2020-04-21

The spatiotemporal spread of human migrations during the European Holocene

Racimo, F

http://hdl.handle.net/10026.1/15421

The spatiotemporal spread of human migrations during the European Holocene

Fernando Racimo[1,*], Jessie Woodbridge[2], Ralph M. Fyfe[2], Martin Sikora[1],
Karl-Göran Sjögren[3], Kristian Kristiansen[3], Marc Vander Linden[4]

[1] Lundbeck GeoGenetics Centre, The Globe Institute, University of Copenhagen, Denmark.

[2] School of Geography, Earth and Environmental Sciences, University of Plymouth, UK.

[3] Department of Historical Studies, University of Gothenburg, Sweden.

[4] Department of Archaeology, University of Cambridge, UK.

* Corresponding author: fracimo@bio.ku.dk

February 26, 2020

Abstract

The European continent was subject to two major migrations of peoples during the Holocene: the northwestward movement of Anatolian farmer populations during the Neolithic and the westward movement of Yamnaya steppe peoples during the Bronze Age. These movements changed the genetic composition of the continent's inhabitants. The Holocene was also characterized by major changes in vegetation composition, which altered the environment occupied by the original hunter-gatherer populations. We aim to test to what extent vegetation change through time is associated with changes in population composition as a consequence of these migrations, or with changes in climate. Using ancient DNA in combination with geostatistical techniques, we produce detailed maps of ancient population movements, which allow us to visualize how these migrations unfolded through time and space. We find that the spread of Neolithic farmer ancestry had a two-pronged wavefront, in agreement with similar findings on the cultural spread of farming from radiocarbon-dated archaeological sites. This movement, however, did not have a strong association with changes in the vegetational landscape. In contrast, the Yamnaya migration speed was at least twice as fast, and coincided with a reduction in the amount of broad-leaf forest and an increase in the amount of pasture and natural grasslands in the continent. We demonstrate the utility of integrating ancient genomes with archaeometric datasets in a spatiotemporal statistical framework, which we foresee will enable future studies of ancient populations movements, and their putative effects on local fauna and flora.

Significance

We present the first study to model the spread of ancestry in ancient genomes through time and space, and the first geostatistical framework for comparing human migrations and land cover changes, while accounting for changes in climate. We show that the two major migrations during the European Holocene had different spatiotemporal structures and expansion rates. In addition, we find that the Yamnaya expansion had a stronger association with vegetational landscape changes than the earlier Neolithic farmer expansion. Our approach paves the way for future work linking paleogenomics with other archaeometric datasets in the study of the past.

Introduction

Up until about 8,500 years before present (BP), Europe was largely populated by groups of hunter-gatherers living at relatively low densities. This scenario changed when a wave of populations from the Middle East entered Europe via Anatolia, as evinced by recent ancient DNA studies (Sikora et al., 2014; Lazaridis, Patterson, et al., 2014; Lazaridis, Nadel, et al., 2016). Studies based on radiocarbon-dated domestic plants, animals and finds from associated contexts suggest that this migration wave spread farming practices into the region, initiating the Neolithic revolution

in Europe (Ammerman and Cavalli-Sforza, 1971; Silva and Steele, 2014; Fort, 2018; Fort, 2015; Pinhasi, Fort, and Ammerman, 2005; Vander Linden and Silva, 2018). A second massive wave of movement occurred later, at the beginning of the Bronze Age, when populations associated with the Yamnaya culture in the Pontic steppe entered the continent from the east (Anthony, 2010; Shishlina, 2008; Kristiansen and Larsson, 2005). These groups may have introduced horse-herding and proto-Indo-European languages as they moved westward, and are associated with the Corded Ware culture in central and northern Europe and, later on, the Bell Beaker phenomenon in northwestern Europe (Haak et al., 2015; Allentoft et al., 2015; Olalde, Brace, et al., 2018; Vandkilde, 2007).

Over the last 10,000 years, the continent also underwent major changes in its land-cover composition, but it is unclear how much the Neolithic and Yamnaya migrations contributed to these changes. Recent pollen-based studies suggest that a dramatic reduction of broad-leaf forests occurred from about 6,000 BP until the present (N. Roberts et al., 2018). This deforestation intensified from around 2,200 BP, resulting in a replacement of these forests by grassland and arable land throughout the continent (Marquer et al., 2017; R. M. Fyfe, Woodbridge, and N. Roberts, 2015). These processes, however, did not occur at the same rate throughout all regions. For example, while considerable decreases in broad-leaf forests occurred in central Europe starting around 4,000 BP, the Atlantic seaboard was predominantly occupied by semi-open vegetation since well before this time, while southern Scandinavia experienced less significant reductions in forest cover, at least until the Middle Ages (R. M. Fyfe, Woodbridge, and N. Roberts, 2015; Anne Birgitte Nielsen et al., 2012; R. M. Fyfe, Twiddle, et al., 2013). Presumably, these phenomena were partly effected by new human land use activities involving forest clearance and the establishment of farming and herding practices, as earlier hunter-gatherer groups likely had limited effects on their surrounding flora and fauna (although see Bishop, Church, and Rowley-Conwy (2015) and Warren et al. (2014)). Changes in climate patterns may have also played a role in vegetation changes. Additionally, changes in vegetation may have also opened up new areas for populations to expand. Until now, however, few efforts have been carried out to explicitly link changes in paleovegetation to particular human population movements, or to distinguish between climatic and human-based factors, assuming these had causal roles in these changes (but see Marquer et al. (2017) and C. N. Roberts et al. (2019)).

In this study, we aim to trace how the major Holocene migrations unfolded across the European continent over time, and to understand how they were associated with changes in the vegetational landscape. We do so by combining ancestry inference on ancient genomes with geostatistical methods, which explicitly account for space and time, and are commonly used to model environmental processes. We use these methods to produce detailed spatiotemporal maps of ancestry movements, and to uncover their relationship with the spread of farming practices and vegetation changes. Additionally, we estimate the front speed of these migrations, and compare our results to reconstructions of cultural dispersal obtained from radiocarbon-dated archaeological sites.

Our modeling approach reveals important factors that may have affected land cover in the past 10,000 years. We find that a decline in broad-leaf forest and an increase in pasture / natural grassland vegetation was concurrent with a decline in hunter-gatherer ancestry, and may have been associated with the fast movement of steppe peoples during the Bronze Age. We also find that natural variations in climate patterns during this period are also associated with these land cover changes. We believe that our approach paves the way for future geostatistical studies integrating paleogenomics with archaeometric datasets, which will yield new insights as information about our past continues to accumulate.

Results

We downloaded publicly available ancient and present-day DNA sequences from human genomic studies (Mathieson, Lazaridis, et al., 2015; Lazaridis, Patterson, et al., 2014; Patterson et al., 2012; Allentoft et al., 2015; Haak et al., 2015; Mathieson, Alpaslan-Roodenberg, et al., 2018; Olalde, Brace, et al., 2018; Lazaridis, Nadel, et al., 2016; Olalde, Schroeder, et al., 2015; Lipson et al., 2017; Fu et al., 2016; Hofmanová et al., 2016; Skoglund, Malmström, Omrak, et al., 2014) (Table S1). We performed unsupervised latent ancestry estimation on these sequences using Ohana (Cheng, Mailund, and R. Nielsen, 2017) with K=4 hidden ancestry clusters. We chose this value of K because, under this scheme, three of the components correspond to the three major ancestral populations that have been previously shown to have resulted - via multiple migration and admixture

events - into the present-day European gene pool: the original Mesolithic hunter-gatherers (HG), the Neolithic farmers who migrated from the Near East (NEOL), and the Yamnaya steppe peoples who entered Europe during the Bronze Age (YAM) (Lazaridis, Patterson, et al., 2014; Allentoft et al., 2015; Haak et al., 2015) (Figure 1). The fourth is an ancestry component that remains largely confined to Northern Africa and the Fertile Crescent throughout most of the Holocene (NAF, Figure S1). We largely focus on the first three components and note that the specification of a larger number of ancestry components could provide further details into more subtle patterns of migration and population expansion, but may also be confounded by bottlenecks and ghost admixture events (Daniel J Lawson, Van Dorp, and Falush, 2018), so we do not pursue finer ancestry estimation here.

As we demonstrate below, the YAM and NEOL ancestries closely parallel the Yamnaya and Neolithic farmer cultural horizons. However, ancestry and culture are distinct concepts that do not always overlap in time and space, so we choose to use the acronym nomenclature when referring to ancestries, and the full name when referring to cultures, unless otherwise specified. Furthermore, there were various, quite differentiated hunter-gatherer populations (Eastern, Western and Scandinavian hunter-gatherers) who migrated into Western Eurasia before the Holocene (Skoglund, Malmström, Omrak, et al., 2014; Skoglund, Malmström, Raghavan, et al., 2012; Sánchez-Quinto et al., 2012; Lazaridis, Nadel, et al., 2016; Fu et al., 2016). The HG ancestry roughly corresponds to the ancestry referred to as "Western hunter-gatherer" in these publications. We note that under our K=4 admixture scheme, Scandinavian hunter-gatherers are modelled as containing high amounts ($\approx 80\%$) of HG ancestry, while the rest of their ancestry is modelled as YAM (which works here as a stand-in for Eastern hunter-gatherer ancestry; for further details about hunter-gatherer ancestry movements see Skoglund, Malmström, Omrak, et al. (2014) and Mathieson, Lazaridis, et al. (2015)). The data for each of these populations is scarcer than for Bronze Age and Neolithic individuals, and, in this work, we chose only to focus on later Holocene ancestry movements.

We first sought to compare the spread of dispersal of NEOL and YAM ancestries over time, using the calibrated C14 dates of each genome. Following Pinhasi, Fort, and Ammerman (2005) and Silva and Steele (2014), we regressed time against distance from the presumed origin of the spread of each of these ancestries, using the ranged major axis (RMA) method. This allows us to obtain an estimate for the migration front speed. We first used samples that had at least 50% of the corresponding ancestry we were studying, as a cutoff value is needed to be able to declare that a particular ancestry was high enough to consider the ancestry had "arrived" at that point in time and space. Using this cutoff, we find that the speed of the YAM migration (4.2 km/year; CI: 3.5-5.2) was at least twice as fast as the NEOL migration (1.8 km/year; CI: 1.6-2.1), assuming an origin of the Yamnaya migration at the center of its historical range (Figure 2.A). A higher ancestry cutoff of 75% to establish "first arrival" yields the same estimate for the NEOL migration (1.8 km/year; CI: 1.6-2.2), but an even faster estimate for the YAM migration (9.3 km/year; CI: 6-20). YAM speed estimates are generally higher than NEOL speed estimates, for almost any choice of minimum ancestry cutoffs, unless these cutoffs are chosen to be very small ($\leq 20\%$) (Table S2).

Given that the original Yamnaya ancestry range was quite large (Anthony, 2010), we also aimed to see how our estimates varied as we altered the point of origin within the Yamnaya range. We obtained estimates of speed assuming a location of origin at the northern-most, eastern-most, western-most and southern-most parts of the Yamnaya range, which yielded similar estimates of speed (Table S3). The magnitude of the negative correlation coefficients between time and distance from origin can also be used to estimate the point of origin (Pinhasi, Fort, and Ammerman, 2005; Hunt et al., 2018), assuming a range expansion for the YAM and NEOL ancestries. Indeed, when we altered the point of origin, we find that the most negative correlation coefficients correspond to Anatolia and the Middle East for the NEOL ancestry and to the Caspian steppe for the YAM ancestry (Figure 2.B).

To be able to compare ancestry through time and space with other variables, we aimed to project our ancestry values to particular times and locations for which we do not necessarily have sampled genomes (Figure 3). To do so, we computed a spatiotemporal variogram and fitted it to a metric covariance function (E. J. Pebesma, 2004; E. Pebesma and Heuvelink, 2016) (Figures S2,S3,S4,S5). We then performed spatiotemporal kriging of the inferred latent ancestry values on a dense grid of spatial points across Europe, over a 10,800-year span, with intervals of 600 years (Figures 4, 5, S6, S7, Movies S1-S3). In practice, however, given the sparseness of the data in the distant past, we restrict our discussion to patterns seen more recently than 8,000 years BP.

We downloaded land cover class (LCC) maps (R. M. Fyfe, Woodbridge, and N. Roberts, 2015)

and paleoclimatic variable maps (Brown et al., 2018) spanning the Holocene, and projected them on the same spatiotemporal grid that we used for our kriged ancestry values (see Methods). The paleovegetation types included needle-leaf forest (LCC1), broad-leaf forest (LCC2), heath / scrubland (LCC5), pasture / natural grassland (LCC6) and arable / disturbed land (LCC7). We computed correlations between each of the spatiotemporally projected ancestry proportions and vegetation types, and between the climate variables and vegetation types. This was done in three different ways. Firstly, we simply obtained the correlation of the raw values of any two variables (Figure S8). Secondly, we obtained the correlation of the differences in these variables between a particular time slice and the immediately previous time slice (Figure S9). Thirdly, we obtained the correlations of the variable anomalies, defined as the value of each ancient variable after subtracting the present-day value from the same location (Figure S10). We note, however, that this approach does not account for autocorrelation in time and space that may exist for all the compared variables, not only because of real autocorrelation in the processes under study, but also as a result of enforced autocorrelation from the smoothing techniques that generated the maps.

The raw correlations reflect spatially static patterns of co-occurrence (Figure S8, Table S4). For example, YAM ancestry is largely prevalent in northeast Europe throughout much of the Holocene, and this coincides with periods of abundant needle-leaf forests, which is why there is strong positive correlation between these variables. Conversely, NEOL ancestry is largely prevalent in southern Europe during this period, which is why there is a negative correlation with needle-leaf forest. In contrast, the correlations in differences and in anomalies reflect spatially dynamic patterns of co-occurrence (Figures S9 and S10, Table S4). Here, temporal increases in one variable that coincide with temporal increases in a second variable at the same location will result in positive correlation. The same will result if there are co-occurring decreases. If, however, a variable decreases while another increases at the same location, this will result in negative correlation. For example, we see that YAM ancestry anomalies are positively correlated with pasture / natural grassland anomalies, but negatively correlated with broad-leaf forest anomalies. We also see that the correlations between ancestry differences and vegetation differences increase when looking at vegetation differences one or two time slices into the future (600 or 1,200 years later, respectively), perhaps suggesting that migrations could have had a role in these vegetational changes (Figure S11).

On a continental level, decreases in broad-leaf forest and increases in pasture / grassland occurred most notably after the arrival of YAM ancestry, not after the arrival of NEOL ancestry. However, vegetation changes behaved in different ways in different parts of the continent (Figures 6, S12). In central France, increases in YAM ancestry coincided with decreases in broad leaf forest cover. In contrast, in southeastern and southwestern Europe, forest cover remained stable (at low levels), even as YAM ancestry was increasing. If humans were responsible for this, it could perhaps due to the development of tree cropping within the agropastoral system in the Mediterranean (C. N. Roberts et al., 2019). Considerable increases in arable land cover occur fairly late in the Holocene throughout the continent, and much later than the incursion of NEOL ancestry during the Neolithic (Figure S12). Interestingly, we observe a decrease in NEOL ancestry that continues even after the incursion of YAM ancestry into Europe, though our resolution for quantifying changes in this ancestry in the last 2,000 years is limited by the scarcity of ancient genomes from the very recent past.

While correlations between vegetation and ancestry are interesting, they do not account for the fact that we are projecting the data to lie in a particular set of spatiotemporal grid points, which have complex auto-correlations in time and space, potentially affecting the correlations we observe between variables. To address this, we used a spatiotemporally explicit hierarchical Bayesian model to better understand the relationships between changes in climate, ancestry and paleovegetation, while accounting for these autocorrelations (Figure 3). We used two models, implemented in the R package spTimer (Bakar, Sahu, et al., 2015). One is a Gaussian process (GP) model that incorporates a spatiotemporal nugget that is independent of time and has a distribution that depends on a spatial correlation matrix (Figure 7, Table 1). The other is an extension of this method that incorporates a temporal autoregressive component (AR; Figure 7, Table 2). We set the kriged ancestry and climate variables to be the explanatory variables, while each of the paleovegetation variables was set as a response variable. We fitted five separate models for each paleovegetation variable. Table S5 lists the goodness of fit score for both models, along with the predictive model choice criteria (PMCC), which accounts for differences in model complexity (Bakar, Sahu, et al., 2015; Gelfand and Ghosh, 1998). In comparison to the GP model, the AR

model results in a more sparse set of posterior coefficients whose credible intervals do not overlap with 0 (Figure 7, Table 2). An evaluation of the root mean squared error of the predictions (see Methods), suggests that the GP model with a uniform prior distribution for the decay parameter of the spatial correlation function generally has the lowest validation error (Figure S13).

We also compared the predictive accuracy of models incorporating ancestry only, climate only, both sets of variables or none of them. In general, adding both climate and ancestry result in a better PMCC score than adding either in isolation or adding none of them (Table 3). However, in all but one of the paleovegetation variables, there is no observable difference in the root mean squared error of the fitted model when adding climate only, ancestry only, or both climate and ancestry. The exception to this pattern is pasture / grassland (LCC6), which has both the lowest error and the lowest PMCC when including both climate and ancestry under the GP model (Figure S13, Table 3).

Because we are using kriged ancestry as an explanatory variable and our ancient genomes are unevenly sampled across space and time, we were mindful that the Bayesian credible intervals (BCI) obtained from the hierarchical model would not accurately reflect uncertainty in particular regions of space-time. For that reason, we performed nonparametric bootstrapping of the parameter estimates. We randomly sampled ancient and present-day genomes with replacement from among the list of all genomes until we had as many genomes as were in the original dataset, then obtained their ancestry assignments, kriged them on the spatiotemporal grid and inputted them into the Bayesian hierarchical model. We did this 100 times to obtain 100 pseudo-samples, which allowed us to obtain 95% bootstrap-based confidence intervals (BBCI) around the mean posterior estimates (Figure S14, Table S6). Below, we discuss results that are supported by one or both models, and that are also supported by the bootstrapping approach.

Regardless of the model used, we find that HG ancestry is positively associated with broad-leaf forest anomalies, but negatively associated with arable land anomalies. YAM ancestry, in turn, is positively associated with pasture / natural grassland (Figures 7, S14). In the fitted AR model and the bootstrapping approach, we also see a negative association of HG ancestry to pasture / grassland and scrubland. In the fitted GP model and the bootstrapping approach, we observe a negative association of YAM ancestry with forest vegetation, which is strongest for broad-leaf forest. We see weak or non-existent associations of NEOL ancestry with any vegetation type. We cannot discard the possibility that we may lack the ability to detect some of associations between ancestry movements and vegetation changes at our current scale of resolution.

We additionally observe associations of different climate variables with the different vegetation anomalies, which become sparser in the AR model (Figure 7). For example, in both the AR and GP models, increases in temperature are related to increases in non-forest vegetation types (scrubland, pasture and arable land). In addition, temperature seasonality may be interpreted as negatively associated with the proportion of arable land, while precipitation during the driest quarter may be interpreted as associated positively with heath / scrubland, under the fitted models.

Finally, we built "first arrival" maps (Pinhasi, Fort, and Ammerman, 2005; Fort, 2015; Vander Linden and Silva, 2018) for both NEOL and YAM, given that changes in these ancestries can be broadly interpreted as incursions of foreign populations into the European continent during the Neolithic and Bronze Age (Allentoft et al., 2015; Haak et al., 2015). The first arrival map of NEOL ancestry shows that this ancestry spread closely parallels the inferred cultural spread of farming, which has been inferred from archaeological sites (Figure 8) (Fort, 2018; Fort, 2015; Pinhasi, Fort, and Ammerman, 2005; Silva and Steele, 2014). When performing the same type of reconstruction for the YAM ancestry, we observe that this spread occurs first via north and central Europe, and only much later begins to spread into southern Europe (Figure 8), reflecting reconstructions from archaeological records for the spread of the Yamnaya, Corded Ware and Bell Beaker phenomena (Anthony, 2010; Shishlina, 2008; Kristiansen and Larsson, 2005).

Discussion

An explicitly geostatistical approach allows us to visualize how movements of ancestry occurred during the Holocene in Europe. The NEOL ancestry expansion followed a two-pronged shape, paralleling the expansion of farming practices estimated from radiocarbon-dated archaeological sites. We observe two wave fronts, one northward across central Europe, and one westward along the Mediterranean coast (Figure 8). In the cultural map, these two wave fronts correspond to the Linear Pottery culture (LBK) (Bickle and Whittle, 2013) and the Impressa / Cardial Pottery

culture (Barnett, 2000; Binder et al., 2018; Manen et al., 2019). Given their close parallels in the ancestry map, this supports the view that these two cultural expansions were probably driven by migrations of people (Allentoft et al., 2015; Haak et al., 2015).

We estimate that the expansion of YAM ancestry occurred faster than the expansion of NEOL ancestry. The reasons for this could be numerous, including the use of horses for long-distance travel (Anthony, 2010). YAM ancestry predominates in individuals associated with the Yamnaya and Corded Ware cultures, and is presumed to have moved into Europe from the Eurasian Steppes (Kristiansen and Larsson, 2005). Another possibility could be the opening of the landscape previous to the arrival of the Yamnaya people, perhaps due to Neolithic agricultural, grazing and mining practices (Schauer et al., 2019), which may have facilitated later movements of people. We do not observe a strong decrease in forest vegetation in Northern and Central Europe until the Bronze Age, however. On the other hand, there is limited evidence that Corded Ware people were horse herders, and evidence from settlements in central Europe suggests they may have practiced mixed agriculture (Müller et al., 2009; Seregély and Müller, 2008; Jacomet, 2008).

We can now begin to understand how these movements of people may have been associated with the European vegetational landscape, while accounting for autocorrelation in time and space (Figure 7). We generally fit HG ancestry as positively associated with broad leaf vegetation, while YAM ancestry as negatively associated with broad leaf forest vegetation, and positively associated with grassland and arable land. We also find associations between climate and changes in land cover type. For example, increases in temperatures are related to increases in scrubland, pasture / grassland and arable land.

We do not find that NEOL ancestry had a strong association to changes in vegetation. One possible explanation is that this association was too minor or localized for us to clearly detect an effect in our model. Earlier studies have shown that Neolithic communities did in fact alter their local environments (Mercuri et al 2019 Holocene) and had a local effect on vegetation to a certain extent, at least in northwestern Europe (Woodbridge, R. M. Fyfe, et al., 2014; Schauer et al., 2019; Lechterbeck et al., 2014; Woodbridge, R. M. Fyfe, et al., 2014). In particular areas, such as northern and northwestern Europe, there is a very minor decline in broad-leaf forest that coincided with the increase in NEOL ancestry, but this is not observed at a continental level (Figure 6). A much more pronounced reduction in broad-leaf forest occurs later on throughout western and northwestern Europe, and coincides with the increase in YAM ancestry. It is important to note that cultivated tree types (olive, chestnut and walnut) - which are pervasive in the Mediterranean - also fall into the category of broad-leaf forest. Thus, our capacity to infer changes in forest types in regions with this type of cultivar (e.g. the Mediterranean, Woodbridge, N. Roberts, and R. Fyfe (2018)) is limited.

The decrease in broad-leaf forest (starting 6,000 years BP) was followed by a minor increase in grassland and disturbed land in some parts of the continent. These vegetation types are naturally present in the Mediterranean and the Black Sea region throughout the earlier part of the Holocene and remain fairly stable until the present (R. M. Fyfe, Woodbridge, and N. Roberts, 2015). In contrast, in western Europe, these vegetation types only reach intermediate levels during the Bronze Age - as YAM ancestry begins to increase - and they continue to increase after the end of this period. Furthermore, increases in YAM ancestry in southern and eastern Europe do not coincide with increases in grassland and disturbed land.

Pasture / natural grassland is the only land cover type that considerably improves in fit as a result of adding ancestry and climate variables into our model. This might be because we currently lack the spatiotemporal resolution to provide much predictive power with the addition of climate or ancestry variables, or because these variables may not be strongly predictive of the other land cover classes. Other unaccounted factors may have had a stronger effect on the landscape. An obvious candidate is the dramatic increase in population density that occurred over the last 3,000 years (Bevan et al., 2017; C. N. Roberts et al., 2019), which likely led to strong changes in land use practices, consequently disturbing vegetation throughout the continent. Earlier population rises and collapses during the Neolithic and Bronze Age could have also influenced the vegetated landscape in significant ways, although on a smaller scale (Shennan et al., 2013; Lechterbeck et al., 2014; Woodbridge, R. M. Fyfe, et al., 2014). Thus, a future study could aim to incorporate estimates of human population density or other measures of human activity into explanatory models for changes in vegetation, together with population movement. A recent approach using human land use estimates, for example, showed that, on a continental scale, climate changes were the main driver of changes in vegetation when the Holocene is considered in its entirety, but the

influence of human land use markedly increased from 4000 BP onwards (Marquer et al., 2017).

There are a number of caveats and assumptions in our modelling procedure that are important to keep in mind. Firstly, we are assuming that changes in ancient ancestries can be used as a proxy for long-distance movement of people. This may be the case for particular periods of time - especially when peoples of highly divergent ancestries first met each other - but this assumption loses validity as we move closer to the present, and the ancestry components tend to become more homogenized due to later migrations within the area of study (Margaryan et al., 2019). Tracing relatively high YAM ancestry in the present day is approximately equivalent to tracing people with high Northern European ancestry, who cannot be equated with ancient "steppe peoples".

Second, our projection of inferred ancestry components to a spatiotemporal grid does not model processes that cause variation in ancestry proportions within a specific region of space-time. Indeed, local departures from the inferred kriged ancestry proportion in a given region is treated as noise in the kriging model. This fails to account for the fact that some regions (e.g. cosmopolitan centers of trade) may have harbored much more variation in proportions than other regions, in which the proportions may have been more homogeneous. These models also do not account for differences in population density which could mean that certain migrations or population expansions may have involved much larger numbers of people than others, even if their consequent changes in ancestry proportions may be inferred to be relatively similar in size.

Third, we are relying on existing ancient DNA data, which has its own idiosyncrasies, due to environmental and historical biases in sampling. For example, North Africa and eastern Scandinavia are sparsely sampled in our dataset, so our ancestry estimates for those regions are much poorer than for the rest of the European continent. We attempted to account for these types of biases via a bootstrapping approach and various estimates of error due to temporal and spatial patchiness, to assess how robust these were to accidents of sampling (see Methods).

Additionally, we are relying on a particular choice of the number of ancestry clusters or components (K) under a latent mixed-membership model. We chose this model and parameter setting to be able to discretize patterns of ancestry into three major population clusters (HG, NEOL and YAM), which have been documented via other, more involved, population genetic analyses (Haak et al., 2015; Allentoft et al., 2015; Lazaridis, Patterson, et al., 2014). We also chose a low number of clusters to have enough data points across extended periods of time, in order to accurately estimate the space-time decay in covariance between ancestries (e.g. Figures S2 and S3). These clusters are, however, an approximation of a very complex genealogical process. Indeed, the clusters cannot be seen as discrete, originally isolated populations, as there may be both isolation-by-distance and hierarchical population structure within each of these groups (Frantz et al., 2009; Janes et al., 2017; Battey, Ralph, and Kern, 2019), and it is unclear how incorporating these phenomena would affect the kriging or the RMA speed estimates. In our downstream analyses, we are also assuming these clusters included groups of people with temporally and spatially self-consistent land change practices. The clusters themselves are also the result of more complex admixture and migration events that occurred before the Holocene (Fu et al., 2016). Additionally, there is likely some differentiation in population structure over time and space, even when looking at the same admixture components (e.g. the NEOL component of a present-day Sardinian is differentiated relative to the NEOL component of a Bronze Age central European). These subtle patterns are hard to pick up by simple latent mixed-membership models (Daniel J Lawson, Van Dorp, and Falush, 2018), although there has recently been some progress in this regard (Joseph and Pe'er, 2018; Bradburd, Coop, and Ralph, 2018). Other types of population genetic frameworks are able to better detect some of these more subtle signals by, for example, modelling patterns of haplotype sharing (Hellenthal et al., 2014; Daniel John Lawson et al., 2012), the full site-frequency spectrum (Excoffier et al., 2013; Kamm et al., 2019) or an approximation to the full ancestral recombination graph (Kelleher et al., 2018; Speidel et al., 2019). Nevertheless, these also have their own limitations and assumptions. For all these reasons, we advice the reader to consider that the ancestry components used in this study are approximations of the true historical admixture process.

Furthermore, in our model relating different ancestries to different land cover types, we are making a unidirectional causality assumption, as we have a priori chosen ancestries and climate as the explanatory variables and land cover as the response variables. In other words, we are testing how migrations and climate may have affected vegetation. It is also possible that people moved to new environments as a consequence of vegetation or climate changes, or of other environmental factors that we are not studying here.

Finally, it is important to remember that the ancestry proportions exist in a simplex, so in-

creases in one ancestry will proportionally lead to decreases in other ancestries. For this reason, a negative contribution to vegetation from one ancestry (e.g. YAM and broad-leaf forest) coupled with a positive contribution from another ancestry (e.g. HG and broad-leaf forest) may be two manifestations of the same process - change in land cover as a result of change in ancestry - rather than two independent processes.

An improvement to our current approach could involve developing a hierarchical dynamical model for explicitly modelling spatiotemporal movements on the genetic data directly, without relying on ancestry assignments estimated from a non-spatiotemporally-aware model (Cressie and Wikle, 2015). This could also help to better deal with boundary constraints that are not accounted for by the kriging methodology. For example, we currently have to correct kriged estimates that are lower than 0 or higher than 1. A generative model of spatiotemporal ancestry would not allow for these types of parameters in the first place, for example, by placing Bayesian priors on ancestry with 0 probability outside of the 0-1 range. This could also be solved by extending compositional interpolation techniques to a spatiotemporal setting (Walvoort and Gruijter, 2001).

Keeping these considerations in mind, the approach developed here is the first attempt at combining in an explicit, quantitative framework, various categories of evidence, which have otherwise either not been considered together (e.g. ancestry and land cover type), or have only been compared in a qualitative way. There is a lot of potential for new geostatistical approaches that could be designed to combine various types of datasets in an integrative approach for the study of the past, including at more local scales than considered here. This could encompass, for example, the combination of strontium and oxygen isotope analyses together with radiocarbon data and contextual archaeological information (Sjögren, Price, and Kristiansen, 2016; Mittnik et al., 2019), the joint analysis of genetic and linguistic changes over time (Kristiansen, Allentoft, et al., 2017) or the study of the interactions between population density and vegetation (Müller, 2015; Kolář et al., 2018).

In summary, although our methodology relies upon several assumptions and could benefit from a myriad of extensions, it provides a robust way to account jointly for space and time in the study of genetic and environmental variables. Our results demonstrate that the two major human migrations recorded in Holocene Europe differ markedly in their expansion rates and, possibly, had distinctive implications for the environment in which they unfolded. By explicitly modelling space and time, researchers can move beyond the mere identification of human migrations: we can begin to understand structural differences between and within human dispersal events, and study local phenomena that may have unravelled in different ways across an area of study. Otherwise, we might run the risk of overlooking important historical processes, by taking an overly global perspective. We should not ignore the forest for the trees, but sometimes, the trees themselves might be hidden by the forest.

Methods

All R code used to perform the analyses in this manuscript has been deposited in: https://github.com/FerRacimo/STAdmix

Kriged ancestry maps over time and space

For our ancestry analyses, we used a combined dataset of 842 ancient and 955 present-day genomic sequences. The present-day sequences were obtained using the Human Origins SNP array, while the ancient sequences were either obtained via this array or via whole-genome sequencing, followed by filtering for SNPs that are in the array (Mathieson, Lazaridis, et al., 2015; Lazaridis, Patterson, et al., 2014; Patterson et al., 2012; Allentoft et al., 2015). We restricted our analyses to modern human genomes obtained from human remains located within an area encompassing most of the European continent: north of 30°N, south of 75°N, east of 15°W and west of 45°E. We inferred latent ancestry components on these genomes using Ohana (Cheng, Mailund, and R. Nielsen, 2017).

We performed ordinary global spatiotemporal kriging using the R libraries gstat (E. J. Pebesma, 2004; E. Pebesma and Heuvelink, 2016) and spacetime (E. Pebesma et al., 2012), to obtain unbiased linear predictions of ancestry for unsampled locations and times. Suppose we have a set of noisy observations of a variable distributed unevenly across space and time. In our case, this will be the inferred proportion of a particular ancestry in each of our ancient genomes. Let s_i be a vector representing the ith site (out of n) in our grid, which is composed of two values: its longitude and

latitude. Following the notation by Cressie and Wikle (2015), suppose we have T_i different temporal samples of a measured variable at site $\boldsymbol{s_i}$. A temporal sample obtained at the jth time (t_{ij}) from this site will be denoted as $Z(\boldsymbol{s_i}, t_{ij})$. Suppose these data are equal to the true spatiotemporal process plus some measurement error ϵ:

$$Z(\boldsymbol{s_i}, t_{ij}) = Y(\boldsymbol{s_i}, t_{ij}) + \epsilon(\boldsymbol{s_i}; t_{ij}) \tag{1}$$

Let $\boldsymbol{Z}^{(i)}$ be the vector containing all values that were measured at different time points in location s_i. Also, let $\boldsymbol{Z} = (Z^{(1)\prime}, ..., Z^{(m)\prime})'$ where m is the number of locations sampled. We can obtain a linear predictor, $Y^*(s_0; t_0)$ for a particular unsampled data point at time t_0 and location s_0:

$$Y^*(s_0; t_0) = \boldsymbol{l'Z} + c \tag{2}$$

where $\boldsymbol{l}$ and c are parameters than can be optimized. In particular, for the case that the true process $Y(;)$ has a constant unknown mean μ, one can show that the linear unbiased predictor that minimizes the mean squared prediction error - also called the ordinary kriging predictor - is equal to:

$$\hat{Y}(\boldsymbol{s_0}, t_0) = \boldsymbol{\lambda'Z} \tag{3}$$

where $\lambda = \{\boldsymbol{c_0} + \boldsymbol{1}(1 - \boldsymbol{1'}C_Z^{-1}\boldsymbol{c_0})/(\boldsymbol{1'}C_Z^{-1}\boldsymbol{1})\}C_Z^{-1}$, $\boldsymbol{c_0} = var(\boldsymbol{Z})$ and $\boldsymbol{C_Z} = cov(Y(\boldsymbol{s_0}, t_0), \boldsymbol{Z})$. The latter can be obtained by fitting a spatiotemporal covariance function for the true process $\boldsymbol{Y}$ to the empirical spatiotemporal variogram of the observed measurements $\boldsymbol{Z}$ (Figures S3, S2, S4, S5). In our case, the variogram was computed over a range of 3,000 years, with 60-year windows, and we used the "metric" variogram model to fit it (Gräler et al. 2015). For a more extensive explanation of spatiotemporal kriging, we refer the reader to Cressie and Wikle (2015).

As our predicted grid, we used a set of spatial points distributed evenly across Europe. We used two types of spatial grids: one containing a dense set of 5,000 points and a sparser set, containing 200 points. We call this our "spatial grid". The dense version of the spatial grid was used for plotting spatiotemporal maps (e.g. Figure 5), while the sparse set was used to fit the Bayesian spatiotemporal model (e.g. Figure 7), for ease of computation. We observed that the ancestry-vegetation and climate-vegetation correlations computed under both schemes were almost identical, suggesting that the use of the sparser grid should not affect inference under the Bayesian model. Potential biases arising from particular grid points possessing few nearby ancient genomes are accounted for in the bootstrapping method described below. Unless otherwise stated, our "temporal" grid had a 10,800-year span, with intervals of 600 years until the present, for a total of 19 time slices. Thus, if our spatial grid had a spatial points, our "spatiotemporal grid" had $19a$ spatiotemporal points. We bounded the kriged ancestry values between 0 and 1, and so kriged values that were negative were set to 0, and those that were larger than 1 were set to 1.

In all analyses below, we did not include a kriged ancestry component that is largely restricted to north Africa and the Fertile Crescent (NAF). The reasons for this are two-fold: 1) this ancestry remains largely spatially static throughout the Holocene, at least with respect to the box we defined to bound our analyses; 2) given that all latent ancestries must add up to 1 in each individual genome, this ancestry is equal to 1 minus the sum of the other three ancestries, and is therefore not linearly independent from them. This component is absent from Europe until the end of our temporal transect, where it surfaces in parts of central Europe, because of the presence of Ashkenazi Jewish genomes in our present-day dataset.

Assessment of quality of kriged maps

To assess the robustness of our kriged maps, we bootstraped our data by sampling with replacement from the set of all genomes 100 times, and recomputed the spatiotemporal kriging each time. This way, we obtained 95% bootstrap-based confidence intervals for each predicted ancestry at all spatiotemporal grid points (Figures S15, S16, S17, S18).

To assess the effects of spatial patchiness in our data, we divided our map into 16 4x4 square sectors. We then computed, for each sector, the mean absolute error (MAE) of the kriged ancestry of the nearest spatiotemporal grid point of each ancient genome inside that sector, relative to the true (Ohana-inferred) ancestry of the genome. In the left panel, the kriged ancestry was obtained by kriging the complete dataset. Here, we observe that our kriging predictions are very

accurate (MAE < 20% across all patches), regardless of the part of the map that we choose to focus on. In the right panel, the kriged ancestry was obtained by kriging a version of the dataset in which all genomes within that sector had been previously removed (Figure S19). Here, our MAE is considerably larger, especially for the YAM and NAF ancestries in northern Europe and Anatolia, suggesting that local genomes are especially important to include in order to derive accurate predictions in these regions.

To assess the effects of temporal patchiness in our data, we also divided our 10,800-year timeline into 10 periods of equal duration (1,080 years). Analogously to the previous analysis, we selected each of the periods in turn, and computed, for each period, the mean absolute error of the kriged ancestry of the nearest spatiotemporal grid point of each genome within that period, relative to the true ancestry of each ancient genome. In the left panel, the kriged ancestry was obtained by kriging using the entire dataset. As in the spatial patch analysis, our predictions are very accurate (MAE < 30% across all slices). In the right panel, the kriged ancestry was obtained by kriging a version of the dataset in which all genomes within that period had been previously removed (Figure S20). In this case, the predictions are less accurate, with particularly inaccurate predictions for NEOL and HG ancestries in the oldest time slices and for all ancestries in the most recent time slice, suggesting there were ancestry changes during these periods that are poorly predicted by using ancestries from adjacent periods.

Paleovegetation maps

We downloaded inferred Holocene paleovegetation spatiotemporal maps from R. M. Fyfe, Woodbridge, and N. Roberts (2015). This paleovegetation reconstruction was built from 982 pollen records across Europe, using the pseudobiomization method (PBM) (R. Fyfe, N. Roberts, and Woodbridge, 2010). It has a 10,800-year span, with intervals of 200 years until the present. To ease computation, we sampled every two time windows, resulting in intervals of 600 years until the present, and for each paleovegetation time slice, we rasterized the map to have 6,540 points (down from 35,856). Then, for each time slice, we inferred the value of each point in our spatial grid by taking the median of the 5 nearest points in the rasterized map.

Paleoclimate maps

We obtained a set of simulation-based Holocene paleoclimate reconstructions for Europe from PaleoClim (Brown et al., 2018), which includes surface temperature and precipitation estimates for the Early (11.7-8.326 kya), Middle (8.326-4.2 kya) and Late Holocene (4.2-0.3 kya), using snapshot-style climate model simulations. These simulations were accessed through PaleoView (Fordham et al., 2017) and come from the TRaCE21ka experiment (Liu, B. Otto-Bliesner, et al., 2009; Liu, Lu, et al., 2014), which used the Community Climate System Model v3 (CCSM3) (Bette L Otto-Bliesner et al., 2006; Collins et al., 2006; Yeager et al., 2006): a general circulation model involving atmosphere, ocean, sea ice and land. The PaleoClim authors refined the simulations from this model, incorporating small-scale topographic nuances of regional climatologies, thus creating high-resolution paleoclimate maps. We projected the three Holocene maps - together with the present-day WorldClim map (Fick and Hijmans, 2017) - onto the previously delineated temporal grid, for each of the 19 climate variables that are present in the PaleoClim database. At each time slice, for each point in the spatial grid, we inferred the value of each climate variable, by taking a weighted average of the values of the two closest bounding paleoclimate time points (past and future) at that spatial point, weighted by their respective temporal distance to our time slice. These allowed us to obtain a spatiotemporal grid of the climate variables at the same locations and times for which we had kriged ancestry and paleovegetation data. In the Bayesian hierarchical model, we excluded one of these variables (temperature annual range) because it is a linear combination of two of the other climate variables.

Computation of correlations

We computed Pearson correlations between the kriged ancestry, climate and vegetation variables in three ways. First, we simply took the vector containing the values of one variable across all points in our spatiotemporal grid and computed its correlation with the values of another variable at all the same spatiotemporal points. We call these the "raw correlations" (Figure S8). Second, starting from the second oldest time slice, we took each of the values of a particular variable of

a time slice and subtracted from them the values of the same variable at the same location but from the immediately previous time slice. We did this for all variables and then computed their pairwise correlations, which we call the "correlations in differences" (Figure S9). Finally, we took each of the values of a particular variable of a time slice and subtracted from them the values of the same variable at the same location but from the last (present-day) time slice. We then computed pairwise correlation between the resulting values for each of the variables, excluding the last time slice from the analysis (as it would just contain zeroes). We call these the "correlations in anomalies", in the sense that the resulting values represent anomalies of a variable with respect to its present-day value at a given location (Figure S10).

We also computed the correlation between the difference in ancestry in a time window and the difference in vegetation one (or two) time window(s) later (Figure S11). In other words, for each time slice i and spatial grid point j, let A_{ij} be the difference in ancestry between time $i+1$ and time i at spatial point j, let B_{ij} be the difference in vegetation between time $i+2$ and time $i+1$ at spatial point j, and let C_{ij} be the difference in vegetation between time $i+3$ and time $i+2$. We computed the correlation between A_{ij} and B_{ij}, and also between A_{ij} and C_{ij} across all spatial grid points j and all time slices i, for each ancestry-vegetation pair.

Spatiotemporal Bayesian modeling of vegetation anomalies

We used two hierarchical spatiotemporal Bayesian models implemented in the R library spTimer (Bakar, Sahu, et al., 2015), in order to jointly model climate and kriged ancestry anomalies as explanatory variables for vegetation type anomalies. To simplify notation, we will now index time with the variable t and assume all sites have observations at the same time slices, i.e. $T_i = T$ for all sites s_i. We will suppose we have n sites, and so nT is the total number of spatiotemporal observations. The first of the spatiotemporal models treats the response variable $\boldsymbol{Z}(t)$ - in our case, containing all vegetation type values in our spatial map at time t - as a noisy observation of a Gaussian process $\boldsymbol{O}(t)$:

$$\boldsymbol{Z}(t) = \boldsymbol{O}(t) + \epsilon_t \tag{4}$$

$$\boldsymbol{O}(t) = \boldsymbol{X}_t\boldsymbol{\beta} + \eta_t \tag{5}$$

Here, $\boldsymbol{\beta}$ is a $p \times 1$ vector of coefficients, $\boldsymbol{X}_t$ is a $n \times p$ matrix of covariates at time t, ϵ_t is an error vector that only depends on an unknown pure error variance σ_ϵ:

$$\epsilon_t = (\epsilon(\boldsymbol{s_1}, t), ..., \epsilon(\boldsymbol{s_n}, t))' \sim N(\boldsymbol{0}, \sigma_\epsilon \boldsymbol{I}_n) \tag{6}$$

while η_t is a spatiotemporal nugget vector that is independent of ϵ_t and whose distribution depends on a site invariant spatial variance σ_η and the spatial correlation matrix S_η:

$$\eta_t = (\eta(s_1, t), ..., \eta(\boldsymbol{s_n}, t))' \sim N(\boldsymbol{0}, \sigma_\eta S_\eta) \tag{7}$$

The correlation matrix S_η is obtained from the general Matérn correlation function (Matérn, 1986), whose shape depends on two unknown parameters - λ and ν. These control the rate of decay of the correlation as the distance between sites increases and the smoothness of the random field, respectively (Bakar, Sahu, et al., 2015).

The second model is a temporal auto-regressive model that works by incorporating a term in equation 7 that depends on the previous instance of the $\boldsymbol{O}()$ process and a temporal correlation parameter ρ:

$$\boldsymbol{O}(t) = \rho\boldsymbol{O}(t-1) + \boldsymbol{X}_t\boldsymbol{\beta} + \eta_t \tag{8}$$

spTimer can fit these models via Gibbs sampling and infer the posterior distribution of the unknown parameters $\boldsymbol{\beta}$, ϵ_t, η_t, ν, ϕ and ρ. Unless otherwise stated, we used spTimer's default prior distributions for these parameters, which are detailed in Bakar, Sahu, et al. (2015). Before inputting all explanatory and response variables into either model, we first centered and scaled them to have mean 0 and variance 1. We tried three different types of prior distributions for the spatial decay parameter of the Matérn correlation function (Figures S13) - a fixed value, a Uniform distribution and a Gamma distribution, each with default hyperparameters - and compared their performance using the root mean squared error of the predictions (see below).

Assessment of error of hierarchical model

We randomly removed 20% of the grid points in the map and fitted a spatiotemporal model to the remaining portion of the data. We computed the root mean squared error by comparing the predicted values across all temporal slices with the previously removed observed values. We then selected the spatiotemporal model (AR vs GP) and the prior distribution for the spatial decay parameter of the Matérn correlation function (fixed value vs. Uniform vs. Gamma) based on visual comparison of the root mean squared error plots for each of these model choices (Figure S13) (Bakar, Sahu, et al., 2015).

Predictive model choice criterion

We used the predictive model choice criterion to compare different hierarchical Bayesian models. The PMCC was developed by Gelfand and Ghosh (1998) and is implemented in spTimer (Bakar, Sahu, et al., 2015). It is based on the concept of the posterior predictive distribution of a model, given a fitted dataset:

$$P(\mathbf{y_r}|\mathbf{y}) = \int \mathbf{P}(\mathbf{y_r}|\gamma)\mathbf{p}(\gamma|\mathbf{y})\mathbf{d}\gamma \tag{9}$$

where γ is a vector containing all the parameters of the model, $\mathbf{y}$ is the dataset used for fitting and $\mathbf{y_r}$ is a new, replicated dataset. In our case, $\mathbf{y}$ constitutes the land cover scores at all fitted spatiotemporal grid points, while γ includes the ancestry and climate coefficients, as well as spatiotemporal decay parameters. The posterior predictive distribution can be estimated by:

$$\hat{P}(\mathbf{y_r}|\mathbf{y}) = \frac{\mathbf{1}}{\mathbf{M}} \sum_{\mathbf{m=1}}^{\mathbf{M}} \mathbf{P}(\mathbf{y}|\hat{\gamma}_{\mathbf{m}}) \tag{10}$$

where $\hat{\gamma}_{\mathbf{m}}$ denotes the mth Monte Carlo sample of γ. The PMCC is then defined as:

$$PMCC = \sum_{i=1}^{n}(\mu_i - y_i)^2 + \sigma_i^2 \tag{11}$$

where μ_i and σ_i^2 are the expectation and the variance of a replicate $y_{r,i}$ coming from the posterior predictive distribution. In practice, these are obtained from the aforementioned estimate of this distribution. The first term of the sum serves as a goodness of fit score, while the second term is a penalty score, which tends to be large for both underfitted and overfitted models.

Nonparametric bootstrapping of parameter estimates

The Gibbs sampler allows us to obtain posterior estimates and 95% posterior credible intervals of the $\boldsymbol{\beta}$ parameters relating the explanatory to the response variables. However, it relies on the kriged ancestry gridpoint maps as input, so it does not account for the uncertainty in the estimation of these maps from the ancient genomes that we currently have. To address this, we derived confidence intervals on the Bayesian posterior estimates using a nonparametric bootstrapping approach. We created 100 pseudo-samples, by randomly sampling ancient and present-day genomes 100 times - with replacement - from among the list of all ancient and present-day genomes, then obtaining their ancestry assignments and kriging them on the spatiotemporal grid. We then fitted the Bayesian spatiotemporal model to each pseudo-sample, and thus obtained a distribution of bootstrapped $\boldsymbol{\beta}$ parameter estimates, from which we obtained 95% confidence intervals.

Arrival time maps

We first created ancestry arrival time maps by recording the time in each cell of the spatial grid at which the spatiotemporal surface map first reaches a value higher than a particular kriged ancestry proportion cutoff (Figure S21). In this case, we used a spatial grid of 5,000 points and 200-year time intervals. We find that these maps contain large proportions of missing data, in regions where an ancestry never reached the ancestry cutoff throughout the duration of the timeline. To correct for this, we instead recorded the times at which a particular ancestry first reached a value higher than $X\% * anc_{MAX}$ where $X\%$ is a chosen percentage cutoff for a particular ancestry and anc_{MAX}

is the maximum value that ancestry reaches at a spatial point throughout the duration of the timeline (Figure 8). Spatial points where anc_{MAX} is less than 10% were kept blank.

To create the cultural arrival maps for the spread of farming, we overlaid a 50x50km map covering Europe and selected, for each square, the oldest radiocarbon date directly associated with early farming. The dataset we used to obtain these dates came from the EUROFARM database, which contains 1,779 records of archaeological farming sites (Vander Linden in prep., downloadable here: https://github.com/mavdlind/Geostat_Farmer). It was then spatially kriged using the spatstat package (Baddeley, Turner, et al., 2004) in R (R Core Team, 2019).

Front speed estimation

To estimate the front speed of the spread of NEOL and YAM ancestries, we used a method developed by Pinhasi, Fort, and Ammerman (2005) and Silva and Steele (2014), in which the authors regress great-circle distances of sampled locations to a hypothesized migration origin against the time at which the migration reached those locations. The negative inverse of the slope is then an estimate of the migration front speed. We restricted to genomes older than 5,000 years BP for the NEOL ancestry spread and to genomes older than 3,000 years BP for the YAM ancestry spread. We used Cayönü (37.38N, 40.39E) as the NEOL ancestry origin, based on estimates of the Neolithic farmer expansion origin (Pinhasi, Fort, and Ammerman, 2005; Silva and Steele, 2014). We set various points at the center and extremes of the hypothesized original Yamnaya distribution in the Eurasian steppe as the YAM ancestry origin (Table S3). We used a ranged major axis (RMA) regression approach implemented in the R package lmodel2 (Borcard, Gillet, and Legendre, 2018), which assumes a symmetrical distribution of measurement error in both distance and time.

Acknowledgments

References

Allentoft, Morten E et al. (2015). "Population genomics of bronze age Eurasia". In: *Nature* 522.7555, p. 167.

Ammerman, Albert J and Luigi L Cavalli-Sforza (1971). "Measuring the rate of spread of early farming in Europe". In: *Man*, pp. 674–688.

Anthony, David W (2010). *The horse, the wheel, and language: how Bronze-Age riders from the Eurasian steppes shaped the modern world.* Princeton University Press.

Baddeley, Adrian J, Rolf Turner, et al. (2004). *Spatstat: An R package for analyzing spatial point pattens.*

Bakar, Khandoker Shuvo, Sujit K Sahu, et al. (2015). "spTimer: Spatio-temporal bayesian modelling using R". In: *Journal of Statistical Software* 63.15, pp. 1–32.

Barnett, William K (2000). "Cardial pottery and the agricultural transition in Mediterranean Europe". In: *Europe's first farmers*, pp. 93–116.

Battey, CJ, Peter L Ralph, and Andrew D Kern (2019). "Space is the Place: Effects of Continuous Spatial Structure on Analysis of Population Genetic Data". In: *BioRxiv*, p. 659235.

Bevan, Andrew et al. (2017). "Holocene fluctuations in human population demonstrate repeated links to food production and climate". In: *Proceedings of the National Academy of Sciences* 114.49, E10524–E10531.

Bickle, Penny and Alasdair Whittle (2013). *The first farmers of central Europe: diversity in LBK lifeways.* Oxbow Books.

Binder, Didier et al. (2018). "Modelling the earliest north-western dispersal of Mediterranean Impressed Wares: new dates and Bayesian chronological model." In: *Documenta praehistorica.* 44, pp. 54–77.

Bishop, Rosie R, Mike J Church, and Peter A Rowley-Conwy (2015). "Firewood, food and human niche construction: the potential role of Mesolithic hunter–gatherers in actively structuring Scotland's woodlands". In: *Quaternary Science Reviews* 108, pp. 51–75.

Borcard, Daniel, François Gillet, and Pierre Legendre (2018). *Numerical ecology with R*. Springer.

Bradburd, Gideon S, Graham M Coop, and Peter L Ralph (2018). "Inferring continuous and discrete population genetic structure across space". In: *Genetics* 210.1, pp. 33–52.

Brown, Jason L et al. (2018). "PaleoClim, high spatial resolution paleoclimate surfaces for global land areas". In: *Scientific data* 5, p. 180254.

Cheng, Jade Yu, Thomas Mailund, and Rasmus Nielsen (2017). "Fast admixture analysis and population tree estimation for SNP and NGS data". In: *Bioinformatics* 33.14, pp. 2148–2155.

Collins, William D et al. (2006). "The community climate system model version 3 (CCSM3)". In: *Journal of Climate* 19.11, pp. 2122–2143.

Cressie, Noel and Christopher K Wikle (2015). *Statistics for spatio-temporal data*. John Wiley & Sons.

Excoffier, Laurent et al. (2013). "Robust demographic inference from genomic and SNP data". In: *PLoS genetics* 9.10, e1003905.

Fick, Stephen E and Robert J Hijmans (2017). "WorldClim 2: new 1-km spatial resolution climate surfaces for global land areas". In: *International journal of climatology* 37.12, pp. 4302–4315.

Fordham, Damien A et al. (2017). "PaleoView: a tool for generating continuous climate projections spanning the last 21 000 years at regional and global scales". In: *Ecography* 40.11, pp. 1348–1358.

Fort, Joaquim (2015). "Demic and cultural diffusion propagated the Neolithic transition across different regions of Europe". In: *Journal of the Royal Society interface* 12.106, p. 20150166.

— (2018). "The Neolithic Transition: Diffusion of People or Diffusion of Culture?" In: *Diffusive spreading in nature, technology and society*. Springer, pp. 313–331.

Frantz, AC et al. (2009). "Using spatial Bayesian methods to determine the genetic structure of a continuously distributed population: clusters or isolation by distance?" In: *Journal of Applied Ecology* 46.2, pp. 493–505.

Fu, Qiaomei et al. (2016). "The genetic history of ice age Europe". In: *Nature* 534.7606, p. 200.

Fyfe, Ralph M, Claire Twiddle, et al. (2013). "The Holocene vegetation cover of Britain and Ireland: overcoming problems of scale and discerning patterns of openness". In: *Quaternary Science Reviews* 73, pp. 132–148.

Fyfe, Ralph M, Jessie Woodbridge, and Neil Roberts (2015). "From forest to farmland: pollen-inferred land cover change across Europe using the pseudobiomization approach". In: *Global Change Biology* 21.3, pp. 1197–1212.

Fyfe, Ralph, Neil Roberts, and Jessie Woodbridge (2010). "A pollen-based pseudobiomisation approach to anthropogenic land-cover change". In: *The Holocene* 20.7, pp. 1165–1171.

Gelfand, Alan E and Sujit K Ghosh (1998). "Model choice: a minimum posterior predictive loss approach". In: *Biometrika* 85.1, pp. 1–11.

Haak, Wolfgang et al. (2015). "Massive migration from the steppe was a source for Indo-European languages in Europe". In: *Nature* 522.7555, p. 207.

Hellenthal, Garrett et al. (2014). "A genetic atlas of human admixture history". In: *Science* 343.6172, pp. 747–751.

Hofmanová, Zuzana et al. (2016). "Early farmers from across Europe directly descended from Neolithic Aegeans". In: *Proceedings of the National Academy of Sciences* 113.25, pp. 6886–6891.

Hunt, Harriet V et al. (2018). "Genetic evidence for a western Chinese origin of broomcorn millet (Panicum miliaceum)". In: *The Holocene* 28.12, pp. 1968–1978.

Jacomet, Stefanie (2008). *Subsistenz und Landnutzung während des 3. Jahrtausends v. Chr. aufgrund von archäobotanischen Daten aus dem südwestlichen Mitteleuropa*. 84. Wachholtz.

Janes, Jasmine K et al. (2017). "The K= 2 conundrum". In: *Molecular Ecology* 26.14, pp. 3594–3602.

Joseph, Tyler A and Itsik Pe'er (2018). "Inference of population structure from ancient DNA". In: *International Conference on Research in Computational Molecular Biology*. Springer, pp. 90–104.

Kamm, Jack et al. (2019). "Efficiently inferring the demographic history of many populations with allele count data". In: *Journal of the American Statistical Association* just-accepted, pp. 1–42.

Kelleher, Jerome et al. (2018). "Inferring the ancestry of everyone". In: *BioRxiv*, p. 458067.

Kolář, Jan et al. (2018). "Population and forest dynamics during the Central European Eneolithic (4500–2000 BC)". In: *Archaeological and anthropological sciences* 10.5, pp. 1153–1164.

Kristiansen, Kristian, Morten E Allentoft, et al. (2017). "Re-theorising mobility and the formation of culture and language among the Corded Ware Culture in Europe". In: *Antiquity* 91.356, pp. 334–347.

Kristiansen, Kristian and Thomas B Larsson (2005). *The rise of Bronze Age society: travels, transmissions and transformations*. Cambridge University Press.

Lawson, Daniel J, Lucy Van Dorp, and Daniel Falush (2018). "A tutorial on how not to over-interpret STRUCTURE and ADMIXTURE bar plots". In: *Nature Communications* 9.1, p. 3258.

Lawson, Daniel John et al. (2012). "Inference of population structure using dense haplotype data". In: *PLoS genetics* 8.1, e1002453.

Lazaridis, Iosif, Dani Nadel, et al. (2016). "Genomic insights into the origin of farming in the ancient Near East". In: *Nature* 536.7617, p. 419.

Lazaridis, Iosif, Nick Patterson, et al. (2014). "Ancient human genomes suggest three ancestral populations for present-day Europeans". In: *Nature* 513.7518, p. 409.

Lechterbeck, Jutta et al. (2014). "Is Neolithic land use correlated with demography? An evaluation of pollen-derived land cover and radiocarbon-inferred demographic change from Central Europe". In: *The Holocene* 24.10, pp. 1297–1307.

Lipson, Mark et al. (2017). "Parallel palaeogenomic transects reveal complex genetic history of early European farmers". In: *Nature* 551.7680, pp. 368–372.

Liu, Zhengyu, Zhengyao Lu, et al. (2014). "Evolution and forcing mechanisms of El Niño over the past 21,000 years". In: *Nature* 515.7528, p. 550.

Liu, Zhengyu, BL Otto-Bliesner, et al. (2009). "Transient simulation of last deglaciation with a new mechanism for Bølling-Allerød warming". In: *Science* 325.5938, pp. 310–314.

Manen, Claire et al. (2019). "The Neolithic Transition in the Western Mediterranean: a Complex and Non-Linear Diffusion Process—The Radiocarbon Record Revisited". In: *Radiocarbon* 61.2, pp. 531–571.

Margaryan, Ashot et al. (2019). "Population genomics of the Viking world". In: *bioRxiv*, p. 703405.

Marquer, Laurent et al. (2017). "Quantifying the effects of land use and climate on Holocene vegetation in Europe". In: *Quaternary Science Reviews* 171, pp. 20–37.

Matérn, Bertil (1986). *Spatial variation, volume 36 of Lecture Notes in Statistics*.

Mathieson, Iain, Songül Alpaslan-Roodenberg, et al. (2018). "The genomic history of southeastern Europe". In: *Nature* 555.7695, p. 197.

Mathieson, Iain, Iosif Lazaridis, et al. (2015). "Genome-wide patterns of selection in 230 ancient Eurasians". In: *Nature* 528.7583, p. 499.

Mittnik, Alissa et al. (2019). "Kinship-based social inequality in Bronze Age Europe". In: *Science*, eaax6219.

Müller, Johannes (2015). "Eight Million Neolithic Europeans: Social Demography and Social Archaeology on the Scope of Change–From the Near East to Scandinavia". In: *Paradigm found: archaeological theory present past and future: essays in honour of Evžen Neustupn. Oxford: Oxbow*, pp. 200–214.

Müller, Johannes et al. (2009). "A revision of Corded Ware settlement pattern–new results from the Central European low mountain range". In: *Proceedings of the Prehistoric Society*. Vol. 75. Cambridge University Press, pp. 125–142.

Nielsen, Anne Birgitte et al. (2012). "Quantitative reconstructions of changes in regional openness in north-central Europe reveal new insights into old questions". In: *Quaternary Science Reviews* 47, pp. 131–149.

Olalde, Iñigo, Selina Brace, et al. (2018). "The Beaker phenomenon and the genomic transformation of northwest Europe". In: *Nature* 555.7695, p. 190.

Olalde, Iñigo, Hannes Schroeder, et al. (2015). "A common genetic origin for early farmers from Mediterranean Cardial and Central European LBK cultures". In: *Molecular biology and evolution* 32.12, pp. 3132–3142.

Otto-Bliesner, Bette L et al. (2006). "Climate sensitivity of moderate-and low-resolution versions of CCSM3 to preindustrial forcings". In: *Journal of Climate* 19.11, pp. 2567–2583.

Patterson, Nick et al. (2012). "Ancient admixture in human history". In: *Genetics* 192.3, pp. 1065–1093.

Pebesma, Edzer et al. (2012). "spacetime: Spatio-temporal data in r". In: *Journal of Statistical Software* 51.7, pp. 1–30.

Pebesma, Edzer J (2004). "Multivariable geostatistics in S: the gstat package". In: *Computers & Geosciences* 30.7, pp. 683–691.

Pebesma, Edzer and Gerard Heuvelink (2016). "Spatio-temporal interpolation using gstat". In: *RFID Journal* 8.1, pp. 204–218.

Pinhasi, Ron, Joaquim Fort, and Albert J Ammerman (2005). "Tracing the origin and spread of agriculture in Europe". In: *PLoS biology* 3.12, e410.

R Core Team (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. URL: https://www.R-project.org.

Roberts, C Neil et al. (2019). "Mediterranean landscape change during the Holocene: Synthesis, comparison and regional trends in population, land cover and climate". In: *The Holocene* 29.5, pp. 923–937.

Roberts, Neil et al. (2018). "Europe's lost forests: a pollen-based synthesis for the last 11,000 years". In: *Scientific reports* 8.1, p. 716.

Sánchez-Quinto, Federico et al. (2012). "Genomic affinities of two 7,000-year-old Iberian hunter-gatherers". In: *Current biology* 22.16, pp. 1494–1499.

Schauer, Peter et al. (2019). "Supply and demand in prehistory? Economics of Neolithic mining in northwest Europe". In: *Journal of Anthropological Archaeology* 54, pp. 149–160.

Seregély, Timo and Johannes Müller (2008). "Endneolithische Siedlungsstrukturen in Oberfranken: Wattendorf-Motzenstein: eine schnurkeramische Siedlung in der nördlichen Frankenalb". In:

Shennan, Stephen et al. (2013). "Regional population collapse followed initial agriculture booms in mid-Holocene Europe". In: *Nature communications* 4, p. 2486.

Shishlina, Natalia I (2008). *Reconstruction of the Bronze Age of the Caspian steppes: life styles and life ways of pastoral nomads*. British Archaeological Reports Ltd.

Sikora, Martin et al. (2014). "Population genomic analysis of ancient and modern genomes yields new insights into the genetic ancestry of the Tyrolean Iceman and the genetic structure of Europe". In: *PLoS genetics* 10.5, e1004353.

Silva, Fabio and James Steele (2014). "New methods for reconstructing geographical effects on dispersal rates and routes from large-scale radiocarbon databases". In: *Journal of Archaeological Science* 52, pp. 609–620.

Sjögren, Karl-Göran, T Douglas Price, and Kristian Kristiansen (2016). "Diet and mobility in the Corded Ware of Central Europe". In: *PloS one* 11.5, e0155083.

Skoglund, Pontus, Helena Malmström, Ayça Omrak, et al. (2014). "Genomic diversity and admixture differs for Stone-Age Scandinavian foragers and farmers". In: *Science* 344.6185, pp. 747–750.

Skoglund, Pontus, Helena Malmström, Maanasa Raghavan, et al. (2012). "Origins and genetic legacy of Neolithic farmers and hunter-gatherers in Europe". In: *Science* 336.6080, pp. 466–469.

Speidel, Leo et al. (2019). "A method for genome-wide genealogy estimation for thousands of samples". In: *BioRxiv*, p. 550558.

Vander Linden, Marc and Fabio Silva (2018). "Comparing and modeling the spread of early farming across Europe." In: *PAGES magazine* 26.1, pp. 28–29.

Vandkilde, Helle (2007). *Culture and change in Central European prehistory*. ISD LLC.

Walvoort, Dennis JJ and Jaap J de Gruijter (2001). "Compositional kriging: a spatial interpolation method for compositional data". In: *Mathematical Geology* 33.8, pp. 951–966.

Warren, Graeme et al. (2014). "The potential role of humans in structuring the wooded landscapes of Mesolithic Ireland: a review of data and discussion of approaches". In: *Vegetation History and Archaeobotany* 23.5, pp. 629–646.

Woodbridge, Jessie, Ralph M Fyfe, et al. (2014). "The impact of the Neolithic agricultural transition in Britain: a comparison of pollen-based land-cover and archaeological 14C date-inferred population change". In: *Journal of Archaeological Science* 51, pp. 216–224.

Woodbridge, Jessie, Neil Roberts, and Ralph Fyfe (2018). "Pan-Mediterranean Holocene vegetation and land-cover dynamics from synthesized pollen data". In: *Journal of Biogeography* 45.9, pp. 2159–2174.

Yeager, Stephen G et al. (2006). "The low-resolution CCSM3". In: *Journal of Climate* 19.11, pp. 2545–2566.
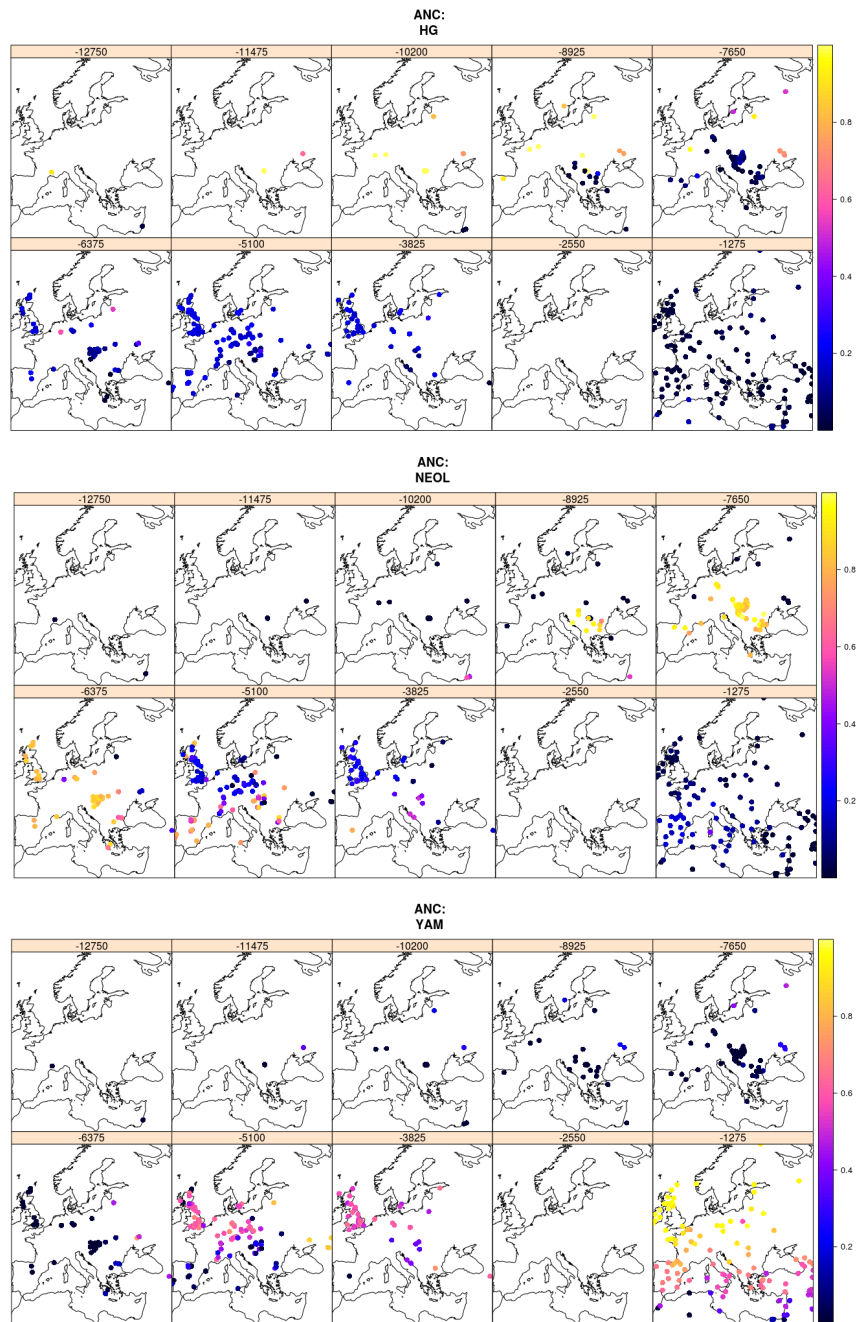
Figures

Figure 1: Spatiotemporal maps of ancestry proportions for ancient and present-day genomes in this study. The year in each panel's title is the year of the most ancient sample in each panel, so not all ancient samples in each panel are strictly contemporaneous with each other. ANC:HG = hunter-gatherer ancestry. ANC:NEOL = Neolithic farmer ancestry. ANC:YAM = Yamnaya steppe ancestry.
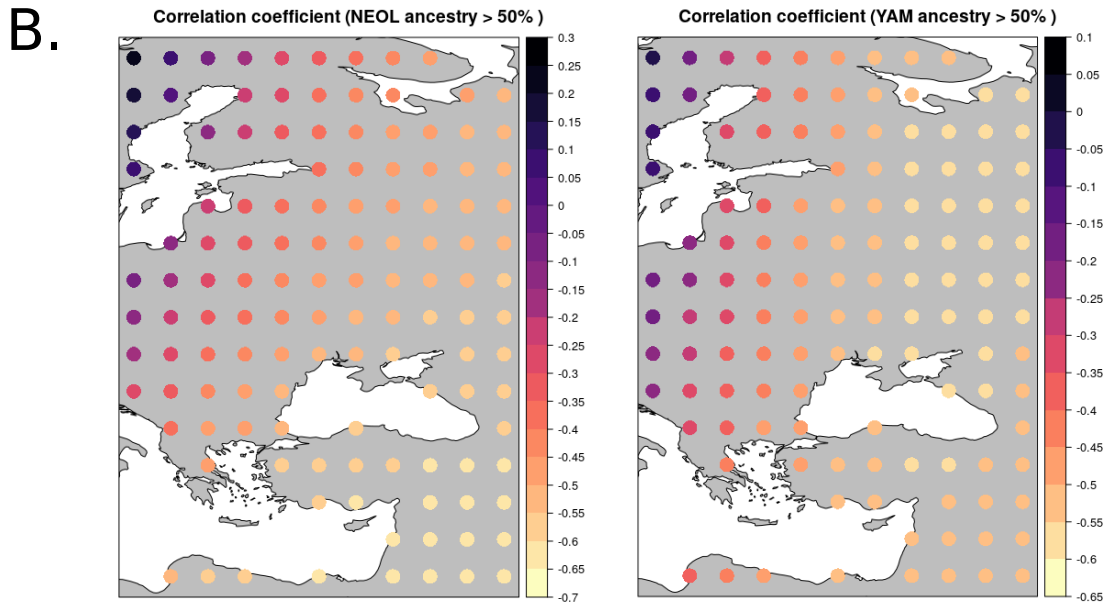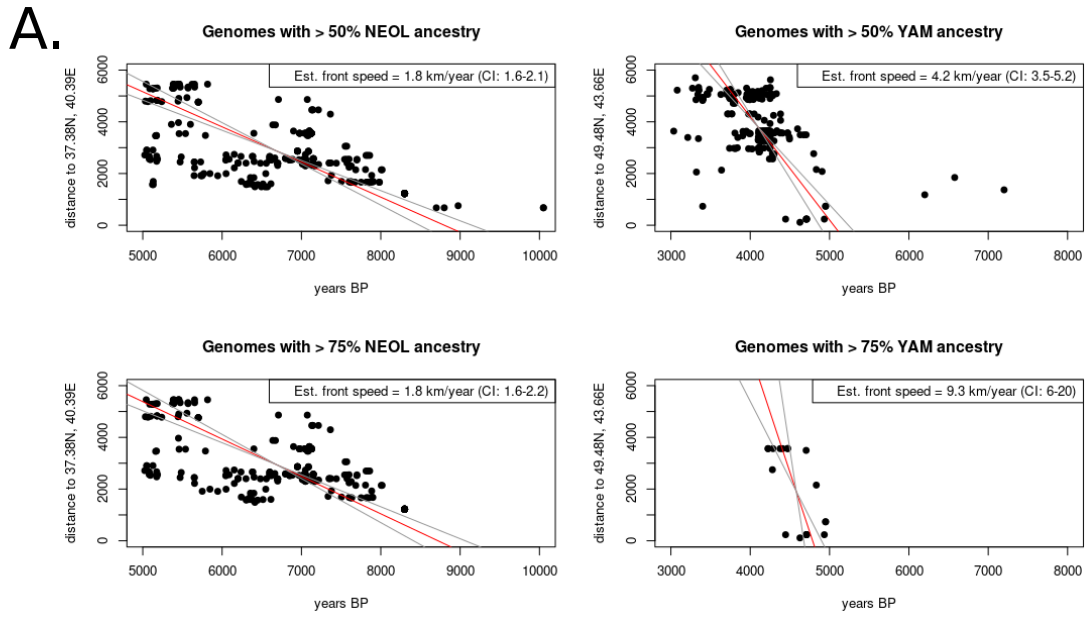
Figure 2: A. Front speed estimation for the Neolithic farmer (left column) and Yamnaya steppe peoples (right column) population movements. We used an RMA regression on time against distance from the hypothesized origin of the spread, to estimate average migration front speed. We used a lenient ($> 50\%$) ancestry cutoff to define genomes as belong to a particular migration wave (top row) and a more conservative ($> 75\%$) ancestry cutoff (bottom row). B. Point of origin estimation. We computed the correlation coefficient between time of sampling and distance from a hypothesized origin, which should be negative for a range expansion. Each dot in the map represents a different hypothesized origin.
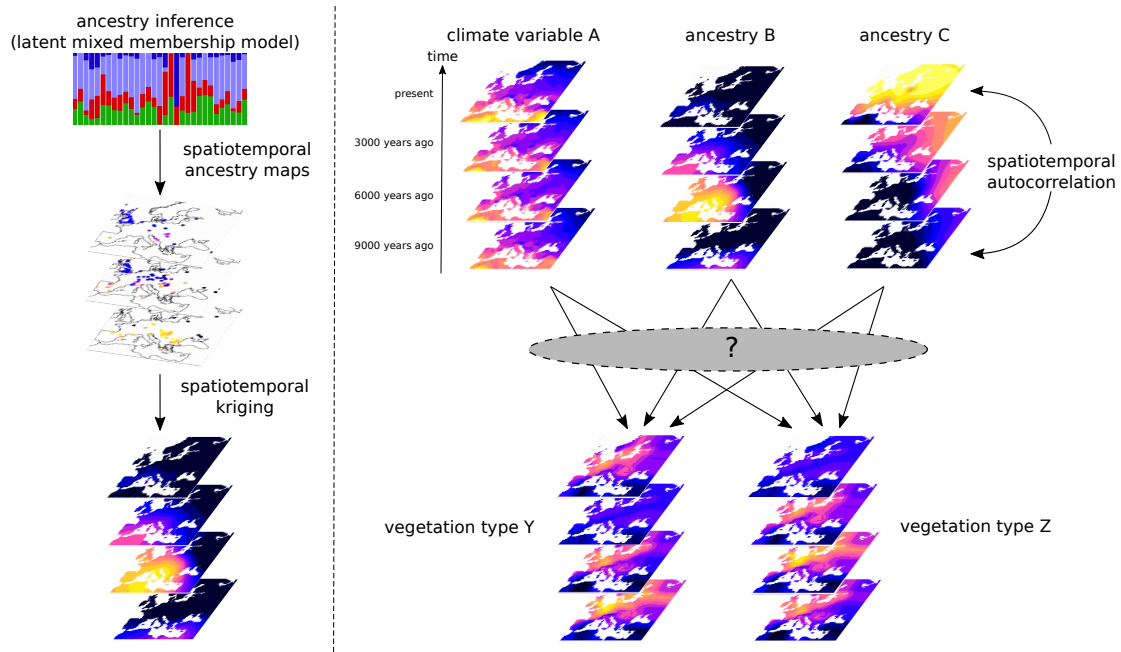
Figure 3: Schematic of methodology for spatiotemporal kriging and vegetation modelling. Left panel: We first fitted a latent mixed membership model to the ancient and present-day genomes. The ancestry proportions are then assigned the temporal and spatial meta-data of their respective genomes, which allows us to perform spatiotemporal kriging to any location and time in the European Holocene. Right panel: We used a spatiotemporally-aware model to understand how patterns of human migration and climate relate to patterns of vegetation type changes during the European Holocene, while accounting for spatiotemporal autocorrelation. We used a bootstrapping method to account for biases due to uneven sampling of ancient genomes. Brighter colors represent higher values of each depicted variable.
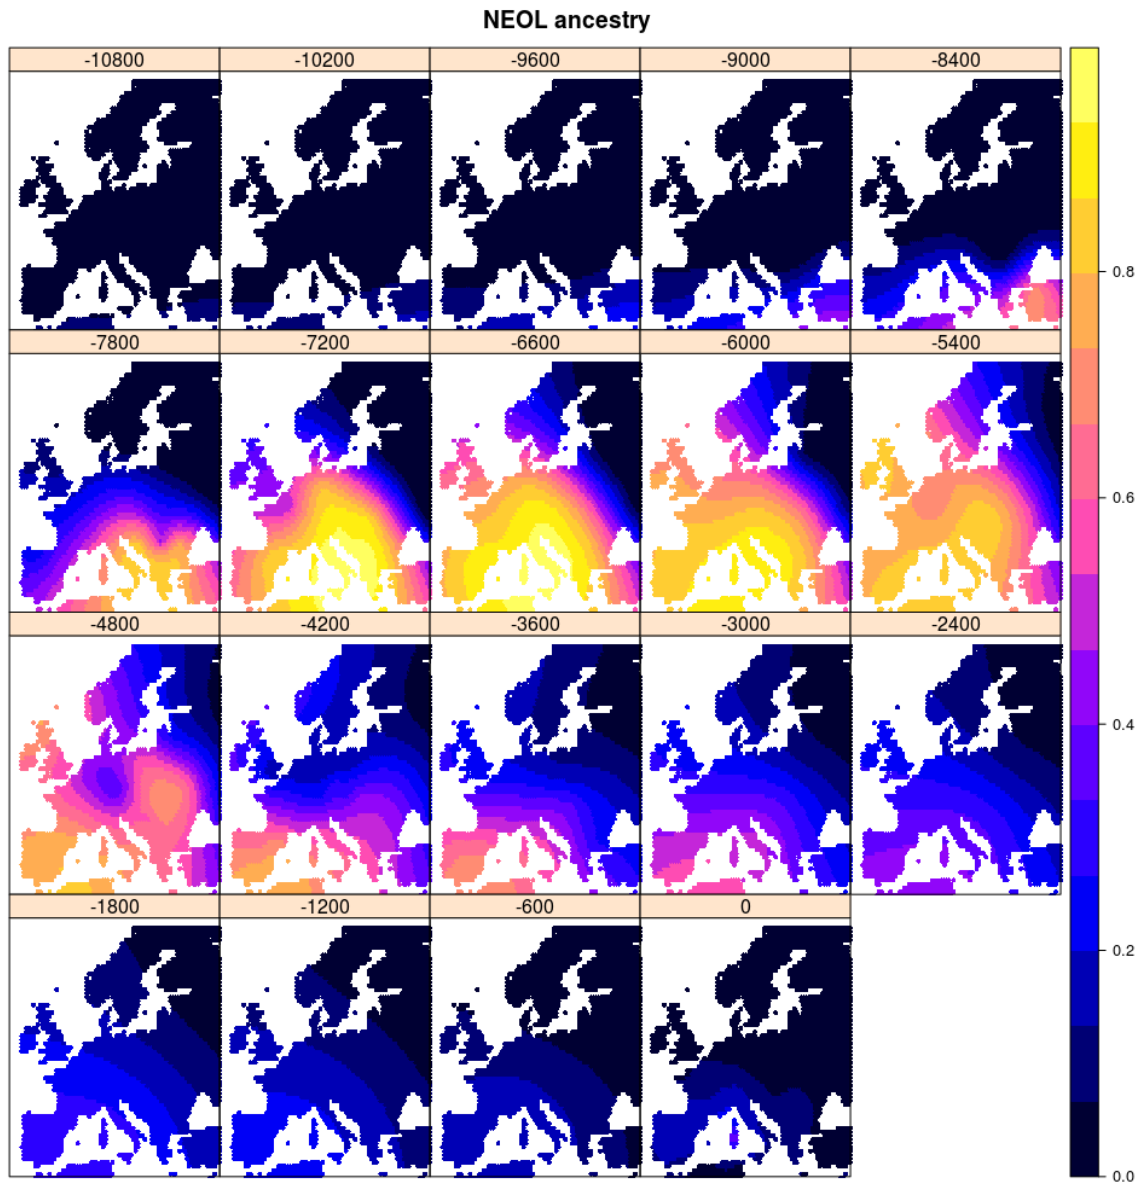
Figure 4: Spatiotemporal kriging of NEOL ancestry during the Holocene, using 5000 spatial grid points. The colors represent the predicted ancestry proportion at each point in the grid.

20
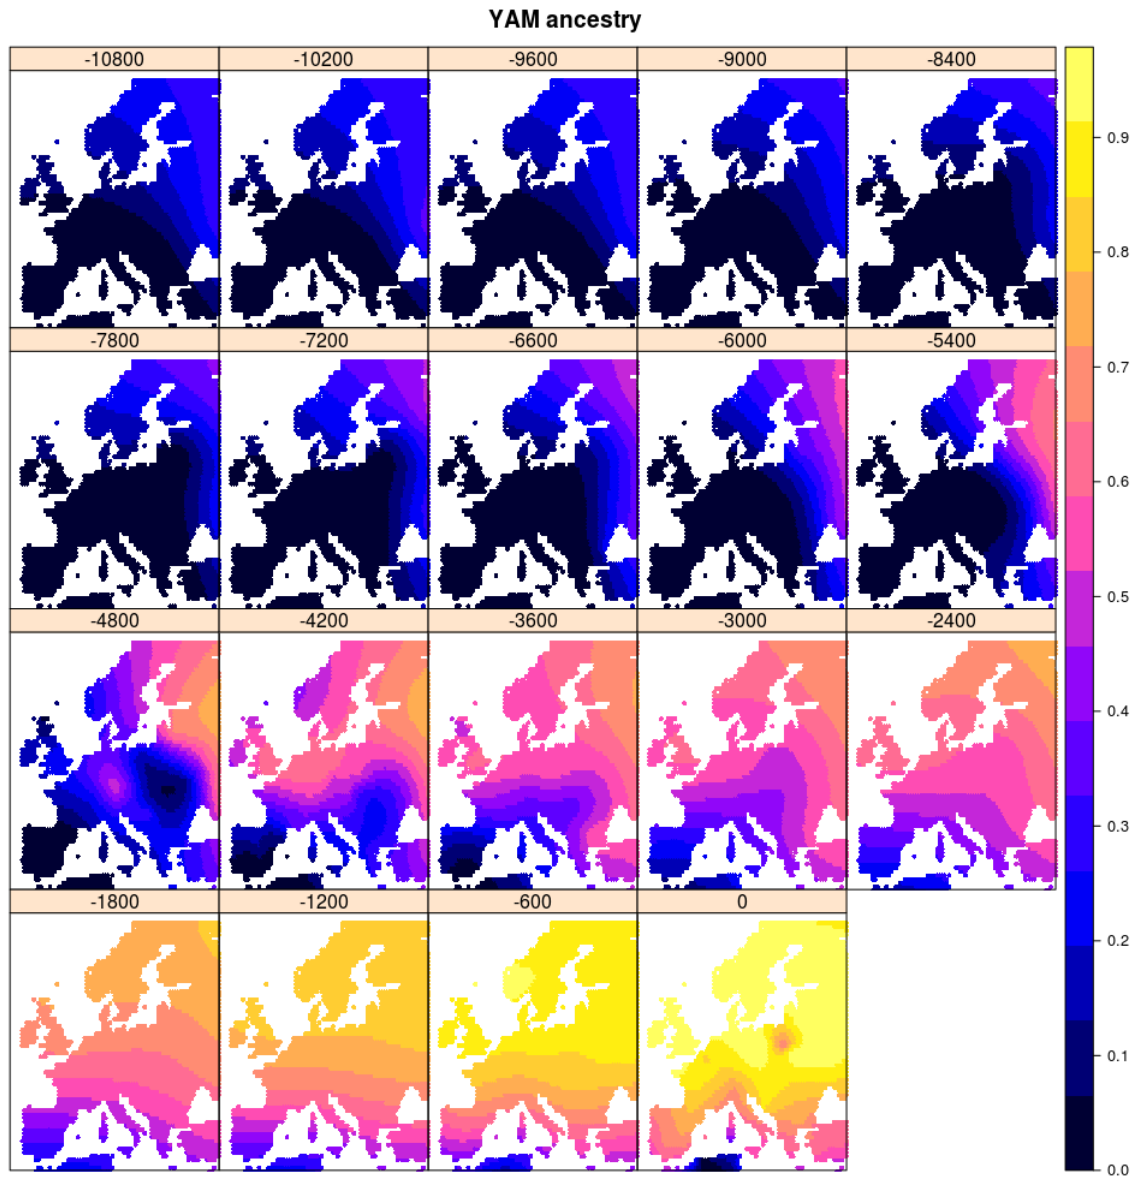
Figure 5: Spatiotemporal kriging of YAM steppe ancestry during the Holocene, using 5000 spatial grid points. The colors represent the predicted ancestry proportion at each point in the grid.
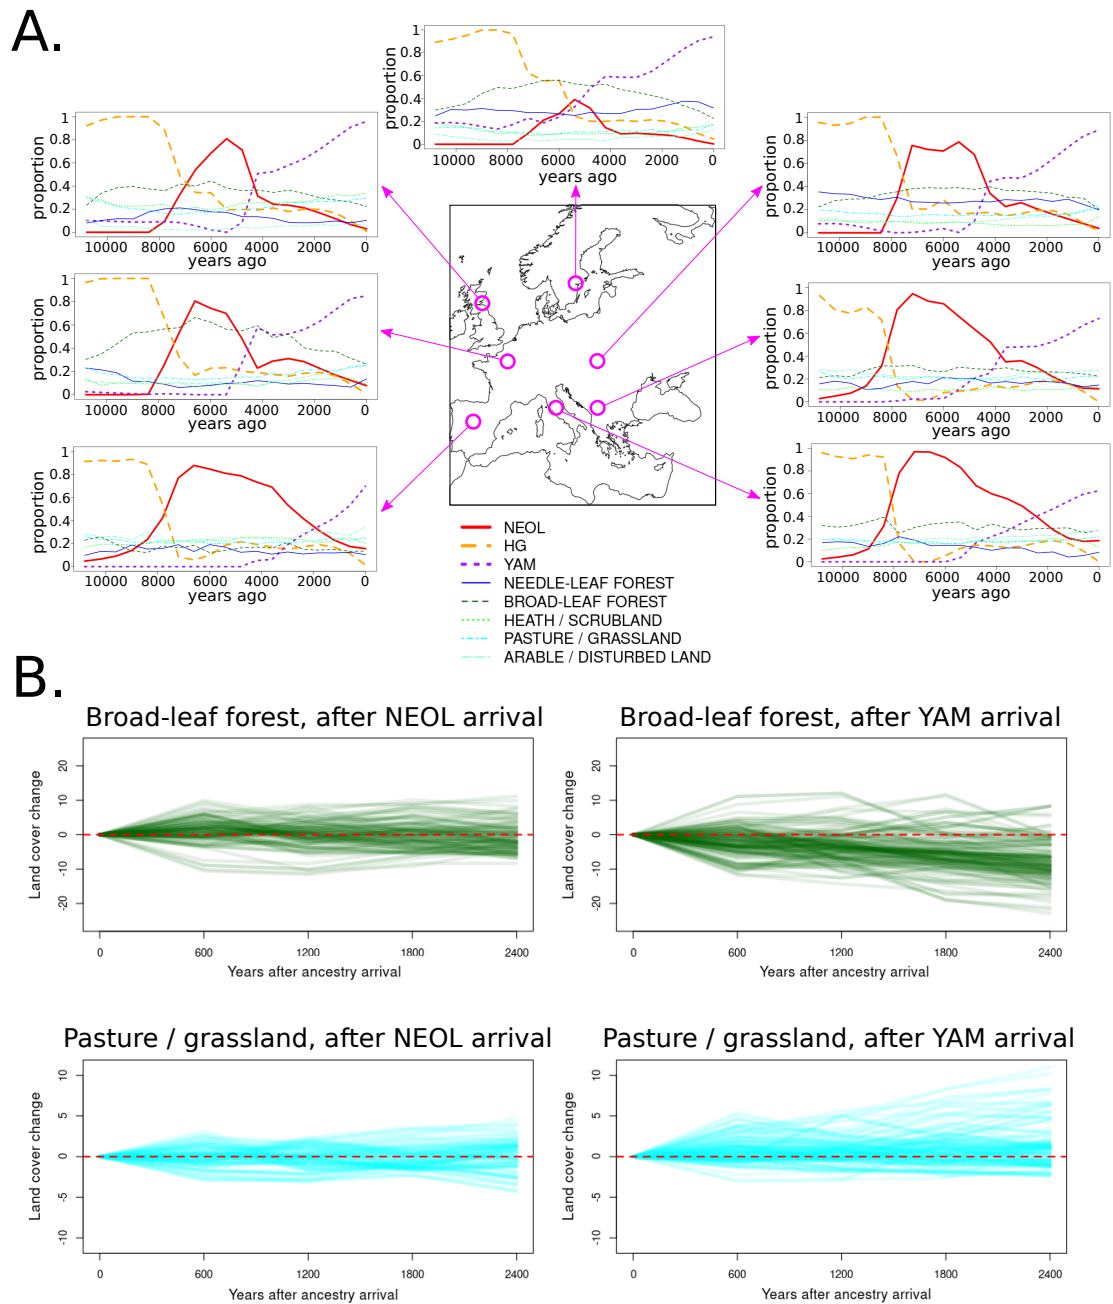
Figure 6: A. Timelines of kriged ancestry and vegetation type proportions at different points in Europe. NEOL = Neolithic farmer ancestry. HG = Hunter-gatherer ancestry. YAM = Yamnaya steppe ancestry. B. Change in pasture / natural grassland and broad-leaf forest cover composition after the arrival (first time there is > 50% ancestry) in each spatial grid point of YAM and NEOL ancestry. Each line corresponds to the post-arrival progression of a different spatial grid point.
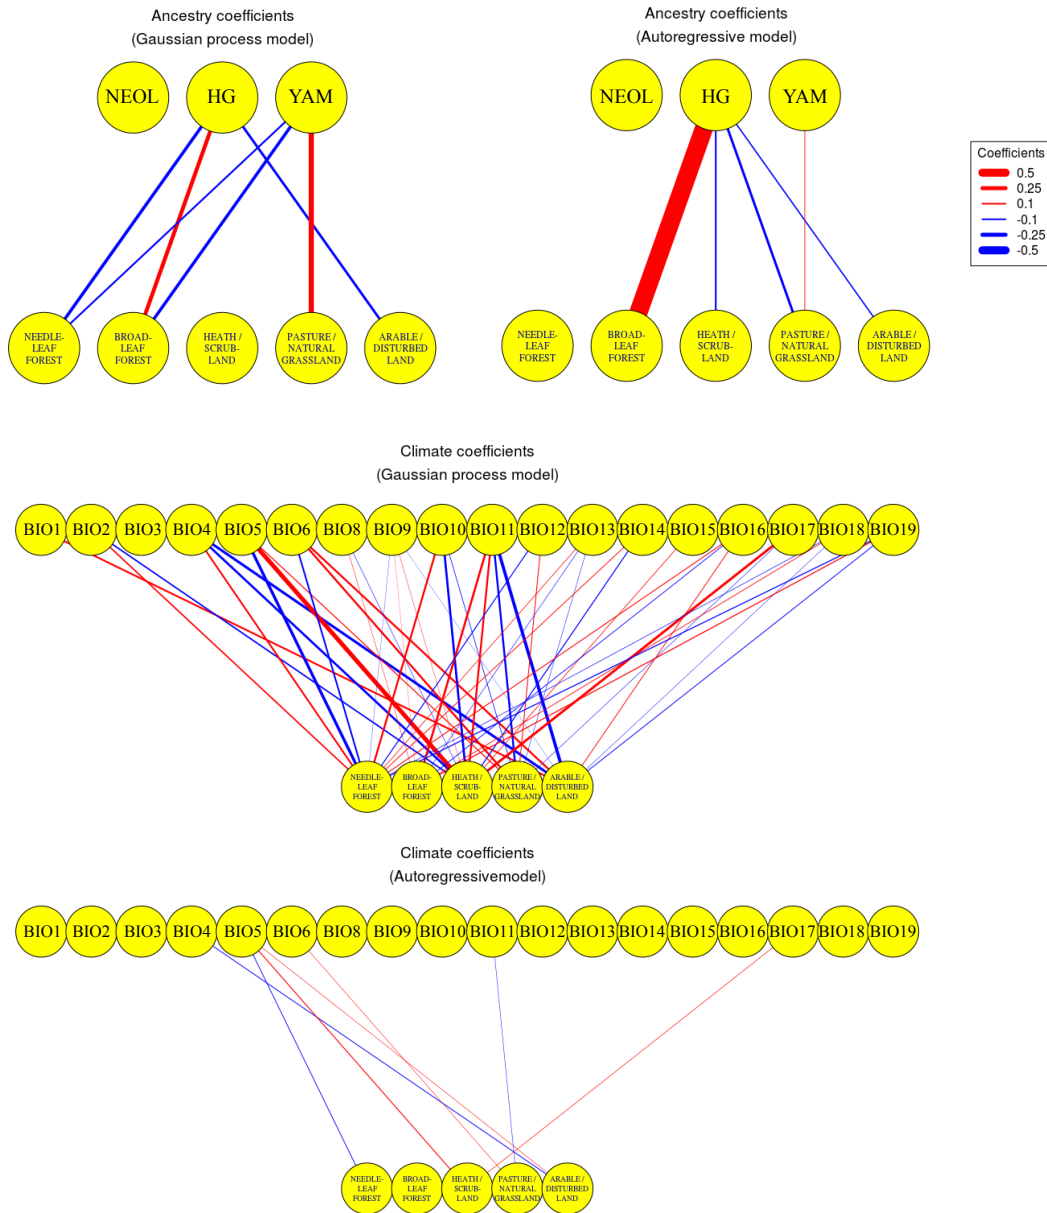
Figure 7: Posterior mean coefficients of spatiotemporal model for paleovegetation anomalies, using kriged ancestry anomalies and anomalies from simulation-based paleoclimate reconstructions as explanatory variables. Upper left and middle panels: posterior coefficients from Gaussian process model. Upper right and bottom panels: coefficients from autoregressive model. Coefficients whose corresponding posterior distribution has a 95% central probability mass interval that spans the value of 0 are not depicted. NEOL = neolithic farmer ancestry. HG = hunter-gatherer ancestry. YAM = Yamnaya steppe ancestry. The climate variables follow the WorldClim nomenclature. BIO1 = Annual Mean Temperature. BIO2 = Mean Diurnal Range (Mean of monthly (max temp - min temp)). BIO3 = Isothermality (BIO2/BIO7) (* 100). BIO4 = Temperature Seasonality (standard deviation *100). BIO5 = Max Temperature of Warmest Month. BIO6 = Min Temperature of Coldest Month. BIO8 = Mean Temperature of Wettest Quarter. BIO9 = Mean Temperature of Driest Quarter. BIO10 = Mean Temperature of Warmest Quarter. BIO11 = Mean Temperature of Coldest Quarter. BIO12 = Annual Precipitation. BIO13 = Precipitation of Wettest Month. BIO14 = Precipitation of Driest Month. BIO15 = Precipitation Seasonality (Coefficient of Variation). BIO16 = Precipitation of Wettest Quarter. BIO17 = Precipitation of Driest Quarter. BIO18 = Precipitation of Warmest Quarter. BIO19 = Precipitation of Coldest Quarter.
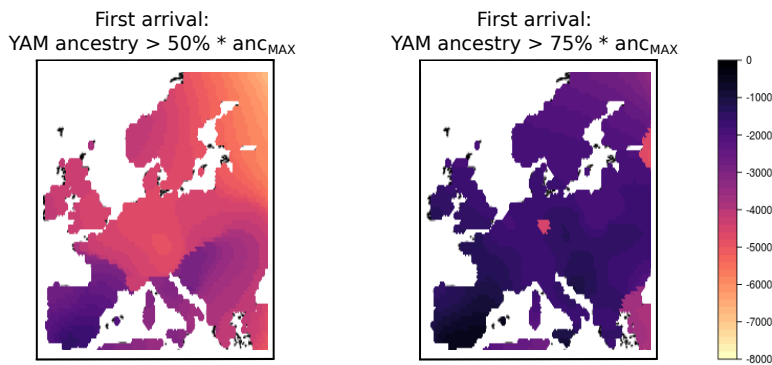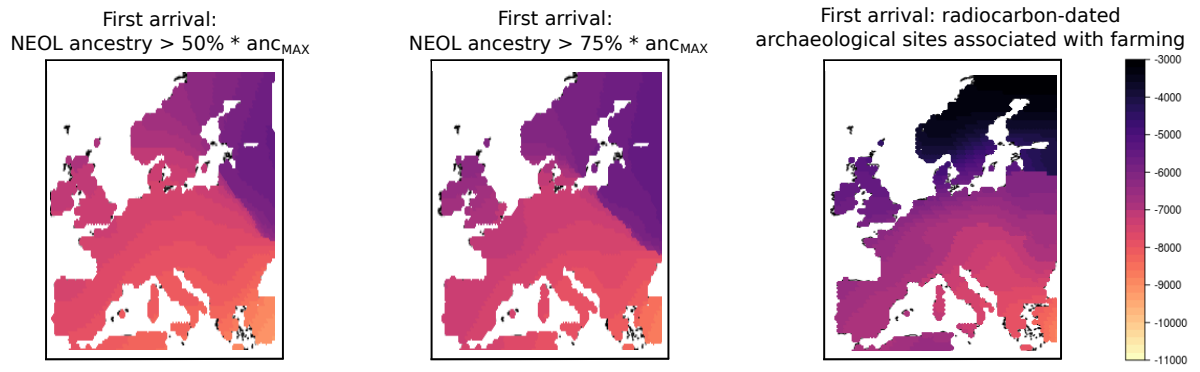
Figure 8: Comparison of inferred spread of farming from archaeological sites and spread of NEOL (upper panel) and YAM ancestries (lower panels). The left panels define first arrival as the first time slice in which a grid point has more than $50\% * anc_{MAX}$ of the ancestry depicted, where anc_{MAX} is the maximum value that ancestry reaches at that point throughout the entire timeline. The center panels are the result of using a more strict cutoff: $75\% * anc_{MAX}$ ancestry. The top right panel is a spatially kriged map of first arrivals of farming practices, based on radiocarbon-dated archaeological sites.

Tables

Coefficient	Posterior median	2.5% lower CI	97.5% upper CI
NEOL → Needle-leaf forest	-0.0567	-0.149	0.035
HG → Needle-leaf forest	-0.196	-0.329	-0.0654
YAM → Needle-leaf forest	-0.114	-0.183	-0.0434
NEOL → Broad-leaf forest	-0.0396	-0.113	0.033
HG → Broad-leaf forest	0.245	0.138	0.344
YAM → Broad-leaf forest	-0.191	-0.244	-0.139
NEOL → Heath / scrubland	0.0545	-0.0584	0.167
HG → Heath / scrubland	0.0426	-0.114	0.2
YAM → Heath / scrubland	0.0534	-0.0262	0.132
NEOL → Pasture / natural grassland	0.0481	-0.0026	0.0992
HG → Pasture / natural grassland	0.0196	-0.0527	0.0947
YAM → Pasture / natural grassland	0.318	0.28	0.356
NEOL → Arable / disturbed land	-0.0603	-0.152	0.0275
HG → Arable / disturbed land	-0.154	-0.282	-0.0285
YAM → Arable / disturbed land	0.0455	-0.016	0.112

Table 1: Coefficients of the spatiotemporal Gaussian process model, using the PaleoClim (Brown et al., 2018) simulation-based paleoclimate variables as covariates, with 95% Bayesian credible intervals. Coefficients whose 95% credible intervals do not overlap with 0 are in bold.

Coefficient	Posterior median	2.5% lower CI	97.5% upper CI
NEOL → Needle-leaf forest	-0.0114	-0.0712	0.0492
HG → Needle-leaf forest	-0.00947	-0.0938	0.0801
YAM → Needle-leaf forest	-0.00891	-0.0533	0.0343
NEOL → Broad-leaf forest	0.324	-0.032	0.105
HG → Broad-leaf forest	0.225	0.121	0.331
YAM → Broad-leaf forest	-0.00494	-0.0462	0.0361
NEOL → Heath / scrubland	-0.0289	-0.0929	0.0361
HG → Heath / scrubland	-0.105	-0.199	-0.0107
YAM → Heath / scrubland	-0.00871	-0.0556	0.0374
NEOL → Pasture / natural grassland	-0.0271	-0.0607	0.00651
HG → Pasture / natural grassland	-0.153	-0.207	-0.0994
YAM → Pasture / natural grassland	0.0436	0.0219	0.0657
NEOL → Arable / disturbed land	-0.0392	-0.0827	0.0043
HG → Arable / disturbed land	-0.0831	-0.146	-0.0208
YAM → Arable / disturbed land	-0.0111	-0.0415	0.0212

Table 2: Coefficients of the spatiotemporal autoregressive model, using the PaleoClim (Brown et al., 2018) simulation-based paleoclimate variables as covariates, with 95% Bayesian credible intervals. Coefficients whose 95% credible intervals do not overlap with 0 are in bold.

Model	None	Climate only	Ancestry only	Both
Needle-leaf forest	2579.09	2514.75	2574.54	2508.52
Broad-leaf forest	1671.87	1651.12	1641.98	1623.55
Heath / scrubland	2759.30	2640.54	2755.95	2635.82
Pasture / natural grassland	1961.25	1408.12	1662.91	1123.51
Arable / disturbed land	2171.43	2151.85	2161.47	2145.55

Table 3: Predictive model choice criteria for the hierarchical Gaussian process model of vegetation anomalies with a Uniform(0.01,0.02) prior for the spatial decay parameter. We compared models including climate variables only, ancestry variables only, neither set of variables or both of them.
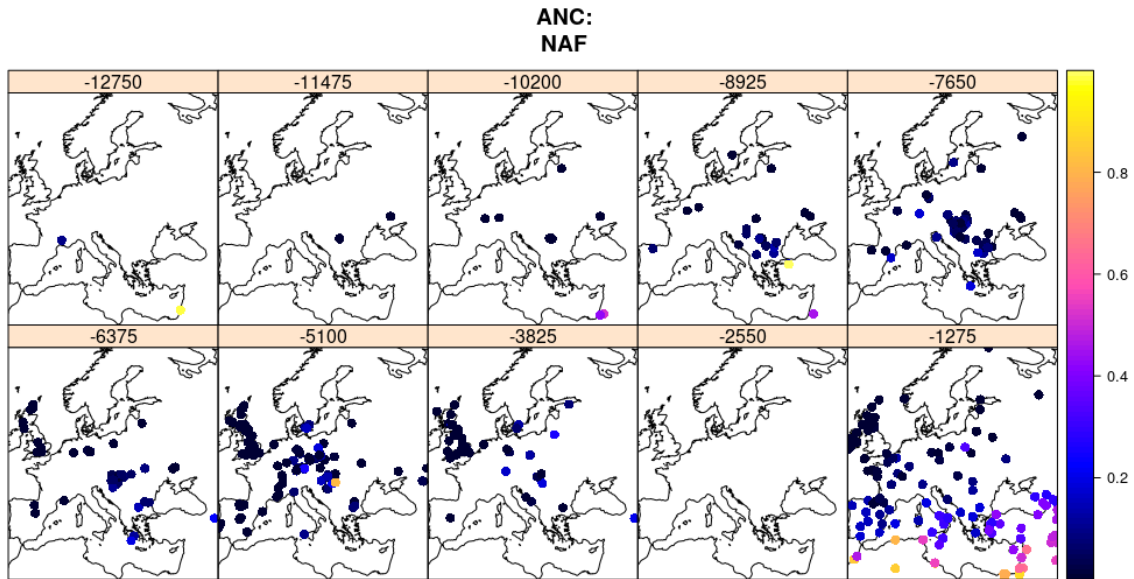
Supplementary Figures

Figure S1: Spatiotemporal maps of NAF ancestry proportion for ancient and present-day genomes in this study. The year in each panel's title is the year of the most ancient sample in each panel, so not all ancient samples in each panel are strictly contemporaneous with each other.
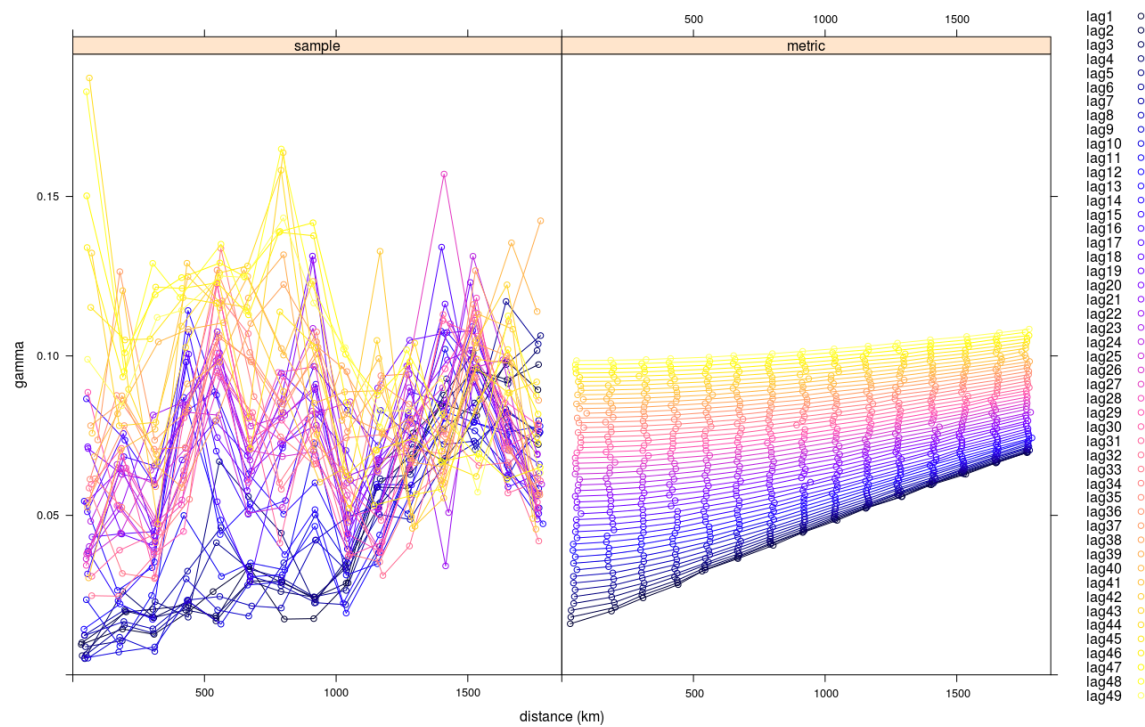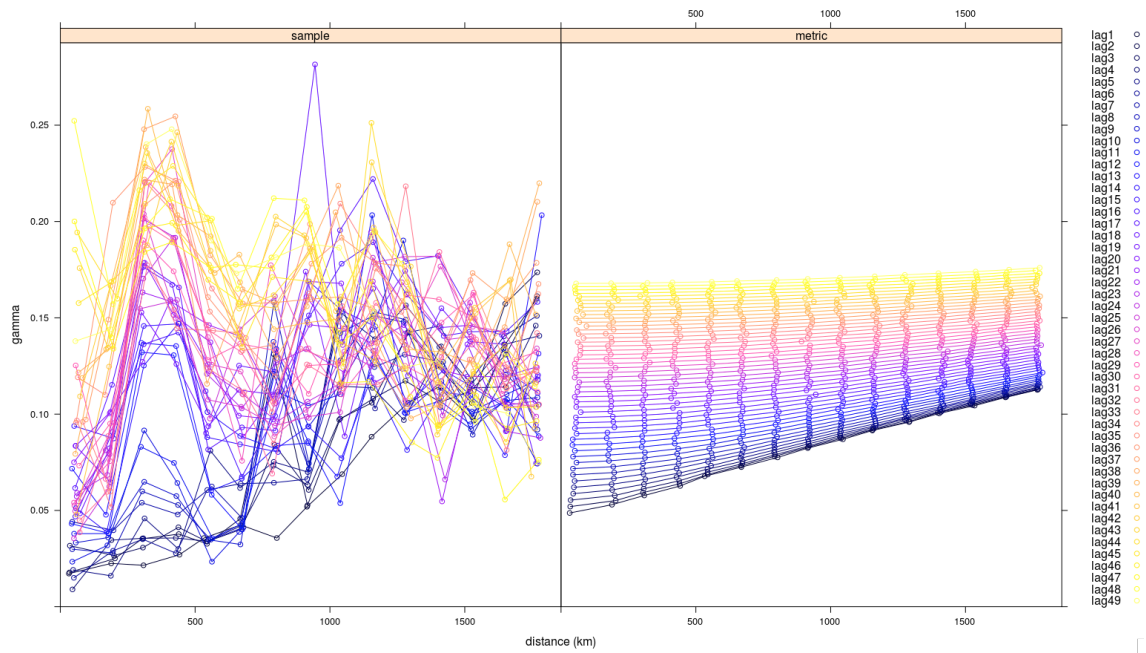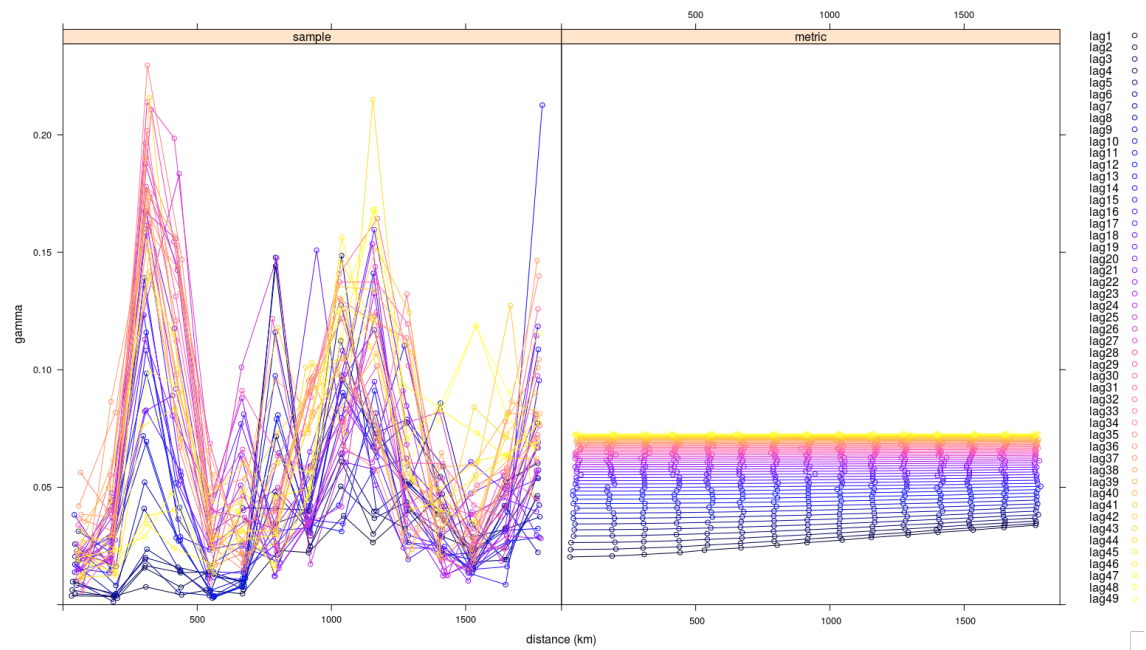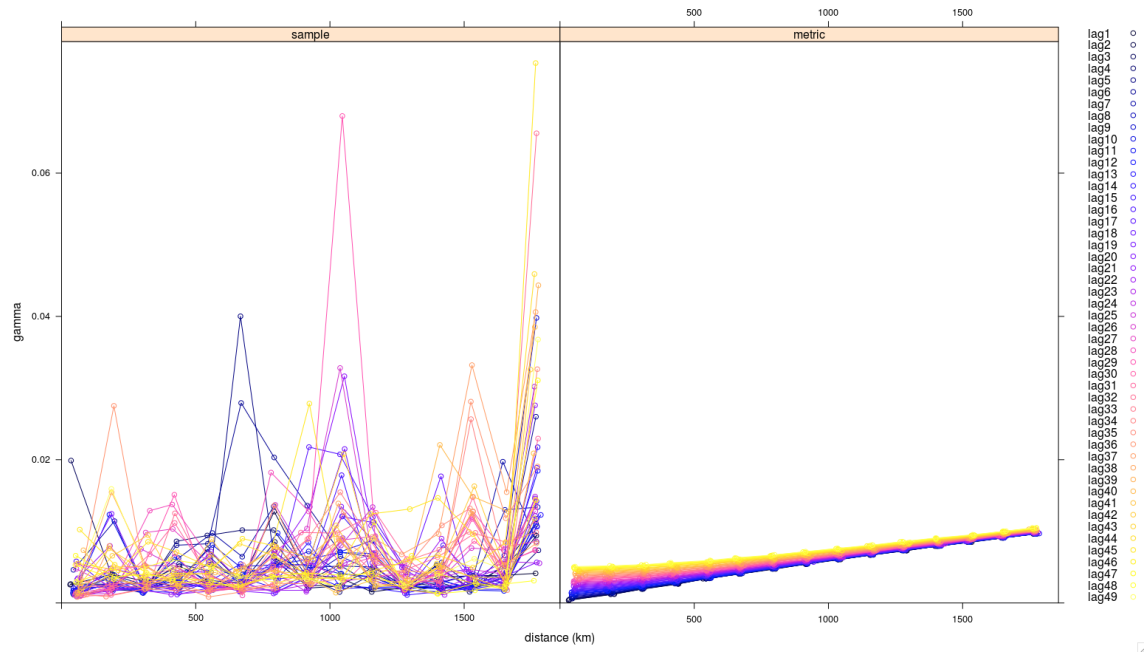

Figure S2: Spatiotemporal semivariogram of YAM ancestry in Europe, with temporal lags of 60-year increments. Left panel: empirical semivariogram. Right panel: fitted semivariogram model using the metric method. The lags are in 60-year increments.

Figure S3: Spatiotemporal semivariogram of NEOL ancestry in Europe, with temporal lags of 60-year increments. Left panel: empirical semivariogram. Right panel: fitted semivariogram model using the metric method. The lags are in 60-year increments.

Figure S4: Spatiotemporal semivariogram of HG ancestry in Europe, with temporal lags of 60-year increments. Left panel: empirical semivariogram. Right panel: fitted semivariogram model using the metric method. The lags are in 60-year increments.

Figure S5: Spatiotemporal semivariogram of NAF ancestry, with temporal lags of 60-year increments. Left panel: empirical semivariogram. Right panel: fitted semivariogram model using the metric method. The lags are in 60-year increments.
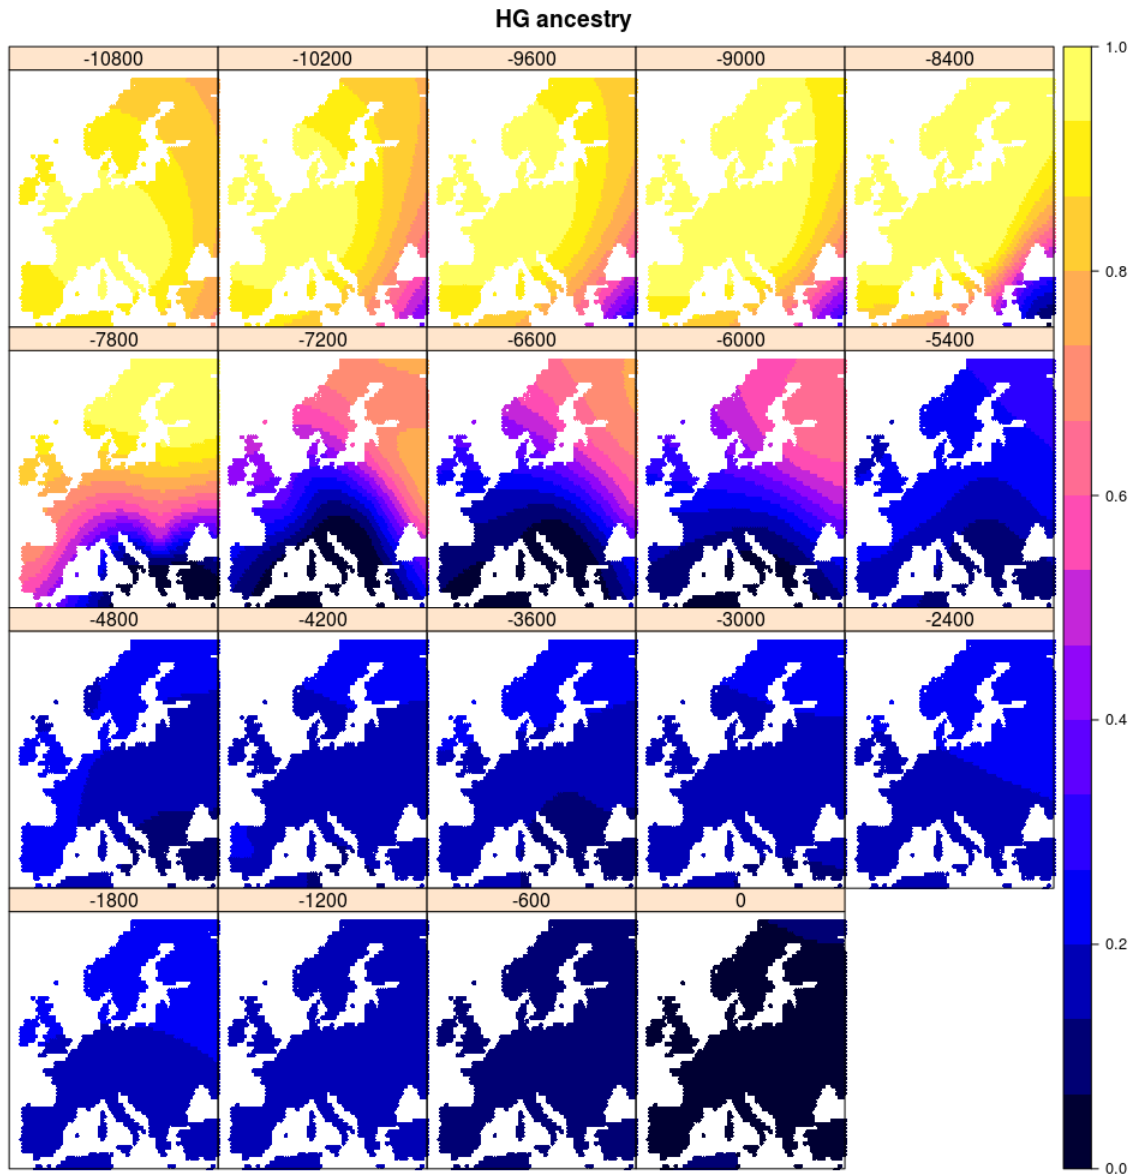
Figure S6: Spatiotemporal kriging of HG ancestry during the Holocene, using 5000 spatial grid points. The colors represent the predicted ancestry proportion at each point in the grid.
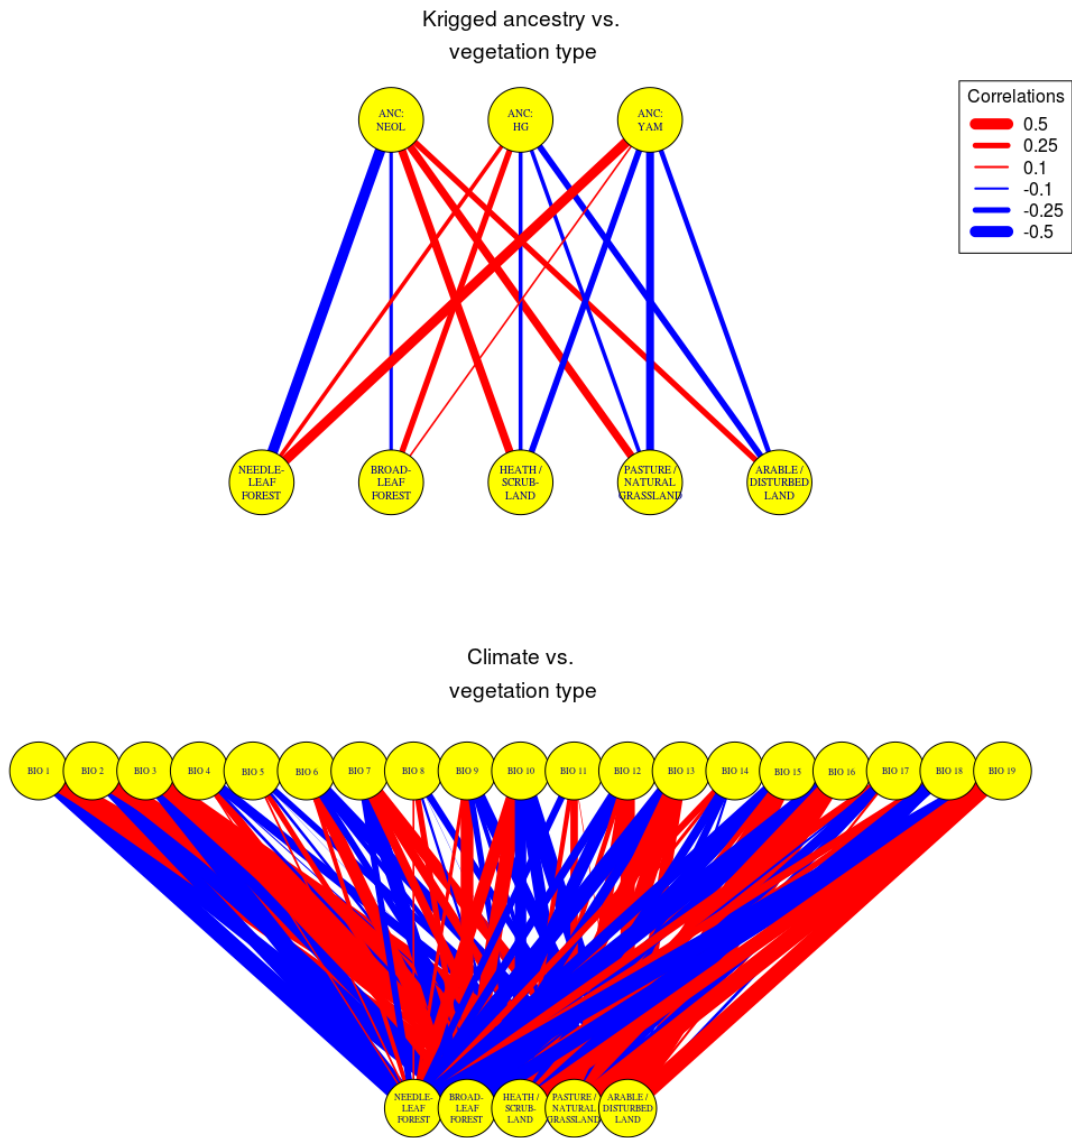
Figure S7: Spatiotemporal kriging of NAF ancestry during the Holocene, using 5000 spatial grid points. The colors represent the predicted ancestry proportion at each point in the grid.

Figure S8: Correlations between kriged ancestry and vegetation type (top) or between climate variables and vegetation type (bottom), computed across time and space. ANC = ancestry. NEOL = neolithic farmer. HG = hunter-gatherer. YAM = Yamnaya. The climate variables follow the WorldClim nomenclature. BIO1 = Annual Mean Temperature. BIO2 = Mean Diurnal Range (Mean of monthly (max temp - min temp)). BIO3 = Isothermality (BIO2/BIO7) (* 100). BIO4 = Temperature Seasonality (standard deviation *100). BIO5 = Max Temperature of Warmest Month. BIO6 = Min Temperature of Coldest Month. BIO7 = Temperature Annual Range (BIO5-BIO6). BIO8 = Mean Temperature of Wettest Quarter. BIO9 = Mean Temperature of Driest Quarter. BIO10 = Mean Temperature of Warmest Quarter. BIO11 = Mean Temperature of Coldest Quarter. BIO12 = Annual Precipitation. BIO13 = Precipitation of Wettest Month. BIO14 = Precipitation of Driest Month. BIO15 = Precipitation Seasonality (Coefficient of Variation). BIO16 = Precipitation of Wettest Quarter. BIO17 = Precipitation of Driest Quarter. BIO18 = Precipitation of Warmest Quarter. BIO19 = Precipitation of Coldest Quarter.
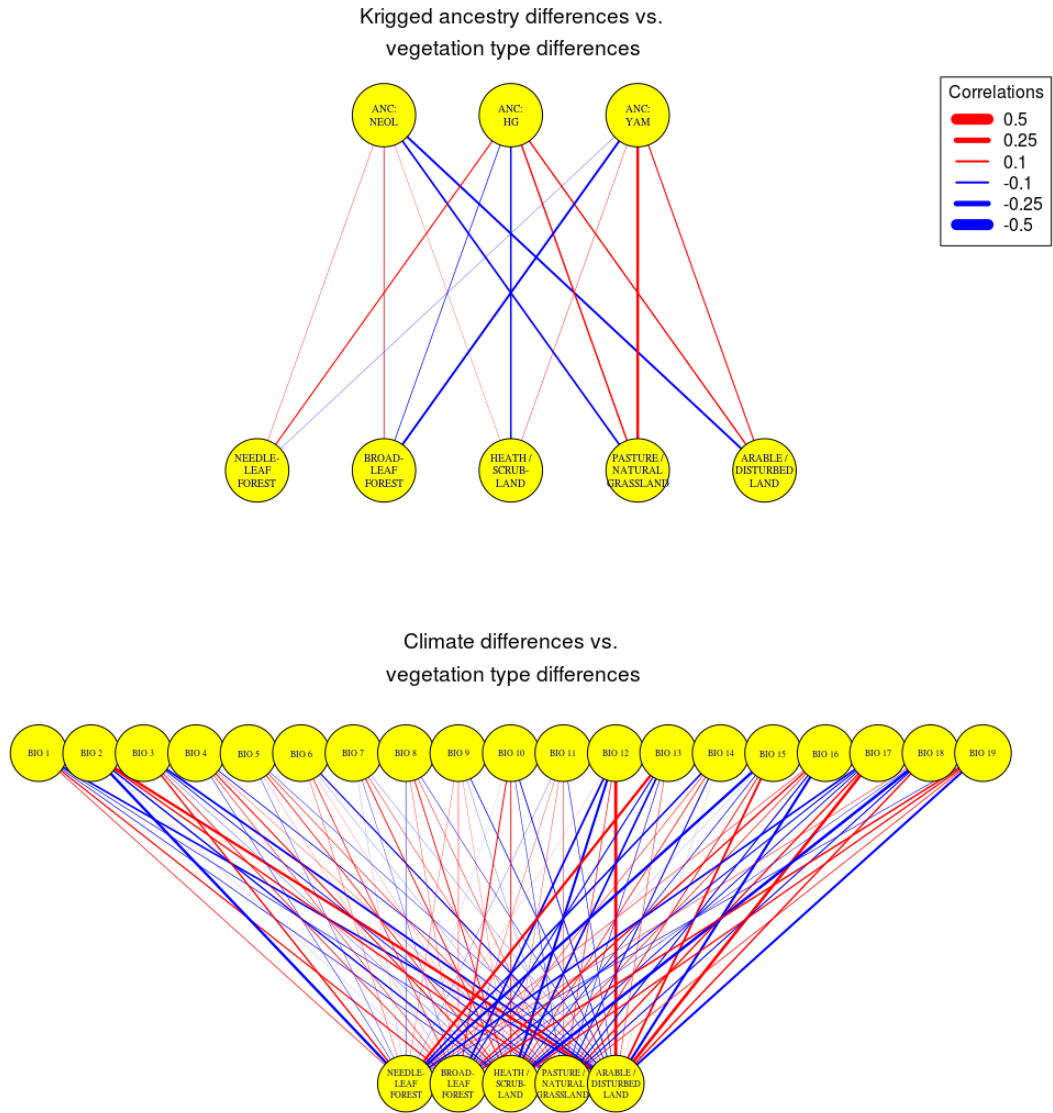
Figure S9: Correlations between differences in kriged ancestry and vegetation type at each time slice (top) or between differences in climate variables and in vegetation type at each time slice (bottom), computed across time and space. ANC = ancestry. ANC:NEOL = neolithic farmer ancestry. ANC:HG = hunter-gatherer ancestry. ANC:YAM = Yamnaya steppe ancestry. The climate variables follow the WorldClim nomenclature. BIO1 = Annual Mean Temperature. BIO2 = Mean Diurnal Range (Mean of monthly (max temp - min temp)). BIO3 = Isothermality (BIO2/BIO7) (* 100). BIO4 = Temperature Seasonality (standard deviation *100). BIO5 = Max Temperature of Warmest Month. BIO6 = Min Temperature of Coldest Month. BIO7 = Temperature Annual Range (BIO5-BIO6). BIO8 = Mean Temperature of Wettest Quarter. BIO9 = Mean Temperature of Driest Quarter. BIO10 = Mean Temperature of Warmest Quarter. BIO11 = Mean Temperature of Coldest Quarter. BIO12 = Annual Precipitation. BIO13 = Precipitation of Wettest Month. BIO14 = Precipitation of Driest Month. BIO15 = Precipitation Seasonality (Coefficient of Variation). BIO16 = Precipitation of Wettest Quarter. BIO17 = Precipitation of Driest Quarter. BIO18 = Precipitation of Warmest Quarter. BIO19 = Precipitation of Coldest Quarter.
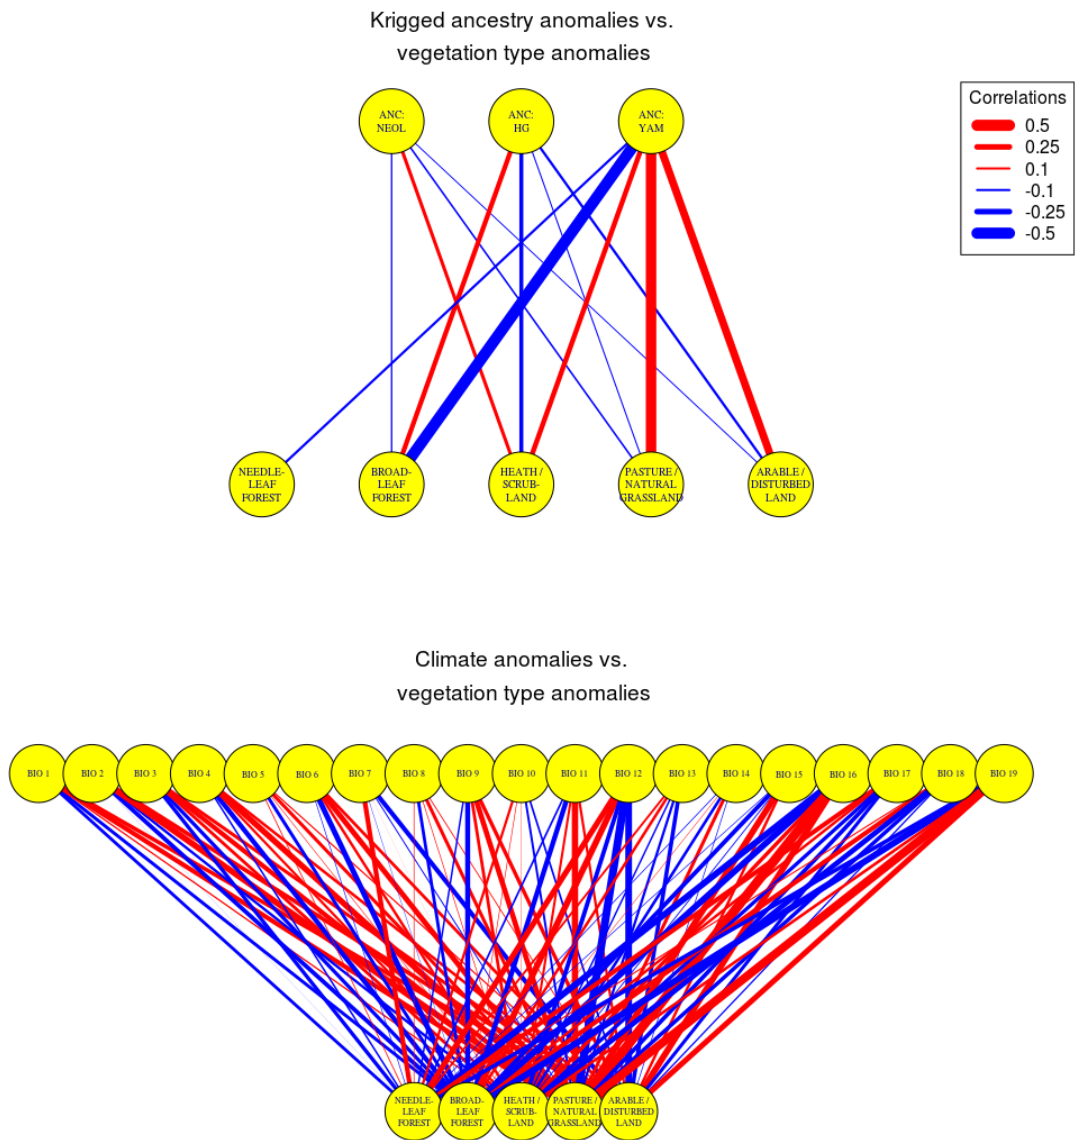
7

Figure S10: Correlations between anomalies, defined as the values of a variable in each past time slice before the present, after subtracting the value of that variable at the same location in the present. Top: Correlations between kriged ancestry anomalies and vegetation type anomalies. Bottom: Correlations between climate anomalies and vegetation type anomalies. ANC = ancestry. NEOL = neolithic farmer. HG = hunter-gatherer. YAM = Yamnaya. The climate variables follow the WorldClim nomenclature. BIO1 = Annual Mean Temperature. BIO2 = Mean Diurnal Range (Mean of monthly (max temp - min temp)). BIO3 = Isothermality (BIO2/BIO7) (* 100). BIO4 = Temperature Seasonality (standard deviation *100). BIO5 = Max Temperature of Warmest Month. BIO6 = Min Temperature of Coldest Month. BIO7 = Temperature Annual Range (BIO5-BIO6). BIO8 = Mean Temperature of Wettest Quarter. BIO9 = Mean Temperature of Driest Quarter. BIO10 = Mean Temperature of Warmest Quarter. BIO11 = Mean Temperature of Coldest Quarter. BIO12 = Annual Precipitation. BIO13 = Precipitation of Wettest Month. BIO14 = Precipitation of Driest Month. BIO15 = Precipitation Seasonality (Coefficient of Variation). BIO16 = Precipitation of Wettest Quarter. BIO17 = Precipitation of Driest Quarter. BIO18 = Precipitation of Warmest Quarter. BIO19 = Precipitation of Coldest Quarter.
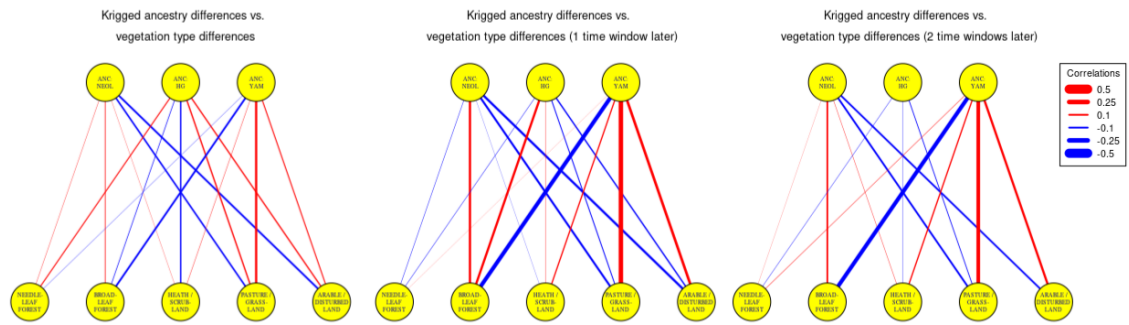
Figure S11: Correlations between differences in kriged ancestry in each time window and differences in vegetation in the same time window (left panel) one time window into the future (middle panel) or two time windows into the future (right panel). ANC = ancestry. NEOL = neolithic farmer. HG = hunter-gatherer. YAM = Yamnaya.
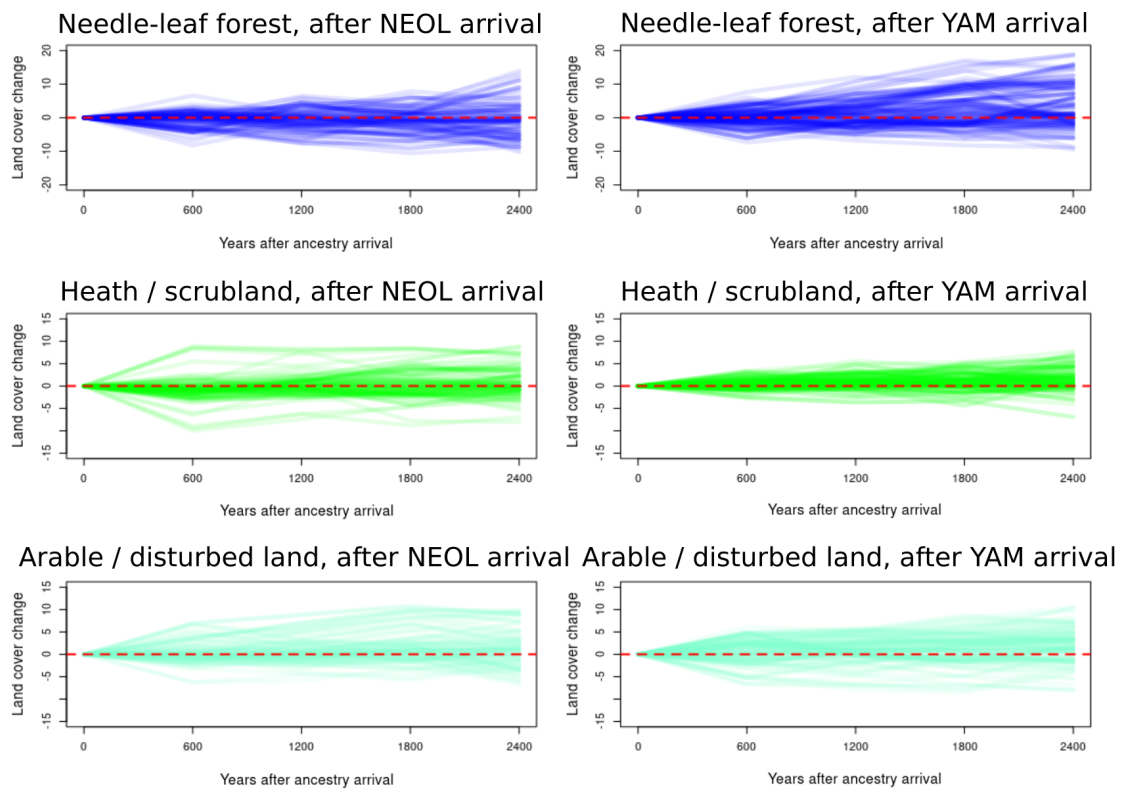

Figure S12: Change in Needle-leaf forest, heath and arable land cover composition after the arrival (first time there is > 50% ancestry) in each spatial grid point of YAM and NEOL ancestry. Each line corresponds to the post-arrival progression of a different spatial grid point.
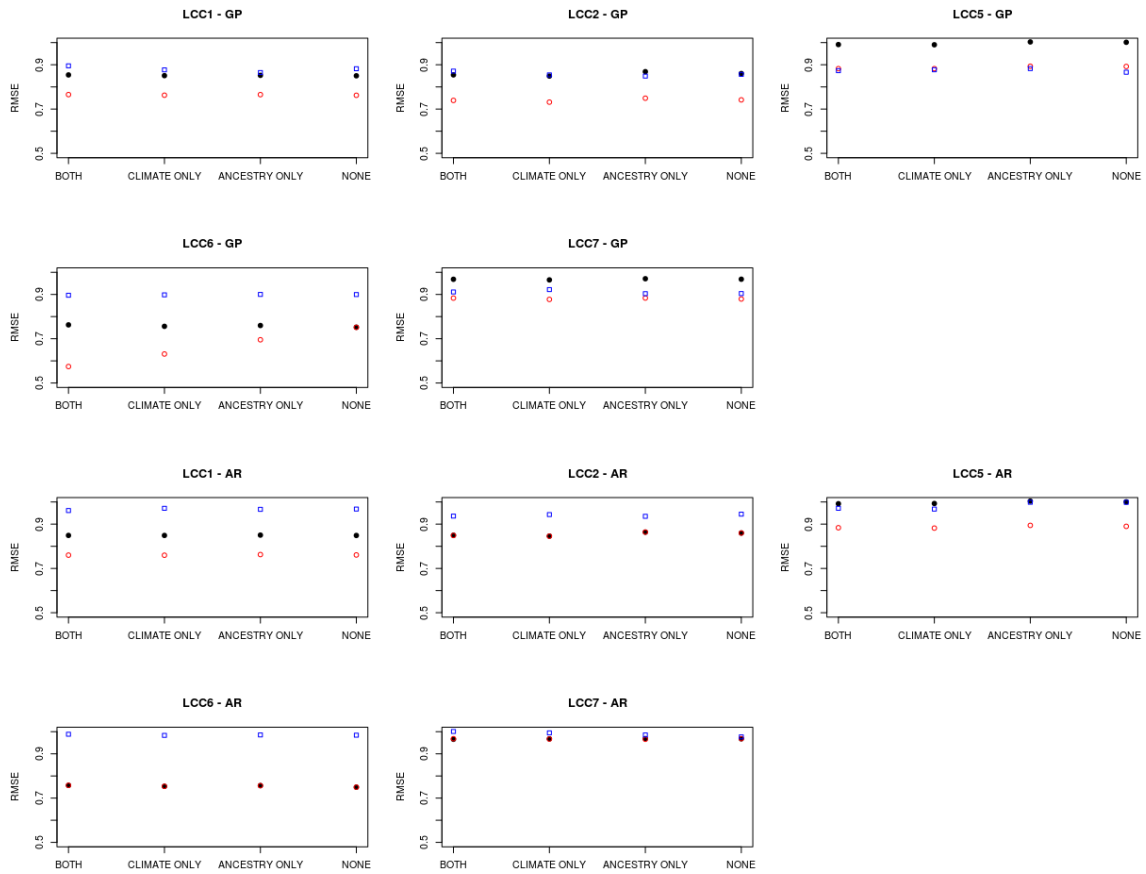
Figure S13: Root mean squared error for each fitted spatiotemporal model for paleovegetation anomalies. Black filled circles: using a fixed spatial decay parameter of the Matérn correlation function (default value in spTimer). Red circles: using a discrete uniform prior distribution for the spatial decay parameter (with default hyperparameters in spTimer). Blue squares: using a Gamma prior distribution for the spatial decay parameter (with default hyperparameters in spTimer). GP = Gaussian process model. AR = autoregressive model. BOTH = model incorporating both climate and ancestry explanatory variables. CLIMATE ONLY = model incorporating climate explanatory variables. CLIMATE ONLY = model incorporating ancestry explanatory variables. NONE = model incorporating neither ancestry nor climate explanatory variables.
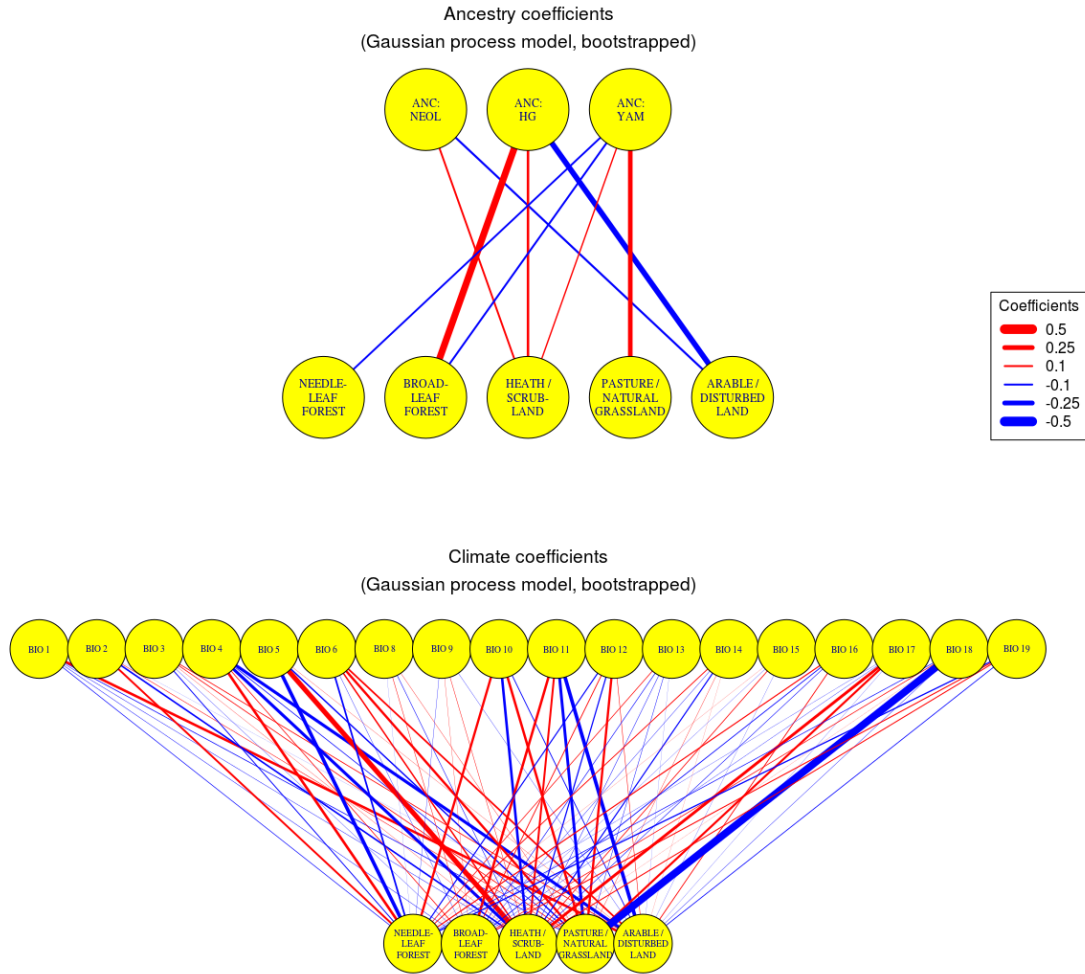
Figure S14: Spatiotemporal Gaussian process model for paleovegetation anomalies, using kriged ancestry anomalies and anomalies from simulation-based paleoclimate reconstructions as explanatory variables. Here, we used bootstrapping to account for uncertainty in the ancestry kriging estimation. The depicted lines denote the mean coefficients from the bootstrap distribution. Coefficients whose bootstrap distribution has a 95% central probability mass interval that spans the value of 0 are not depicted. ANC = ancestry. NEOL = neolithic farmer. HG = hunter-gatherer. YAM = Yamnaya. CLIM = Climate. The climate variables follow the WorldClim nomenclature.
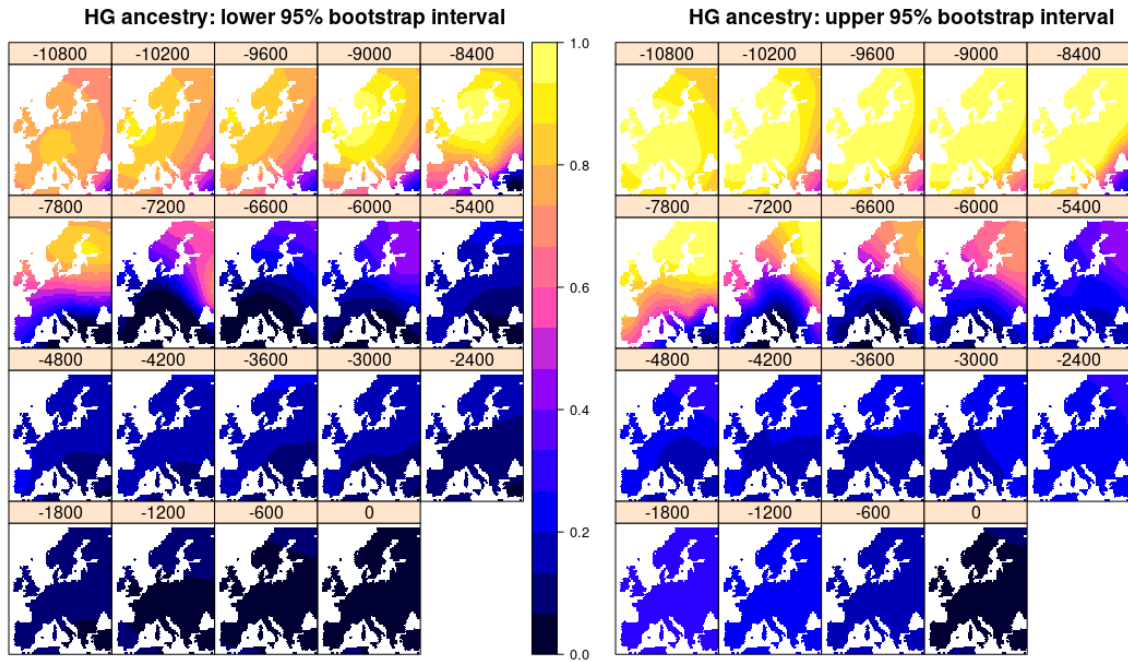
Figure S15: Lower and upper 95% quantiles of the bootstrap distribution of the spatiotemporal kriging of HG ancestry.
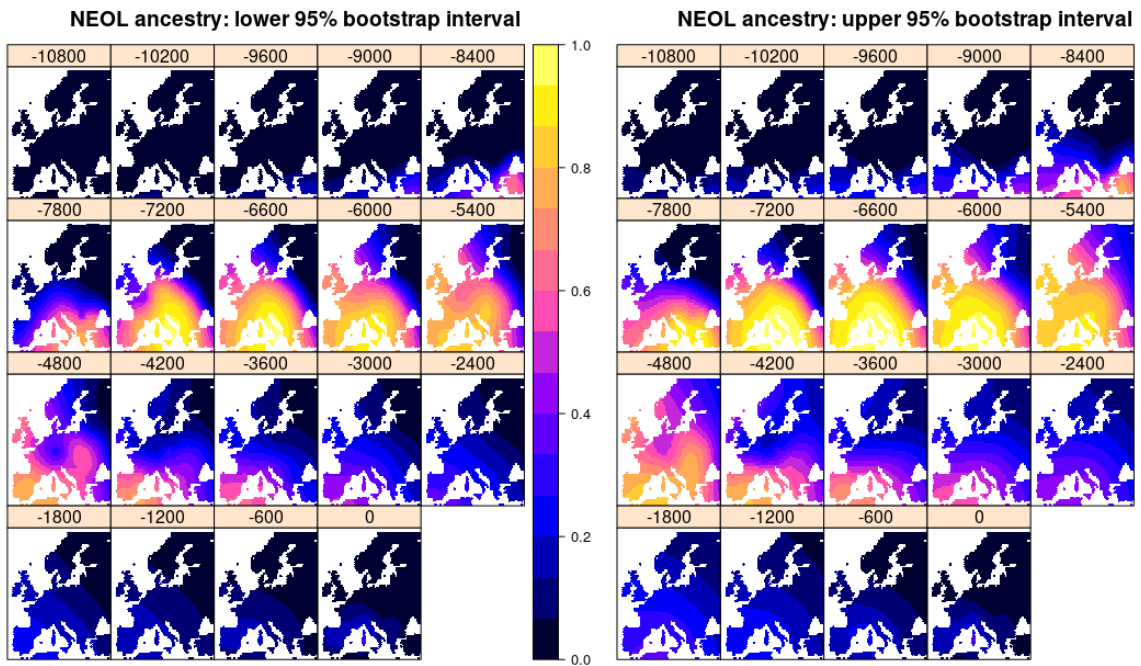

Figure S16: Lower and upper 95% quantiles of the bootstrap distribution of the spatiotemporal kriging of NEOL ancestry.
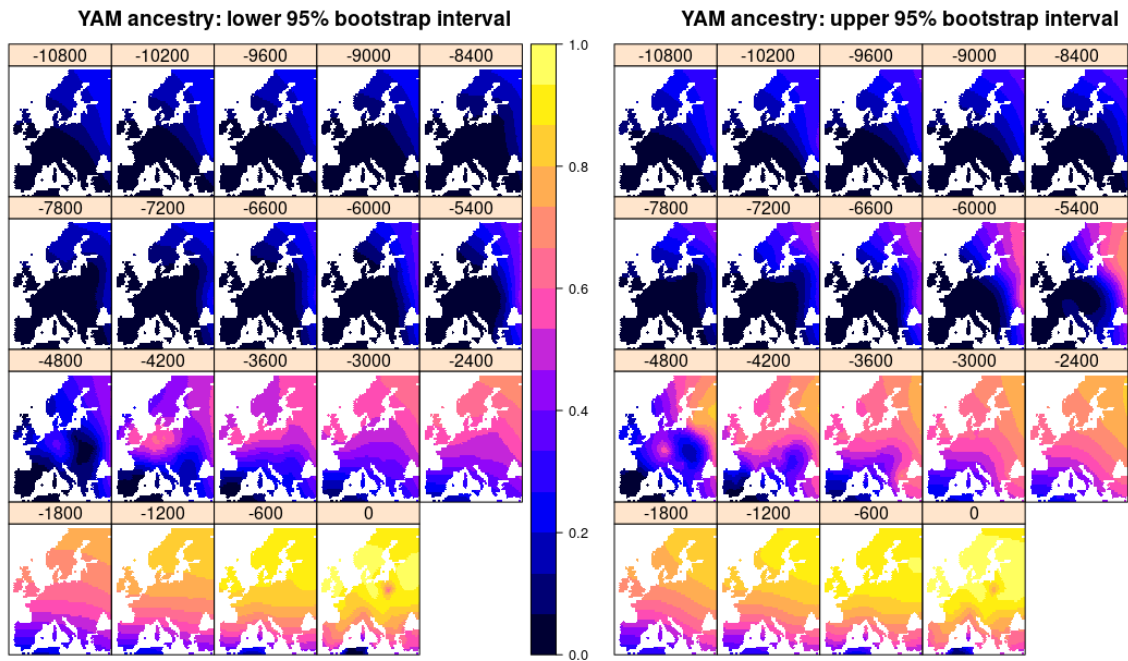
Figure S17: Lower and upper 95% quantiles of the bootstrap distribution of the spatiotemporal kriging of YAM ancestry.
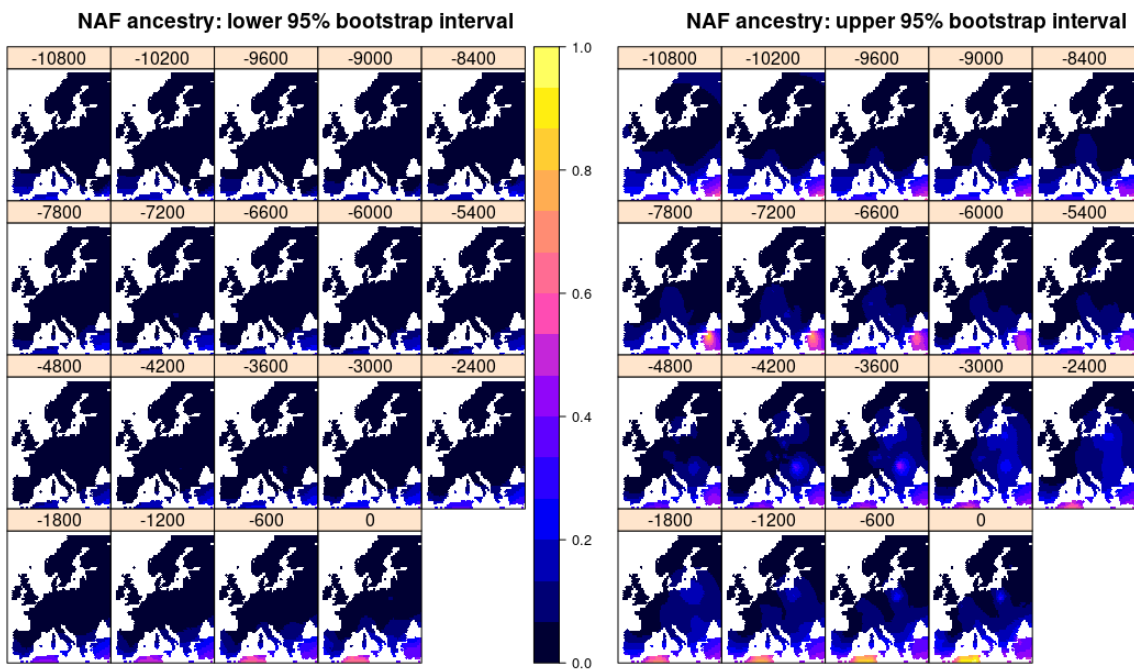

Figure S18: Lower and upper 95% quantiles of the bootstrap distribution of the spatiotemporal kriging of NAF ancestry.

Figure S19: Robustness of kriging to missing data across space. We divided our map into 16 different square sectors. We then computed, for each sector, the mean absolute error (MAE) of the kriged ancestry at the nearest spatiotemporal grid point to each genome within the sector, relative to the true ancestry of that genome. In the left panel, the kriging was performed while including the genome in the sector. In the right panel, the kriging was performed after excluding all genomes within the sector.
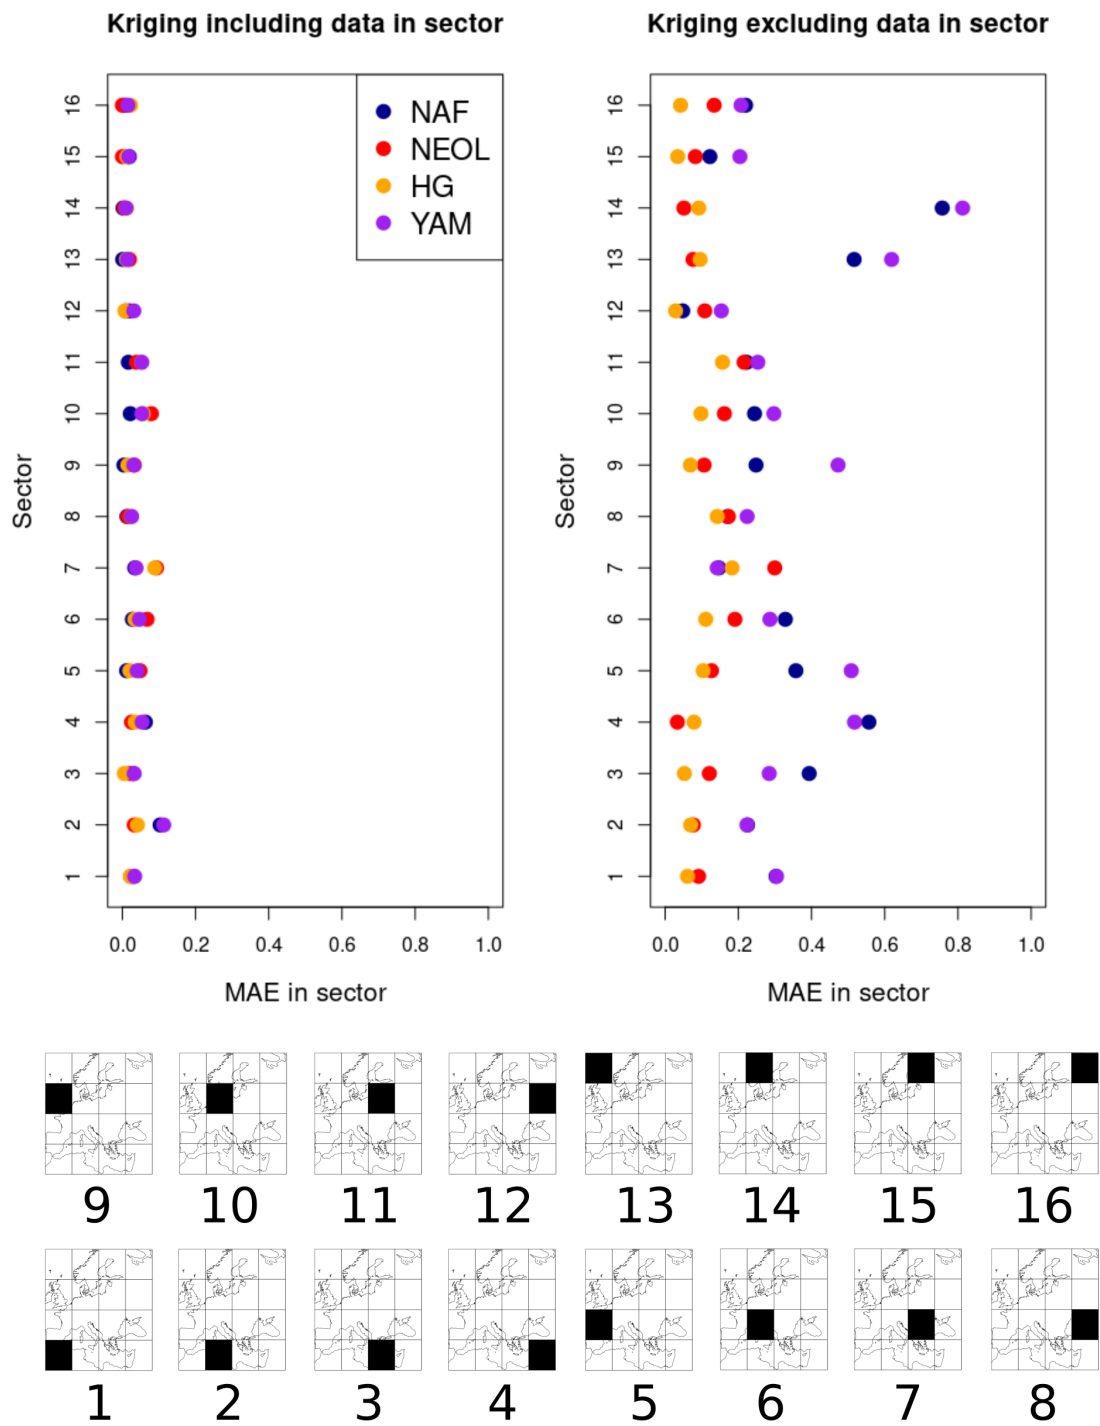
Figure S20: Robustness of kriging to missing data across time. We divided our timeline into 10 different periods. We then computed, for each period, the mean absolute error (MAE) of the kriged ancestry at the nearest spatiotemporal grid point to each genome within the period, relative to the true ancestry of that genome. In the left panel, the kriging was performed while including the genome in the period. In the right panel, the kriging was performed after excluding all genomes within the period.
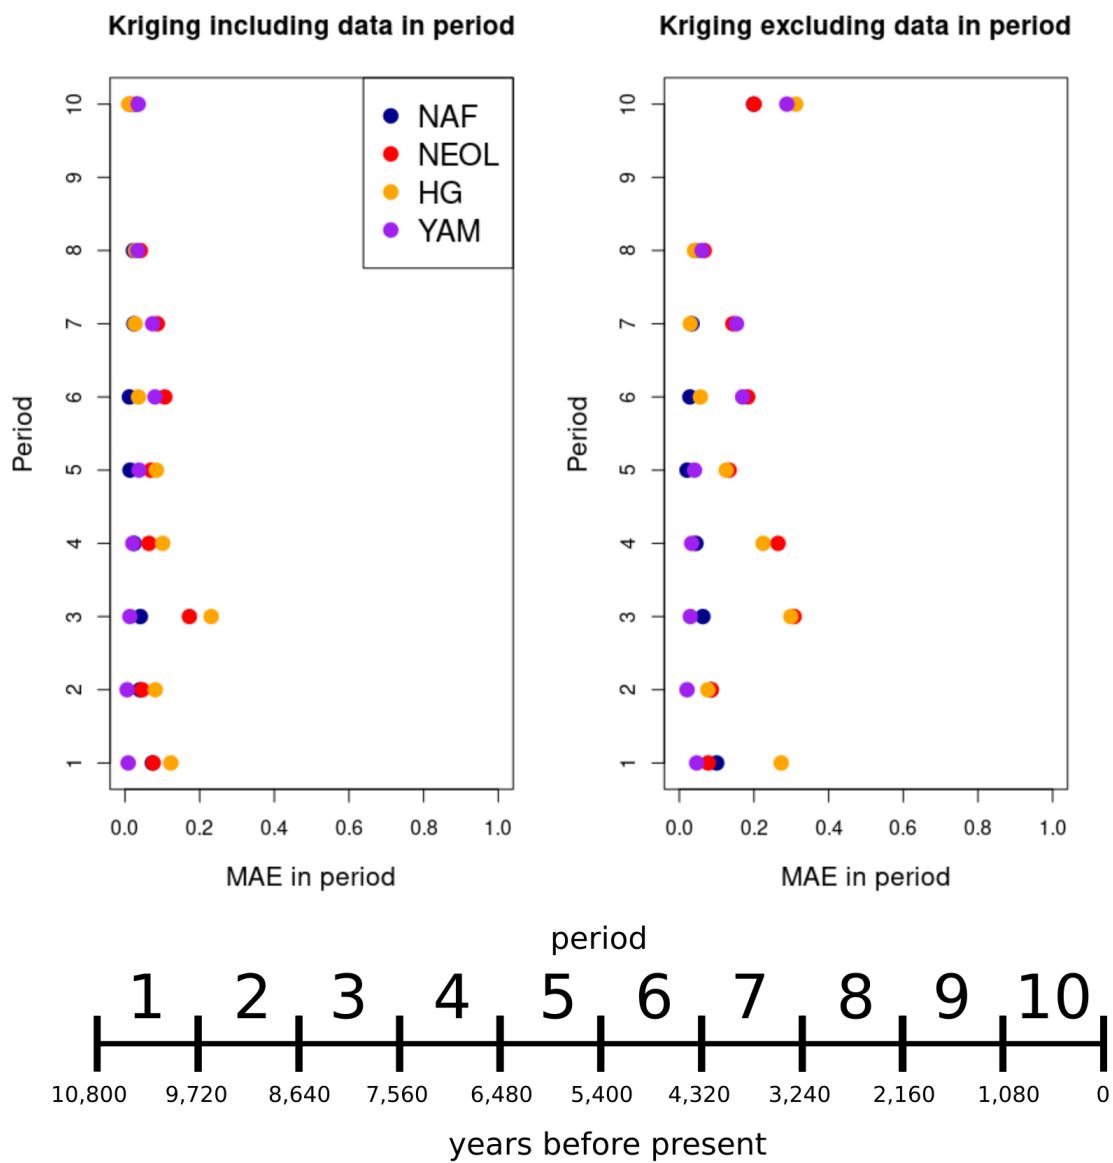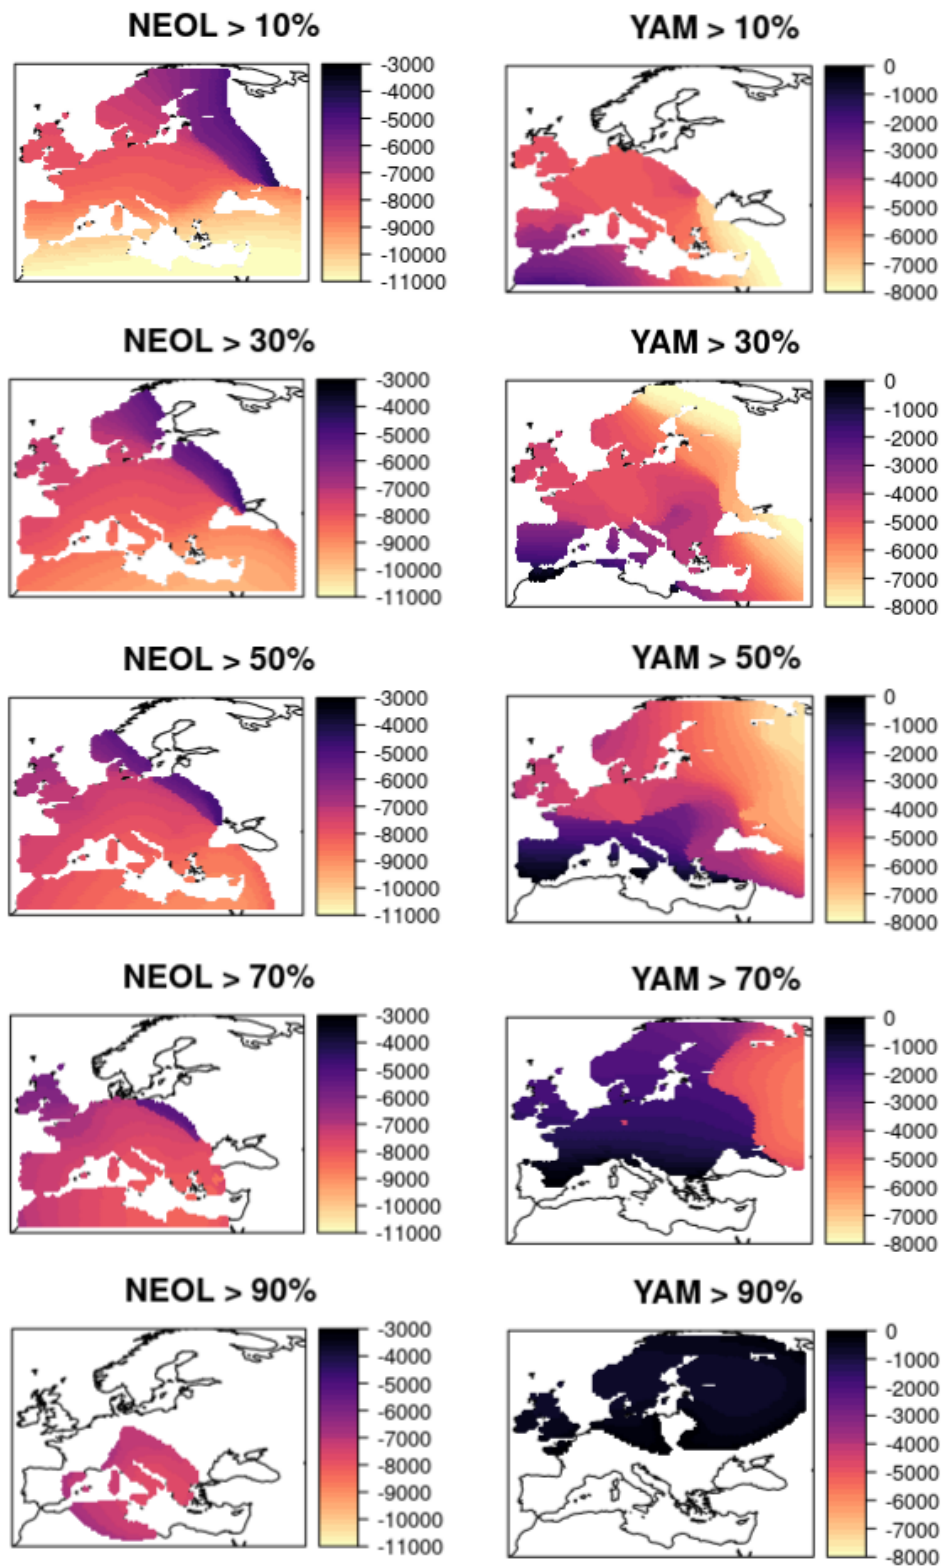
Figure S21: First-arrival maps for the spread of NEOL (left column) and YAM (right column) ancestries. First arrival is defined as the first time slice in which a grid point has more than a particular ancestry. We chose different minimum ancestry cutoffs, ranging from 10% to 90%.

Supplementary Tables

Table S1: [AVAILABLE ONLINE] Metadata for all ancient and present day genome sequences used, including regions and countries of sampling, three types of archaeological groupings (groupA, groupB and groupC), latitude, longitude and age in years before present.

MAP	Estimated NEOL speed	n(NEOL)	Estimated YAM speed	n(YAM)
10%	1.7 km/year (CI: 1.5-2.1)	285	1.3 km/year (CI: 1.2-1.5)	409
20%	1.7 km/year (CI: 1.5-2.1)	281	1.4 km/year (CI: 1.3-1.6)	378
30%	1.7 km/year (CI: 1.5-2)	279	2 km/year (CI: 1.7-2.3)	337
40%	1.7 km/year (CI: 1.5-2)	278	3.3 km/year (CI: 2.8-3.9)	301
50%	1.8 km/year (CI: 1.6-2.1)	274	4.2 km/year (CI: 3.5-5.2)	239
60%	1.9 km/year (CI: 1.6-2.3)	264	5.3 km/year (CI: 4-7.9)	116
70%	1.9 km/year (CI: 1.6-2.2)	255	6.8 km/year (CI: 4.1-18)	30
80%	1.8 km/year (CI: 1.6-2.2)	211	NA	13
90%	4.8 km/year (CI: 2.7-21)	65	NA	0

Table S2: Estimated speed of the YAM and NEOL ancestry expansion based on ancient genomes that have more than a minimum proportion of each ancestry, for various choices of minimum proportion. These were obtained from an RMA regression of the point of origin against the time of sampling. n(NEOL) and n(YAM) are the number of genomes used in each analysis. MAP = minimum ancestry proportion. CI = 95% confidence interval.

Origin	Estimated speed	95% confidence interval	Correlation coefficient (r)
49.48N, 43.66E	4.2 km/year	3.5-5.2 km/year	-0.55
43N, 44E	4.3 km/year	3.6-5.3 km/year	-0.543
48.4N, 31.4E	3.5 km/year	2.9-4.3 km/year	-0.534
50.7N, 59.1E	4.3 km/year	3.5-5.2 km/year	-0.546
54.1N, 43.1E	3.9 km/year	3.3-4.9 km/year	-0.556

Table S3: Estimated speed of the YAM ancestry expansion based on the genomes that have more than 50% of this ancestry, obtained from an RMA regression of the point of origin against the time of sampling. We tried different points of origin at the center and extremes of the Yamnaya historical range. We also list the correlation coefficients obtained from the RMA analysis.

Correlation	Raw	Differences	Anomalies
NEOL vs. Needle-leaf forest	-0.473	0.0182	0.000753
NEOL vs. Broad-leaf forest	-0.153	0.0392	-0.065
NEOL vs. Heath / scrubland	0.362	0.0135	0.152
NEOL vs. Pasture / natural grassland	0.35	-0.0859	-0.0777
NEOL vs. Arable / disturbed land	0.261	-0.0978	-0.0472
HG vs. Needle-leaf forest	0.176	0.0566	-0.000407
HG vs. Broad-leaf forest	0.266	-0.0356	0.212
HG vs. Heath / scrubland	-0.175	-0.0762	-0.176
HG vs. Pasture / natural grassland	-0.172	0.0723	-0.0527
HG vs. Arable / disturbed land	-0.28	0.0688	-0.113
YAM vs. Needle-leaf forest	0.439	-0.0148	-0.109
YAM vs. Broad-leaf forest	0.0806	-0.098	-0.503
YAM vs. Heath / scrubland	-0.279	0.0208	0.221
YAM vs. Pasture / natural grassland	-0.352	0.12	0.467
YAM vs. Arable / disturbed land	-0.22	0.0558	0.343

Table S4: Correlations between kriged ancestries and paleovegetation variables. Raw = correlation of raw variables corresponding to each spatiotemporal grid point. Differences = correlation of differences in raw variables at the same spatial grid point, between consecutive time slices. Anomalies = correlation in anomalies in variables, computed by subtracting the raw variable from its present-day value at the same spatial grid point.

Model	GF (F)	GF (U)	GF (G)	PMCC (F)	PMCC (U)	PMCC (G)
LCC1 - GP	2138.51	1662.44	2167.07	3178.34	**2508.52**	3465.82
LCC1 - AR	2135.14	2342.86	2464.87	2419.45	2237156.92	2835.11
LCC2 - GP	1549.44	1113.71	1791.07	2161.70	**1623.55**	2613.91
LCC2 - AR	1548.76	1547.14	2036.32	1706.15	1709.85	2293.76
LCC5 - GP	1884.91	1514.57	2006.90	3270.48	2635.82	3680.99
LCC5 - AR	1881.86	1505.86	2523.08	2191.77	**1724.91**	100472.82
LCC6 - GP	1652.28	849.61	1897.16	1930.61	**1123.51**	2247.83
LCC6 - AR	1644.27	1644.98	2327.76	1721.44	2531.40	2434.21
LCC7 - GP	1789.41	1393.14	1925.59	2701.84	2145.55	3030.07
LCC7 - AR	1784.27	1784.73	2370.32	**1929.59**	1929.61	2575.64

Table S5: Goodness of fit (GF) score and predictive model choice criteria (PMCC) (Bakar, Sahu, et al., 2015; Gelfand and Ghosh, 1998) for the Gaussian process (GP) and autoregressive (AR) models . F = Fixed decay parameter at 0.01. U = Uniform(0.01,0.02) prior for spatial decay parameter. G = Gamma(2,1) prior for spatial decay parameter. The lowest PMCC scores for each vegetation type are in bold.

Coefficient	Bootstrap median	2.5% quantile	97.5% quantile
NEOL → Needle-leaf forest	-0.0903	-0.153	0.0127
HG → Needle-leaf forest	-0.171	-0.284	0.0267
YAM → Needle-leaf forest	-0.101	-0.134	-0.0606
NEOL → Broad-leaf forest	0.00774	-0.0752	0.0788
HG → Broad-leaf forest	0.358	0.246	0.436
YAM → Broad-leaf forest	-0.0981	-0.16	-0.049
NEOL → Heath / scrubland	0.0963	0.0122	0.188
HG → Heath / scrubland	0.134	0.0183	0.278
YAM → Heath / scrubland	0.0719	0.0172	0.136
NEOL → Pasture / natural grassland	0.031	-0.0995	0.0919
HG → Pasture / natural grassland	-0.108	-0.321	0.00488
YAM → Pasture / natural grassland	0.21	0.148	0.272
NEOL → Arable / disturbed land	-0.103	-0.252	-0.0244
HG → Arable / disturbed land	-0.271	-0.554	-0.141
YAM → Arable / disturbed land	-0.023	-0.105	0.068

Table S6: Coefficients of the spatiotemporal autoregressive model, using the PaleoClim (Brown et al., 2018) simulation-based paleoclimate variables as covariates, with 95% bootstrap-based confidence intervals. Coefficients whose 95% confidence intervals do not overlap with 0 are in bold.

Movie S1 (separate file): Animation depicting the change in HG ancestry in Europe during the Holocene.

Movie S2 (separate file): Animation depicting the change in NEOL ancestry in Europe during the Holocene.

Movie S3 (separate file): Animation depicting the change in YAM ancestry in Europe during the Holocene.