, D. J. Richter, P. Sabeti, R. Saxena, S. F. Schaffner, P. C. Sham, P. Varily, D. Altshuler, L. D. Stein, L. Krishnan, A. V. Smith, M. K. Tello-Ruiz, G. A. Thorisson, A. Chakravarti, P. E. Chen, D. J. Cutler, C. S. Kashuk, S. Lin, G. R. Abecasis, W. Guan, Y. Li, H. M. Munro, Z. S. Qin, D. J. Thomas, G. M. Vean, A. Auton, L. Bottolo, N. Cardin, S. Eyheramendy, C. Freeman, J. Marchini, S. Myers, C. Spencer, M. Stephens, P. Donnelly, L. R. Cardon, G. Clarke, D. M. Evans, A. P. Morris, B. S. Weir, T. Tsunoda, J. C. Mullikin, S. T. Sherry, M. Feolo, A. Skol, H. Zhang, C. Zeng, H. Zhao, I. Matsuda, Y. Fukushima, D. R. Macer, E. Suda, C. N. Rotimi, C. A. Adebanowo, I. Ajayi, T. Aniagwu, P. A. Marshall, C. Nkwodimmah, C. D. M. Royal, M. F. Leppert, M. Dixon, A. Peiffer, R. Qiu, A. Kent, K. Kato, N. Niikawa, I. F. Adewole, B. M. Knoppers, M. W. Foster, E. W. Clayton, J. Watkin, R. A. Gibbs, J. W. Belmont, D. Muzny, L. Nazareth, E. Sodergren, G. M. Weinstock, D. A. Wheeler, I. Yakub, S. B. Gabriel, R. C. Onofrio, D. J. Richter, L. Ziaugra, B. W. Birren, M. J. Daly, D. Altshuler, R. K. Wilson, L. L. Fulton, J. Rogers, J. Burton, N. P. Carter, C. M. Clee, M. Griffiths, M. C. Jones, K. M. Lay, R. W. Plumb, M. T. Ross, S. K. Sims, D. L. Willey, Z. Chen, H. Han, L. Kang, M. Godbout, J. C. Wallenburg, P. L'Archevêque, G. Bellemare, K. Saeki, H. Wang, D. An, H. Fu, Q. Li, Z. Wang, R. Wang, A. L. Holden, L. D. Brooks, J. E. M. Ewen, M. S. Guyer, V. O. Wang, J. L. Peterson, M. Shi, J. Spiegel, L. M. Sung, L. F. Zacharia, F. S. Collins, K. Kennedy, R. Jamieson, J. Stewart, A second generation human haplotype map of over 3.1 million SNPs. *Nature* **449**, 851–861 (2007).
- R. E. Green, J. Krause, A. W. Briggs, T. Maricic, U. Stenzel, M. Kircher, N. Patterson, H. Li, W. Zhai, M. H. Y. Fritz, N. F. Hansen, E. Y. Durand, A. S. Malaspinas, J. D. Jensen, T. Marques-Bonet, C. Alkan, K. Prüfer, M. Meyer, H. A. Burbano, J. M. Good, R. Schultz, A. Aximu-Petri, B. Hober, B. Hoffner, M. Siegemund, A. Weihmann, C. Nussbaum, E. S. Lander, C. Russ, N. Novod, J. Affourtit, M. Egholm, C. Verna, P. Rudan, D. Brajkovic, Z. Kucan, I. Gusic, V. B. Doronichev, L. V. Golovanova, C. Lalueza-Fox, M. de la Rasilla, J. Fortea, A. Rosas, R. W. Schmitz, P. L. F. Johnson, E. E. Eichler, D. Falush, E. Birney, J. C. Mullikin, M. Slatkin, R. Nielsen, J. Kelso, M. Lachmann, D. Reich, S. Paabo, A draft sequence of the Neanderthal genome. *Science* **328**, 710–722 (2010).
- E. Y. Durand, N. Patterson, D. Reich, M. Slatkin, Testing for ancient admixture between closely related populations. *Mol. Biol. Evol.* **28**, 2239–2252 (2011).
- L. Chen, A. B. Wolf, W. Fu, L. Li, J. M. Akey, Identifying and interpreting apparent Neanderthal ancestry in African individuals. *Cell* **180**, 677–687.e16 (2020).
- B. Vernot, J. M. Akey, Resurrecting surviving Neanderthal lineages from modern human genomes. *Science* **343**, 1017–1021 (2014).
- S. Sankararaman, S. Mallick, M. Dannemann, K. Prüfer, J. Kelso, S. Pääbo, N. Patterson, D. Reich, The genomic landscape of Neanderthal ancestry in present-day humans. *Nature* **507**, 354–357 (2014).
- B. Vernot, S. Tucci, J. G. Schraiber, A. B. Wolf, R. M. Gitterman, M. Dannemann, S. Grote, R. C. McCoy, H. Norton, L. B. Scheinfeldt, D. A. Merriwether, G. Koki, J. S. Friedlaender, J. Wakefield, S. Paabo, J. M. Akey, Excavating Neanderthal and Denisovan DNA from the genomes of Melanesian individuals. *Science* **352**, 235–239 (2016).
- 1000 Genomes Project Consortium, A. Auton, L. D. Brooks, R. M. Durbin, E. P. Garrison, H. M. Kang, J. O. Korbel, J. L. Marchini, S. M.Carthy, G. A. M. Vean, G. R. Abecasis, A global reference for human genetic variation. *Nature* **526**, 68–74 (2015).
- M. F. Hammer, A. E. Woerner, F. L. Mendez, J. C. Watkins, J. D. Wall, Genetic evidence for archaic admixture in Africa. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 15123–15128 (2011).
- J. K. Pickrell, N. Patterson, P. R. Loh, M. Lipson, B. Berger, M. Stoneking, B. Pakendorf, D. Reich, Ancient west Eurasian ancestry in southern and eastern Africa. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 2632–2637 (2014).

27. C. M. Schlebusch, H. Malmström, T. Günther, P. Sjödin, A. Coutinho, H. Edlund, A. R. Munters, M. Vicente, M. Steyn, H. Soodyall, M. Lombard, M. Jakobsson, Southern African ancient genomes estimate modern human divergence to 350,000 to 260,000 years ago. *Science* **358**, 652–655 (2017).
28. F. Mafessoni, S. Grote, C. de Filippo, V. Slon, K. A. Kolobova, B. Viola, S. V. Markin, M. Chintalapati, S. Peyrégne, L. Skov, P. Skoglund, A. I. Krivoschapkin, A. P. Derevianko, M. Meyer, J. Kelso, B. Peter, K. Prüfer, S. Pääbo, A high-coverage Neandertal genome from Chagyrskaya Cave. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 15132–15136 (2020).
29. B. Vernot, J. M. Akey, Complex history of admixture between modern humans and Neandertals. *Am. J. Hum. Genet.* **96**, 448–453 (2015).
30. F. A. Villanea, J. G. Schraiber, Multiple episodes of interbreeding between Neandertal and modern humans. *Nat. Ecol. Evol.* **3**, 39–44 (2019).
31. S. Sankararaman, S. Mallick, N. Patterson, D. Reich, The combined landscape of Denisovan and Neandertal ancestry in present-day humans. *Curr. Biol.* **26**, 1241–1247 (2016).
32. K. Harris, R. Nielsen, The genetic cost of Neandertal introgression. *Genetics* **203**, 881–891 (2016).
33. P. Moorjani, C. E. G. Amorim, P. F. Arndt, M. Przeworski, Variation in the molecular clock of primates. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 10607–10612 (2016).
34. M. Kuhlwiilm, C. Boeckx, A catalog of single nucleotide changes distinguishing modern humans from archaic hominins. *Sci. Rep.* **9**, 8463 (2019).
35. D. Szklarczyk, J. H. Morris, H. Cook, M. Kuhn, S. Wyder, M. Simonovic, A. Santos, N. T. Doncheva, A. Roth, P. Bork, L. J. Jensen, C. von Mering, The STRING database in 2017: Quality-controlled protein-protein association networks, made broadly accessible. *Nucleic Acids Res.* **45**, D362–D368 (2017).
36. S. Fan, D. E. Kelly, M. H. Beltrame, M. E. B. Hansen, S. Mallick, A. Ranciaro, J. Hirbo, S. Thompson, W. Beggs, T. Nyambo, S. A. Omar, D. W. Meskel, G. Belay, A. Froment, N. Patterson, D. Reich, S. A. Tishkoff, African evolutionary history inferred from whole genome sequence data of 44 indigenous African populations. *Genome Biol.* **20**, 82 (2019).
37. E. K. F. Chan, A. Timmermann, B. F. Baldi, A. E. Moore, R. J. Lyons, S. S. Lee, A. M. F. Kalsbeek, D. C. Petersen, H. Rautenbach, H. E. A. Förtsch, M. S. R. Bornman, V. M. Hayes, Human origins in a southern African palaeo-wetland and first migrations. *Nature* **575**, 185–189 (2019).
38. J.-J. Hublin, A. Ben-Ncer, S. E. Bailey, S. E. Freidline, S. Neubauer, M. M. Skinner, I. Bergmann, A. le Cabec, S. Benazzi, K. Harvati, P. Gunz, New fossils from Jebel Irhoud, Morocco and the pan-African origin of Homo sapiens. *Nature* **546**, 289–292 (2017).
39. G. Bräuer, The origin of modern anatomy: By speciation or intraspecific evolution? *Evol. Anthropol. Issues News Rev.* **17**, 22–37 (2008).
40. J. J. Vitti, S. R. Grossman, P. C. Sabeti, Detecting natural selection in genomic data. *Annu. Rev. Genet.* **47**, 97–120 (2013).
41. Gene Ontology Consortium, Gene Ontology Consortium: Going forward. *Nucleic Acids Res.* **43**, D1049–D1056 (2015).
42. S. Peyrégne, M. J. Boyle, M. Dannemann, K. Prüfer, Detecting ancient positive selection in humans using extended lineage sorting. *Genome Res.* **27**, 1563–1572 (2017).
43. J. Ule, A. Ule, J. Spencer, A. Williams, J. S. Hu, M. Cline, H. Wang, T. Clark, C. Fraser, M. Ruggiu, B. R. Zeeberg, D. Kane, J. N. Weinstein, J. Blume, R. B. Darnell, Nova regulates brain-specific splicing to shape the synapse. *Nat. Genet.* **37**, 844–852 (2005).
44. B.-I. Bae, D. Jayaraman, C. A. Walsh, Genetic changes shaping the human brain. *Dev. Cell* **32**, 423–434 (2015).
45. J. Montalbano, W. Jin, M. S. Sheikh, Y. Huang, RBEL1 is a novel gene that encodes a nucleocytoplasmic Ras superfamily GTP-binding protein and is overexpressed in breast cancer. *J. Biol. Chem.* **282**, 37640–37649 (2007).
46. S. M. Cutts, K. J. Fowler, B. T. Kile, L. L. Hii, R. A. O'Dowd, D. F. Hudson, R. Saffery, P. Kalitsis, E. Earle, K. H. Choo, Defective chromosome segregation, microtubule bundling and nuclear bridging in inner centromere protein gene (Incenp)-disrupted mice. *Hum. Mol. Genet.* **8**, 1145–1155 (1999).
47. D. E. Koser, A. J. Thompson, S. K. Foster, A. Dwivedy, E. K. Pillai, G. K. Sheridan, H. Svoboda, M. Viana, L. d. F. Costa, J. Guck, C. E. Holt, K. Franze, Mechanosensing is critical for axon growth in the developing brain. *Nat. Neurosci.* **19**, 1592–1598 (2016).
48. J. Brasch, K. M. Goodman, A. J. Noble, M. Rapp, S. Mannepalli, F. Bahna, V. P. Dandey, T. Beppler, B. Berger, T. Maniatis, C. S. Potter, B. Carragher, B. Honig, L. Shapiro, Visualization of clustered protocadherin neuronal self-recognition complexes. *Nature* **569**, 280–283 (2019).
49. B. Y. Kim, K. E. Lohmueller, Selection and reduced population size cannot explain higher amounts of Neandertal ancestry in East Asian than in European human populations. *Am. J. Hum. Genet.* **96**, 454–261 (2015).
50. O. Delaneau, J. Marchini; 1000 Genomes Project Consortium; 1000 Genomes Project Consortium, Integrating sequence and array data to create an improved 1000 Genomes Project haplotype reference panel. *Nat. Commun.* **5**, 3934 (2014).
51. K. Prüfer, snpAD: An ancient DNA genotype caller. *Bioinformatics* **34**, 4165–4171 (2018).
52. M. Patterson, T. Marschall, N. Pisanti, L. van Iersel, L. Stougie, G. W. Klau, A. Schönhuth, WhatsHap: Weighted haplotype assembly for future-generation sequencing reads. *J. Comput. Biol.* **22**, 498–509 (2015).
53. H. Li, A statistical framework for SNP calling, mutation discovery, association mapping and population genetic parameter estimation from sequencing data. *Bioinformatics* **27**, 2987–2993 (2011).
54. P.-R. Loh, P. Danecek, P. F. Palamara, C. Fuchsberger, Y. A. Reshef, H. K. Finucane, S. Schoenherr, L. Forer, S. McCarthy, G. R. Abecasis, R. Durbin, A. L. Price, Reference-based phasing using the Haplotype Reference Consortium panel. *Nat. Genet.* **48**, 1443–1448 (2016).
55. Chimpanzee Sequencing and Analysis Consortium, Initial sequence of the chimpanzee genome and comparison with the human genome. *Nature* **437**, 69–87 (2005).
56. J. Casper, A. S. Zweig, C. Villarreal, C. Tyner, M. L. Speir, K. R. Rosenbloom, B. J. Raney, C. M. Lee, B. T. Lee, D. Karolchik, A. S. Hinrichs, M. Haussler, L. Gurusvadoo, J. Navarro Gonzalez, D. Gibson, I. T. Fiddes, C. Eisenhart, M. Diekhans, H. Clawson, G. P. Barber, J. Armstrong, D. Haussler, R. M. Kuhn, W. J. Kent, The UCSC Genome Browser database: 2018 update. *Nucleic Acids Res.* **46**, D762–D769 (2018).
57. R. M. Durbin, D. Altshuler, A. Auton, A map of human genome variation from population-scale sequencing. *Nature* **467**, 1061–1073 (2010).
58. M. Ehrlich, R. Y. Wang, 5-Methylcytosine in eukaryotic DNA. *Science* **212**, 1350–1357 (1981).
59. J. Kelleher, A. M. Etheridge, G. McVean, Efficient coalescent simulation and genealogical analysis for large sample sizes. *PLoS Comput. Biol.* **12**, e1004842 (2016).
60. P. R. Staab, S. Zhu, D. Metzler, G. Lunter, scrm: Efficiently simulating long sequences using the approximated coalescent with recombination. *Bioinformatics* **31**, 1680–1682 (2015).
61. R. N. Gutenkunst, R. D. Hernandez, S. H. Williamson, C. D. Bustamante, Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genet.* **5**, e1000695 (2009).
62. A. R. Quinlan, I. M. Hall, BEDTools: A flexible suite of utilities for comparing genomic features. *Bioinformatics* **26**, 841–842 (2010).
63. A. Khan, A. Mathelier, Intervene: A tool for intersection and visualization of multiple gene or genomic region sets. *BMC Bioinformatics* **18**, 287 (2017).
64. K. J. Karczewski, L. C. Francioli, G. Tiao, B. B. Cummings, J. Alfoldi, Q. Wang, R. L. Collins, K. M. Laricchia, A. Ganna, D. P. Birnbaum, L. D. Gauthier, H. Brand, M. Solomonson, N. A. Watts, D. Rhodes, M. Singer-Berk, E. M. England, E. G. Seaby, J. A. Kosmicki, R. K. Walters, K. Tashman, Y. Farjoun, E. Banks, T. Poterba, A. Wang, C. Seed, N. Whiffin, J. X. Chong, K. E. Samocha, E. Pierce-Hoffman, Z. Zappala, Anne H. O'Donnell-Luria, E. V. Minikel, B. Weisburd, M. Lek, J. S. Ware, C. Vittal, I. M. Armean, H. Bergelson, K. Cibulskis, K. M. Connolly, M. Covarrubias, S. Donnelly, S. Ferreira, S. Gabriel, J. Gentry, N. Gupta, T. Jeandet, D. Kaplan, C. Llanwarne, R. Mounshi, S. Novod, N. Petrillo, D. Roazen, V. Ruano-Rubio, A. Saltzman, M. Schleicher, J. Soto, K. Tibbetts, C. Tolonen, G. Wade, M. E. Talkowski; The Genome Aggregation Database Consortium, B. M. Neale, M. J. Daly, D. G. Mac Arthur, Variation across 141,456 human exomes and genomes reveals the spectrum of loss-of-function intolerance across human protein-coding genes. bioRxiv 531210 [Preprint]. 13 August 2019. <https://doi.org/10.1101/531210>.
65. A. Favorov, L. Mularoni, A. M. Cope, Y. Medvedeva, A. A. Mironov, V. J. Makeev, S. J. Wheelan, Exploring massive, genome scale datasets with the GenometriCorr package. *PLoS Comput. Biol.* **8**, e1002529 (2012).
66. A. Frankish, M. Diekhans, A. M. Ferreira, R. Johnson, I. Jungreis, J. Loveland, J. M. Mudge, C. Sisu, J. Wright, J. Armstrong, I. Barnes, A. Berry, A. Bignell, S. Carbonell Sala, J. Chrast, F. Cunningham, T. di Domenico, S. Donaldson, I. T. Fiddes, C. García Girón, J. M. Gonzalez, T. Grego, M. Hardy, T. Hourlier, T. Hunt, O. G. Izuogu, J. Lagarde, F. J. Martin, L. Martínez, S. Mohanan, P. Muir, F. C. P. Navarro, A. Parker, B. Pei, F. Pozo, M. Ruffier, B. M. Schmitt, E. Stapleton, M. M. Suner, I. Sycheva, B. Uszczyńska-Ratajczak, J. Xu, A. Yates, D. Zerbino, Y. Zhang, B. Aken, J. S. Choudhary, M. Gerstein, R. Guigó, T. J. P. Hubbard, M. Kellis, B. Paten, A. Reymond, M. L. Tress, P. Flicek, GENCODE reference annotation for the human and mouse genomes. *Nucleic Acids Res.* **47**, D766–D773 (2019).
67. S. Fishilevich, R. Nudel, N. Rappaport, R. Hadar, I. Plaschkes, T. I. Stein, N. Rosen, A. Kohn, M. Twik, M. Safran, D. Lancet, D. Cohen, GeneHancer: Genome-wide integration of enhancers and target genes in GeneCards. *Database (Oxford)* **2017**, (2017). <https://doi.org/10.1093/database/bax028>
68. M. Ashburner, C. A. Ball, J. A. Blake, D. Botstein, H. Butler, J. M. Cherry, A. P. Davis, K. Dolinski, S. S. Dwight, J. T. Eppig, M. A. Harris, D. P. Hill, L. Issel-Tarver, A. Kasarskis, S. Lewis, J. C. Matese, J. E. Richardson, M. Ringwald, G. M. Rubin, G. Sherlock, Gene ontology: Tool for the unification of biology. The Gene Ontology Consortium. *Nat. Genet.* **25**, 25–29 (2000).
69. G. T. E. Consortium; Laboratory, Data Analysis & Coordinating Center (LDACC)—Analysis Working Group; Statistical Methods Group—Analysis Working Group; Enhancing GTEX (eGTEX) groups; NIH Common Fund; NIH/NCI; NIH/NHGRI; NIH/NIMH; NIH/NIDA; Biospecimen Collection Source Site—NDR; Biospecimen Collection Source Site—RP; Biospecimen Core Resource—VARI; Brain Bank Repository—University of Miami Brain

- Endowment Bank; Leidos Biomedical—Project Management; ELSI Study; Genome Browser Data Integration & Visualization—EBI; Genome Browser Data Integration & Visualization—UCSC Genomics Institute, University of California Santa Cruz; Lead analysts; Laboratory, Data Analysis & Coordinating Center (LDACC); NIH program management; Biospecimen collection; Pathology; eQTL manuscript working group, A. Battle, C. D. Brown, B. E. Engelhardt, S. B. Montgomery, Genetic effects on gene expression across human tissues. *Nature* **550**, 204–213 (2017).
70. N. Kryuchkova-Mostacci, M. Robinson-Rechavi, A benchmark of gene expression tissue-specificity metrics. *Brief. Bioinform.* **18**, 205–214 (2017).
 71. K. Prüfer, B. Muetzel, H. H. Do, G. Weiss, P. Khaitovich, E. Rahm, S. Pääbo, M. Lachmann, W. Enard, FUNC: A package for detecting significant associations between gene sets and ontological annotations. *BMC Bioinformatics* **8**, 41 (2007).
 72. A. W. Wohns, Y. Wong, B. Jeffery, A. Akbari, S. Mallick, R. Pinhasi, N. Patterson, D. Reich, J. Kelleher, G. M. Vean, A unified genealogy of modern and ancient genomes. *bioRxiv* 2021.02.16.431497 [Preprint]. 17 February 2021. <https://doi.org/10.1101/2021.02.16.431497>.
 73. V. M. Narasimhan, R. Rahbari, A. Scally, A. Wuster, D. Mason, Y. Xue, J. Wright, R. C. Trembath, E. R. Maher, D. A. van Heel, A. Auton, M. E. Hurles, C. Tyler-Smith, R. Durbin, Estimating the human mutation rate from autozygous segments reveals population differences in human mutational processes. *Nat. Commun.* **8**, 303 (2017).
 74. K. Harris, Evidence for recent, population-specific evolution of the human mutation rate. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 3439–3444 (2015).
 75. K. E. Langergraber, K. Prüfer, C. Rowney, C. Boesch, C. Crockford, K. Fawcett, E. Inoue, M. Inoue-Muruyama, J. C. Mitani, M. N. Muller, M. M. Robbins, G. Schubert, T. S. Stoinski, B. Viola, D. Watts, R. M. Wittig, R. W. Wrangham, K. Zuberbühler, S. Paabo, L. Vigilant, Generation times in wild chimpanzees and gorillas suggest earlier divergence times in great ape and human evolution. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 15716–15721 (2012).
 76. G. C. Gibb, S. F. K. Hills, Intergenerational mutation rate does not equal long-term evolutionary substitution rate. *Proc. Natl. Acad. Sci. U.S.A.* **110**, E611–E611 (2013).
 77. N. Patterson, D. J. Richter, S. Gnerre, E. S. Lander, D. Reich, Genetic evidence for complex speciation of humans and chimpanzees. *Nature* **441**, 1103–1108 (2006).
 78. S. Ramachandran, O. Deshpande, C. C. Roseman, N. A. Rosenberg, M. W. Feldman, L. L. Cavalli-Sforza, Support from the relationship of genetic and geographic distance in human populations for a serial founder effect originating in Africa. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 15942–15947 (2005).
 79. A. Durvasula, S. Sankararaman, Recovering signals of ghost archaic introgression in African populations. *Sci. Adv.* **6**, eaax5097 (2020).
 80. E. Ravindran, H. Hu, S. A. Yuzwa, L. R. Hernandez-Miranda, N. Kraemer, O. Ninnemann, L. Musante, E. Boltshauser, D. Schindler, A. Hübner, H. C. Reinecker, H. H. Ropers, C. Birchmeier, F. D. Miller, T. F. Wienker, C. Hübner, A. M. Kaindl, Homozygous ARHGEF2 mutation causes intellectual disability and midbrain-hindbrain malformation. *PLOS Genet.* **13**, e1006746 (2017).
 81. C. Hashizume, A. Moyori, A. Kobayashi, N. Yamakoshi, A. Endo, R. W. Wong, Nucleoporin Nup62 maintains centrosome homeostasis. *Cell Cycle* **12**, 3804–3816 (2013).
 82. K. I. Swenson, K. E. Winkler, A. R. Means, A new identity for MLK3 as an NIMA-related, cell cycle-regulated kinase that is localized near centrosomes and influences microtubule organization. *Mol. Biol. Cell* **14**, 156–172 (2003).
 83. T. K. Sears, J. M. Angelastro, The transcription factor ATF5: Role in cellular differentiation, stress responses, and cancer, and cancer. *Oncotarget* **8**, 84595–84609 (2017).
 84. O. D. Gil, G. Zanazzi, A. F. Struyk, J. L. Salzer, Neurotrimin mediates bifunctional effects on neurite outgrowth via homophilic and heterophilic interactions. *J. Neurosci.* **18**, 9312–9325 (1998).
 85. T. Maes, A. Barceló, C. Buesa, Neuron navigator: A human gene family with homology to unc-53, a cell guidance gene from *Caenorhabditis elegans*. *Genomics* **80**, 21–30 (2002).
 86. D. L. C. van den Berg, R. Azzarelli, K. Oishi, B. Martynoga, N. Urbán, D. H. W. Dekkers, J. A. Demmers, F. Guillemot, Nipbl interacts with Zfp609 and the integrator complex to regulate cortical neuron migration. *Neuron* **93**, 348–361 (2017).
 87. B. B. Muhoberac, R. Vidal, Abnormal iron homeostasis and neurodegeneration. *Front. Aging Neurosci.* **5**, 32 (2013).
 88. E. Shokri-Kojori, G. J. Wang, C. E. Wiers, S. B. Demiral, M. Guo, S. W. Kim, E. Lindgren, V. Ramirez, A. Zehra, C. Freeman, G. Miller, P. Manza, T. Srivastava, S. de Santi, D. Tomasi, H. Benveniste, N. D. Volkow, β -Amyloid accumulation in the human brain after one night of sleep deprivation. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 4483–4488 (2018).
 89. S. A. Kaeser, M. C. Herzig, J. Coomaraswamy, E. Kilger, M. L. Selenica, D. T. Winkler, M. Staufenbiel, E. Levy, A. Grubb, M. Jucker, Cystatin C modulates cerebral beta-amyloidosis. *Nat. Genet.* **39**, 1437–1439 (2007).
 90. D. M. Holtzman, J. Herz, G. Bu, Apolipoprotein E and apolipoprotein E receptors: Normal biology and roles in Alzheimer disease. *Cold Spring Harb. Perspect. Med.* **2**, a006312 (2012).
 91. H. A. Pearson, C. Peers, Physiological roles for amyloid beta peptides. *J. Physiol.* **575**, 5–10 (2006).
 92. H. B. Nygaard, C. H. van Dyck, S. M. Strittmatter, Fyn kinase inhibition as a novel therapy for Alzheimer's disease. *Alzheimers. Res. Ther.* **6**, 8 (2014).
 93. I. Letunic, P. Bork, Interactive tree of life (iTOL) v3: An online tool for the display and annotation of phylogenetic and other trees. *Nucleic Acids Res.* **44**, W242–W245 (2016).
- Acknowledgments:** We thank S. Vohr and R. Corbett-Detig for helpful input. **Funding:** This work was funded by the NIH T32 grant HG00834, NSF grant DEB-1754451, and by the Howard Hughes Medical Institute. **Author contributions:** All authors conceived of and helped direct the study. N.K.S. performed the analyses. All authors contributed to writing this paper. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Code is available at <https://github.com/nkschaefer/sarge>. Additional data related to this paper may be requested from the authors.
- Submitted 4 April 2020
Accepted 3 June 2021
Published 16 July 2021
10.1126/sciadv.abc0776
- Citation:** N. K. Schaefer, B. Shapiro, R. E. Green, An ancestral recombination graph of human, Neanderthal, and Denisovan genomes. *Sci. Adv.* **7**, eabc0776 (2021).