

Fast principal components analysis reveals convergent evolution of ADH1B gene in Europe and East Asia

Kevin J. Galinsky^{1,2}, Gaurav Bhatia^{2,3}, Po-Ru Loh^{2,3}, Stoyan Georgiev⁴, Sayan Mukherjee⁵,
Nick J. Patterson^{2,*}, Alkes L. Price^{1,2,3,*}

¹Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, USA.

²Program in Medical and Population Genetics, Broad Institute of MIT and Harvard, Cambridge, USA. ³Department of Epidemiology, Harvard T.H. Chan School of Public Health, Boston, USA. ⁴Google, Palo Alto, CA. ⁵Departments of Statistical Science, Computer Science, and Mathematics, Duke University, Durham, NC.

Correspondence should be addressed to K.J.G. (galinsky@fas.harvard.edu), N.J.P. (nickp@broadinstitute.org) or A.L.P. (aprice@hsph.harvard.edu). *These authors contributed equally to this work.

Abstract

Searching for genetic variants with unusual differentiation between subpopulations is an established approach for identifying signals of natural selection. However, existing methods generally require discrete subpopulations. We introduce a method that infers selection using principal components (PCs) by identifying variants whose differentiation along top PCs is significantly greater than the null distribution of genetic drift. To enable the application of this method to large data sets, we developed the FastPCA software, which employs recent advances in random matrix theory to accurately approximate top PCs while reducing time and memory cost from quadratic to linear in the number of individuals, a computational improvement of many orders of magnitude. We apply FastPCA to a cohort of 54,734 European Americans, identifying 5 distinct subpopulations spanning the top 4 PCs. Using the PC-based test for natural selection, we replicate previously known selected loci and identify three new

genome-wide significant signals of selection, including selection in Europeans at the ADH1B gene. The coding variant rs1229984*T has previously been associated to a decreased risk of alcoholism and shown to be under selection in East Asians; we show that it is a rare example of independent evolution on two continents. We also detect new selection signals at IGFBP3 and IGH, which have also previously been associated to human disease.

Introduction

Searching for genetic variants with unusual differentiation between populations is an established approach for identifying signals of natural selection¹⁻⁶. We and others have employed this approach to identify signals of selection in a wide range of settings, informing our understanding of genes under evolutionary adaptation⁷⁻²⁴. Examples includes genes linked to lactase persistence^{9,11}, starch hydrolysis¹², fatty acid decomposition²⁴, red blood cell abundance¹⁷, hypoxia response¹⁸, alcoholism¹⁴, kidney disease²¹, malaria^{7,13,19,23}, HIV/AIDS¹⁶, autoimmune disease²⁰, cancer¹⁹, cystic fibrosis⁸ and hypertension²³. However, the signals of selection identified thus far may represent “only the tip of the iceberg²⁵”, implying that further research on selection will provide additional insights about human disease. Unlike extended haplotype homozygosity (EHH) or allele frequency spectrum based tests for selection, the population differentiation approach is able to detect older selection events and selection on standing variation^{1,3}. In addition, signals of selection detected using population differentiation can flag stratified genetic variants that are susceptible to false-positive associations in genome-wide association studies¹⁵.

Previous work on detecting selection using population differentiation has focused on methods that evaluate deviations from genome-wide patterns of genetic drift between discrete populations, such as Locus-Specific Branch Length (LSBL)⁶, Population Branch Statistic (PBS)¹⁷ and TreeSelect¹⁹. The population differentiation approach has greatest power when comparing very closely related

populations with very large sample size¹⁹. The increasing availability of very large population cohorts for genetic analysis provides strong prospects for analyzing subtle differences in ancestry in large sample sizes, but raises the challenge of how to select subpopulations to compare; a population cohort with a single continental ancestry may be better represented by continuous clines rather than discrete clusters^{26–28}, and/or may contain a large number of discrete subpopulations corresponding to a large number of possible population comparisons^{29,30}. Principal components analysis (PCA)^{26,31} offers an appealing alternative to model-based clustering methods^{32,33} for modeling human genetic diversity, and has been applied to infer population structure in many settings^{27,28,31,34–40}. One advantage of PCA is that results for top PCs are not sensitive to the number of PCs analyzed, whereas results of model-based clustering methods often vary with the number of clusters. Another advantage of PCA is its low computational cost, as top PCs can be inferred in time only linear in the number of samples by drawing upon recent advances in random matrix theory^{41–43}, implemented in the FastPCA software that we introduce here. We thus developed a test for selection that uses the SNP weights from PCA to calculate the differentiation of each locus along top PCs. Specifically, the squared correlation of each SNP to a PC, rescaled to account for genetic drift, follows a chi-square (1 d.o.f.) distribution under the null hypothesis of no selection. Our PC-based test produces a *p*-value at each locus and is able to detect novel signals at genome-wide significance, a key consideration in genome scans for selection¹⁹.

We ran FastPCA on 54,734 individuals of European descent from the Genetic Epidemiology Research on Adult Health and Aging (GERA) cohort; FastPCA required only 57 minutes of compute time and 2.6GB of RAM for this analysis, orders of magnitude better than any other publicly available software. We detected evidence of population structure along the top 4 PCs, which separated samples into several subpopulations. Using our PC-based test for selection, we replicate previously known selected loci *LCT* [MIM 603202], *HLA* [MIM 142800], *OCA2* [MIM 611409] and *IRF4* [MIM 601900] and identify three new signals of selection at *IGH* [MIM 147100], *IGFBP3* [MIM 146732] and *ADH1B* [MIM 103720]. The signal

in *ADH1B* at coding variant rs1229984 has previously been associated to alcoholism^{44–47} and shown to be under selection in East Asians^{14,46,48,49}; we show that it is a rare example of independent evolution on two continents^{11,12}.

Methods

Overview of methods

We first describe the FastPCA algorithm, which is an implementation of the *blanczos* method from Rokhlin *et al.*^{41–43}. As with our previous work on PCA^{26,31}, FastPCA makes use of existing computational literature and does not contain any new computational ideas; nonetheless, we anticipate that the software will be widely used, since to our knowledge it is the only publicly available software for computing top PCs on genetic data in linear time. The algorithm generalizes the method of power iteration⁵⁰, a technique to estimate the largest eigenvalue and corresponding eigenvector of a matrix. Multiplying a random vector by a square matrix projects that vector onto the eigenvectors of that matrix and then scales it according the respective eigenvalues of that matrix. After repeating, the projection along the eigenvector with the largest eigenvalue grows faster than the rest and the repeated matrix by vector product converges to this eigenvector. Additional eigenvectors can be found by repeating this process and orthogonalizing to previously-found PCs. The *blanczos* method improves on this method by initially estimating more PCs than ultimately desired. The original estimates are perturbed from the true PCs, but this missing variation is captured by estimating the extra PCs. The genotype matrix is then projected onto this set of eigenvectors, reducing its dimension while preserving the variation along the top PCs. Traditional PCA methods are applied to this reduced matrix to find accurate estimates of the top PCs of the original matrix.

We next describe our PC-based selection statistic, which generalizes a previous selection statistic developed for discrete populations¹⁹. We detect unusual allele frequency differences along inferred PCs

by making use of the fact that the squared correlation of each SNP to a PC, rescaled to account for genetic drift (Selection statistic), follows a chi-square (1 d.o.f.) distribution under the null hypothesis of no selection. We have released open-source software implementing the FastPCA algorithm and PC-based selection statistic (see Web Resources).

FastPCA algorithm

We are given an input $M \times N$ genotype matrix \mathbf{X} , where M is the number of SNPs and N is the number of individuals (e.g. each row is a SNP, each column is a sample). Each entry in this matrix takes its values from $\{0,1,2\}$ indicating the count of variant alleles for a sample at a SNP. From this matrix we can generate the normalized genomic matrix $\mathbf{Y}_{M \times N} = (\mathbf{y}_1^T, \mathbf{y}_2^T, \dots, \mathbf{y}_M^T)^T$ where each row \mathbf{y}_i has approximately mean 0 and variance 1 for SNPs in Hardy-Weinberg equilibrium.

$$\begin{aligned}\hat{p}_i &= \frac{\sum_{j=1}^N x_{ij}}{2N_i} = \frac{\mathbf{x}_i \mathbf{1}}{2\mathbf{1}^T \mathbf{1}} \\ y_{ij} &= \frac{x_{ij} - 2\hat{p}_i}{\sqrt{2\hat{p}_i(1 - \hat{p}_i)}} \\ \mathbf{y}_i &= (y_{i1}, y_{i2}, \dots, y_{iN}) = \frac{\mathbf{x}_i - 2\hat{p}_i \mathbf{1}}{\sqrt{2\hat{p}_i(1 - \hat{p}_i)}}\end{aligned}\tag{1}$$

Here, \mathbf{x}_i is the row vector of genotypes for SNP i and \mathbf{y}_i is the normalized row vector. x_{ij} and y_{ij} are the genotype/normalized genotype at SNP i for sample j . N_i is the number of valid genotypes at SNP i . All this is used to calculate \hat{p}_i , the sample allele frequency for SNP i , which is used to normalize the genotypes. In practice, the genotype matrix is normalized through the use of a lookup table mapping from genotypes (stored as 0, 1 or 2 copies of the alternate allele, or missing data) to normalized genotypes (using the above formula, with missing data having a normalized value of 0).

We are seeking the top K PCs for the normalized genomic matrix \mathbf{Y} . Traditional PCA algorithms compute the PCs by performing the eigendecomposition of the genetic relationship matrix ($GRM = \mathbf{Y}^T \mathbf{Y} / M$), a costly procedure which returns all the principal components. FastPCA, which makes use of recent advances in random matrix theory⁴¹⁻⁴³, speeds this process up by only approximating the top K PCs.

FastPCA is seeded with a random $N \times L$ matrix \mathbf{G}_0 composed of values drawn from a standard Gaussian distribution. L affects the accuracy of the result and L should be greater than K . For $K = 10$, $L = 20$ is a good choice. Then, for I iterations, we calculate $\mathbf{H}_i = \mathbf{Y} \times \mathbf{G}_i$ and $\mathbf{G}_{i+1} = \mathbf{Y}^T \times \mathbf{H}_i / M$, where the \mathbf{H}_i s are $M \times L$ matrices and \mathbf{G}_i s are $N \times L$ matrices like \mathbf{G}_0 . In simulated samples with discrete populations, $I = 3$ was sufficient, but in real datasets, $I = 10$ was found to provide accurate results.

After the iterative step completes, we stack the \mathbf{H}_i matrices to produce the matrix $\mathbf{H}_{M \times (I+1)L} = (\mathbf{H}_0, \mathbf{H}_1, \dots, \mathbf{H}_I)$, and the singular value decomposition of matrix \mathbf{H} is taken: $\mathbf{H} = \mathbf{U}_H \mathbf{\Sigma}_H \mathbf{V}_H^T$. \mathbf{U}_H is a low-rank approximation to the column-space of \mathbf{Y} with dimension $M \times (I + 1)L$, where $\mathbf{Y} \approx \mathbf{U}_H \mathbf{U}_H^T \mathbf{Y}$. \mathbf{Y} is then projected onto \mathbf{U}_H to produce $\mathbf{T}_{(I+1)L \times N} = \mathbf{U}_H \mathbf{Y}$. The SVD of $\mathbf{T} = \mathbf{U}_T \mathbf{\Sigma}_T \mathbf{V}_T^T$ can be computed efficiently and approximates the SVD of \mathbf{Y} since $\mathbf{Y} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \approx \mathbf{U}_H \mathbf{T} = \mathbf{U}_H \mathbf{U}_T \mathbf{\Sigma}_T \mathbf{V}_T^T$. For the PCA, we are only interested in the left K columns of \mathbf{V}_T and the first K entries along the diagonal of $\mathbf{\Sigma}_T$.

FastPCA runs in linear time and memory relative to M and N . There are $O(I)$ matrix multiplications where each multiplication takes $O(MNL)$ time. Then, the SVD of H takes $O(MI^2L^2)$ and the SVD of T takes $O(NI^2L^2)$ time. Taking I and L to be constants, the overall running time simplifies to $O(MN)$.

Selection statistic

We first consider the simple case of an ancestral population that split into two extant populations with genetic distance F_{ST} . We consider the allele frequencies at SNP i for the ancestral population (p_i) and the two extant populations (p_{i1} and p_{i2}). If there is no selection and SNPs are randomly ascertained,

$p_{i1} - p_{i2}$ has expectation 0 (because allele frequencies can drift either up or down in each population) and variance $2p_i(1 - p_i)F_{ST}$ ⁵¹. In the case where p_i is not close to 0 or 1 and F_{ST} is small, the distribution of this difference approximately follows a normal distribution:

$$\begin{aligned} E[p_{i1} - p_{i2}] &= 0 \\ \text{Var}[p_{i1} - p_{i2}] &= 2p_i(1 - p_i)F_{ST} \\ p_{i1} - p_{i2} &\sim N[0, 2p_i(1 - p_i)F_{ST}], \quad F_{ST} \ll 1, 0 \ll p_i \ll 1 \end{aligned} \quad (2)$$

In practice, we do not have access to either the ancestral allele frequency or the extant population allele frequencies. Instead, we have sample allele frequencies for the two extant populations, \hat{p}_{i1} and \hat{p}_{i2} . Assuming a large enough sample size from each population (N_1 and N_2) and that the true population allele frequency is not close to 0 or 1, these sample allele frequency estimates approximately follow a normal distribution with respect to the true allele frequencies. If we additionally assume that the ancestral allele frequency can be approximated by averaging the sample allele frequencies and that the true population allele frequencies are not that different, the sample allele frequency difference also follows a normal distribution^{13,15,19}:

$$\begin{aligned} \hat{p}_{i1} &\sim N\left[p_{i1}, \frac{p_{i1}(1 - p_{i1})}{2N_1}\right], \hat{p}_{i2} \sim N\left[p_{i2}, \frac{p_{i2}(1 - p_{i2})}{2N_2}\right], \quad N_1, N_2 \gg 0, 0 \ll p_{i1}, p_{i2} \ll 1 \\ D_i = \hat{p}_{i1} - \hat{p}_{i2} &\sim N\left[0, \sigma_D^2\right] = N\left[0, \hat{p}_i(1 - \hat{p}_i)\left(2F_{ST} + \frac{1}{2N_1} + \frac{1}{2N_2}\right)\right], \\ p_i \approx \hat{p}_i &= \frac{\hat{p}_{i1} + \hat{p}_{i2}}{2}, p_{i1} \approx p_{i2} \end{aligned} \quad (3)$$

Below, we build the intuition behind our PC-based statistic by rewriting the discrete-population statistic using vector notation, then extending this statistic to individuals with fractional ancestries, and then to continuous-valued PCs.

In the case with two discrete populations, we define a vector α where α_j indicates the ancestry in population 1 (e.g. $\alpha_j = 1$ if sample j is in population 1 and 0 if sample j is in population 2). D_i can be rewritten as:

$$\hat{p}_1 = \frac{\mathbf{x}_i \alpha}{2\mathbf{1}^T \alpha}, \quad \hat{p}_2 = \frac{\mathbf{x}_i (\mathbf{1} - \alpha)}{2\mathbf{1}^T (\mathbf{1} - \alpha)}, \quad D_i = \frac{\mathbf{x}_i \alpha}{2\mathbf{1}^T \alpha} - \frac{\mathbf{x}_i (\mathbf{1} - \alpha)}{2\mathbf{1}^T (\mathbf{1} - \alpha)} \quad (4)$$

If we run PCA on this sample, we would ideally get an eigenvector \mathbf{v} that has value v_1 for individuals in population 1 and $-v_2$ for individuals in population 2, where (since $\mathbf{v}^T \mathbf{1} = 0$, $\mathbf{v}^T \mathbf{v} = 1$)

$$v_q = \frac{1}{N_q} \sqrt{\frac{N_1 N_2}{N}} \quad (5)$$

In this case, D_i can be rewritten as:

$$D_i = \frac{1}{2} \sqrt{\frac{N_1 N_2}{N}} \mathbf{x}_i \mathbf{v} \quad (6)$$

In the limiting case where F_{ST} approaches 0, the statistic becomes:

$$\frac{D_i^2}{\sigma_D^2} = \frac{\frac{1}{4} \frac{N_1 N_2}{N} (\mathbf{x}_i \mathbf{v})^2}{\hat{p}_i (1 - \hat{p}_i) \left(\frac{1}{2N_1} + \frac{1}{2N_2} \right)} = \left[\left(\frac{\mathbf{x}_i - 2\hat{p}_i \mathbf{1}^T}{2\hat{p}_i (1 - \hat{p}_i)} \right) \mathbf{v} \right]^2 = [\mathbf{y}_i \mathbf{v}]^2 \quad (7)$$

Thus, the square of the SNP weight follows a chi-square 1-d.o.f. distribution in the case where $F_{ST} \rightarrow 0$.

In the case where $F_{ST} \neq 0$, then the scaling parameter has to be changed, but D_i still follows a normal distribution.

In the case with fractional ancestry ($\alpha_j \in [0, 1]$), \hat{p}_1 , \hat{p}_2 and D_i can still be estimated using equation (4).

The individual \hat{p}_q s will still asymptotically follow a normal distribution (because of the Lyapunov central

limit theorem⁵²), but will be correlated due to individuals with fractional ancestry contributing to both estimates. Thus, D_i will still follow a normal distribution, but the variance of equation (3) will not hold.

Now consider the case where we do not have fractional ancestries, but rather an eigenvector that separates individuals along some axis of variation. (We assume that extreme outlier individuals detected by PCA have been removed²⁶, as PCs dominated by such outliers may violate normality assumptions.) We can treat the eigenvector as a linear transformation of the ancestry vector:

$$\boldsymbol{\alpha} = \beta_0 + \beta_1 \boldsymbol{v} \quad (8)$$

Substituting these values into (4), we find:

$$D_i = \frac{\beta_1}{2N\beta_0(1-\beta_0)} \boldsymbol{x}_i \boldsymbol{v} \propto \boldsymbol{y}_i \boldsymbol{v} \quad (9)$$

Thus, our new selection statistic D_i is based on the dot product of the normalized genotypes and the eigenvector. Since the variance of D_i is not known, it will need to be rescaled in order to follow a $N(0,1^2)$ distribution.

If we are operating on the same set of SNPs that we used for PCA, then the rescaling of $\boldsymbol{y}_i \boldsymbol{v}$ is straightforward. Because PCA is the same as SVD, we see that:

$$\begin{aligned} \boldsymbol{Y} &= \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T \\ \boldsymbol{U} &= \boldsymbol{Y}\boldsymbol{V}\boldsymbol{\Sigma}^{-1} \end{aligned} \quad (10)$$

Here, \boldsymbol{V} contains the right singular vectors which are equivalent to the PCs, \boldsymbol{U} contains the left singular vectors which are rescaled SNP weights and $\boldsymbol{\Sigma}$ contains the singular values which are the square roots of the eigenvalues of the GRM. \boldsymbol{V} and \boldsymbol{U} are unitary, so the columns of \boldsymbol{U} are guaranteed to have a norm of 1. Multiplying \boldsymbol{U} by \sqrt{M} will then produce a properly normalized vector of differences $\boldsymbol{D} = (D_1, D_2, \dots, D_M)^T$. In other words:

$$\frac{\sqrt{M}}{\Sigma_k} \mathbf{y}_i \mathbf{v}_k \sim N(0,1)$$
$$\frac{M}{\Sigma_k^2} (\mathbf{y}_i \mathbf{v}_k)^2 \sim \chi_1^2$$
(11)

In the case of non-random SNP ascertainment and non-random choice of reference and variant allele, the expectation of D_i may be non-zero. However, if we randomly flip the reference and variant alleles in such a situation, the resulting principal components and values of D_i remain unchanged up to a factor of -1 and the expectation of D_i becomes 0. As a result, even if there are systematically positive or negative SNP loadings, D_i^2 still follows a chi-square 1-d.o.f distribution.

In the case where we are computing selection statistics on a different set of SNPs than the one for which we computed PCs, then the above property is not guaranteed to hold. Specifically, inflation can occur if SNPs with higher differentiation tend to have higher LD, which can occur as a consequence of true selection signals⁵³.

One assumption underlying the statistic is that the true minor allele frequency is not extremely small, otherwise the assumption of normality will not hold¹⁹. For this reason, the selection statistic was only computed for those SNPs containing minor allele frequency greater than 1% in our sample.

Simulation framework

Genotypes were simulated at M independent SNPs and N independent individuals in four steps:

1. The ancestral allele frequency (p_i) for a given SNP i was sampled from a *Uniform*(0.05,0.95) distribution.
2. Allele frequencies for Q populations ($\mathbf{P}_i = (p_{i1}, p_{i2}, \dots, p_{iQ})^T$) were generated by simulating random drift (see below).
3. Admixture (α_j) for individual j was sampled from a *Dirichlet*(\mathbf{a}) distribution.

4. Genotype g_{ij} was sampled from a $Binomial(2, \alpha_j^T \mathbf{P}_i)$ distribution.

Population allele frequencies were generated by simulating random drift in Q populations of fixed size N_e for τ generations and stored in $Q \times 1$ vector $\mathbf{P}_i = (p_{i1}, p_{i2}, \dots, p_{iQ})^T$. The number of alternate alleles z_{iqt} at SNP i in population q at generation t were sampled from a $Binomial(2N_e, p_{i,q,t-1})$ distribution, where p_{iq0} is the ancestral allele frequency p_i . The population allele frequency at this generation was then calculated as $p_{iqt} = \frac{z_{iqt}}{2N_e}$. For most simulations, population allele frequency simulations were run for $\tau = 200$ total generations and population size N_e was calculated for a target F_{ST} by using the formula $F_{ST} = -\log\left(1 - \frac{\tau}{2N_e}\right)^{19}$. For $F_{ST} \approx 0.1, 0.01$ and 0.001 , $N_e = 1k, 10k$ and $100k$ respectively. To detect the effect of population bottlenecks at the same level of F_{ST} , simulations were also run for $\tau = 20$ and $N_e = 100, 1k$ and $10k$, again producing populations with genetic distance $F_{ST} \approx 0.1, 0.01$ and 0.001 . Most simulations were run with two populations, but we also simulated 5 populations with a phylogenetic structure as follows. We set $N_e = 10k$ and $\tau = 200$ for populations 1 and 2, and $\tau = 180$ for an intermediary ancestral population of populations 3, 4 and 5, yielding allele frequency p_i^* . This was then fed back into the random drift model for an additional 20 generations for populations 3, 4 and 5. The pairwise genetic distance between populations 3, 4 and 5 is $F_{ST} \approx 0.001$ while the genetic distance between any other pair of populations is $F_{ST} \approx 0.01$.

We also considered simulations with admixed samples. In these simulations, the $Q \times 1$ population membership vector α_j for individual j was sampled from a $Dirichlet(\mathbf{a})$ distribution, where \mathbf{a} is a vector containing ancestry weightings. In the most simple case $\mathbf{a} = a\mathbf{1}$, where a is the admixture coefficient. For $a = 0$, this does not form a proper distribution and instead ancestry was selected by alternating individual ancestry between each of the populations. Increasing this coefficient increases admixture. When $a = 1$, this is effectively a uniform distribution and when $a > 1$, the mode of the distribution is one containing even admixture between all the populations.

The individual ancestries α_j make up the rows of ancestry matrix \mathbf{A} , which has dimension $N \times Q$.

Multiplying this ancestry matrix by the population allele frequency vector (\mathbf{P}_i), which (for a given SNP i) has length Q , generated an $N \times 1$ vector of allele frequencies for each individual ($\mathbf{P}'_i = \mathbf{A}\mathbf{P}_i$). Individual genotypes g_{ij} were generated from a *Binomial*(2, P'_{ij}) distribution.

To assess running time, the simulated datasets had $F_{ST} = 0.01$, $M = 100k$ SNPs and $N \approx \{1k, 1.5k, 2k, 3k, 5k, 7k, 10k, 15k, 20k, 30k, 50k, 70k, 100k\}$ individuals (since we used 6 populations of equal sample size, we rounded N to multiples of 6). Throughout this paper we report CPU time, but due to multi-threading present in the GSL⁵⁴ and OpenBLAS libraries, run time was about 60% of CPU time. FastPCA accuracy was assessed using $M = 50k$ SNPs and $N \approx 10k$ individuals at $F_{ST} = \{0.001, 0.002, \dots, 0.010\}$. Calibration and power of the selection statistic was assessed using 2 populations at $F_{ST} = \{0.1, 0.05, 0.02, 0.01, 0.005, 0.002, 0.001, 0.0005\}$, and also using 5 populations with the tree structure described above. We set $M = 60k$, the effective number of independent SNPs in genotype array data⁵⁵. When testing the power of the statistic, we wished to control the absolute difference in allele frequencies (D) between pairs of populations. For this purpose, SNPs under selection were generated in a similar manner as the above, except population allele frequencies were fixed at $p_{iq*} = 0.5 + \frac{D}{2}$ for one population and $p_{iq} = 0.5 - \frac{D}{2}$ for the remaining population(s); this approximates allele frequency differences under a population genetic selection model with strong selection in one population, because the magnitude of allele frequency differences caused by strong selection is much larger than the magnitude of allele frequency differences caused by genetic drift.

Assessing PC accuracy

Accuracy was assessed via the Mean of Explained Variances (MEV) of eigenvectors. Two different sets of K N -dimensional principal components each produce a K -dimensional column space. A metric for the performance of a PCA algorithm against some baseline is to see how much the column spaces overlap.

This is done by projecting the eigenvectors of one subspace onto the other and finding the mean lengths of the projected eigenvectors. If we have a reference set of PCs ($\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K$) against which we wish evaluate the performance a set of computed PCs ($\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_K$), then the performance calculation becomes:

$$MEV = K^{-1} \sum_{j=1}^K \sqrt{\sum_{k=1}^K (\mathbf{v}_k \cdot \mathbf{u}_j)^2} = K^{-1} \sum_{j=1}^K \|\mathbf{U}^T \mathbf{v}_k\| \quad (12)$$

Here, \mathbf{U} is a matrix whose column vectors are the PCs which we are testing. The test matrix can either be the result of another computation or the truth for a simulated sample. K eigenvectors can describe the population structure in a dataset with $K + 1$ populations. They can be constructed by first creating a vector $\mathbf{v}_k^* = (v_{k,1}^*, v_{k,2}^*, \dots, v_{k,N}^*)$ where $v_{k,j}^* = 1$ if individual j is in population k and 0 otherwise. The set of eigenvectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_K\}$ are constructed by taking K of these vectors, normalizing them to have mean 0, and scaling/orthogonalizing them via the Gram-Schmidt process.

GERA data set

The GERA dataset includes 62,318 individuals from Northern California typed on a European-specific 670,176-SNP array⁵⁶. This dataset underwent two levels of filtration: a quality control step to produce the QC set of SNPs used to detect natural selection, and a second step used to produce the LD-pruned set of SNPs for PCA.

For the QC step, individuals were filtered to remove those with missing sex information, individuals related according to the provided pedigree data or with observed genomic relatedness greater than 0.05 in the GRM⁵⁸ and individuals with less than 90% European ancestry as predicted by SNPweights⁵⁷ using a worldwide dataset containing European, African, and Asian ancestry. After filtering, 54,734 individuals

remained. Additionally, SNPs were initially filtered to remove non-autosomal SNPs, SNPs with minor allele frequency less than 1%, and SNPs with >1% missing data, leaving 608,981 SNPs.

The second stage of filtering removed SNPs that failed PLINK's Hardy-Weinberg Equilibrium test⁵⁸ with $p < 10^{-6}$, and performed LD-pruning using PLINK. Due to regions of long-range LD, LD persisted even after one filtering run. Multiple rounds of LD filtering were performed using an r^2 cutoff of 0.2 until additional rounds of LD filtering did not remove additional SNPs, leaving 162,335 SNPs.

FastPCA was run on the pruned set of 162,335 SNPs, while selection statistics were computed on the full set of 608,981 SNPs, prior to H-W filtering and LD-pruning. We note that many of the SNPs producing signals of selection generated significant H-W p -values (see Results - e.g. H-W $p = 1.37 \times 10^{-79}$ for LCT SNP rs6754311), which is an expected consequence of unusual population differentiation.

SNPweights⁵⁷ was used to predict fractional Northwest European, Southeast European, and Ashkenazi Jewish ancestry for each individual. For plotting purposes, percentage ancestry in each of these three populations was mapped to an integer in [0,255], which was then used for the RGB color value for that sample, so a NW sample would appear red, SE would appear green and AJ would appear blue.

PC Projection

POPRES⁵⁹ individuals were projected onto these PCs. The left singular vectors (\mathbf{U}) were generated by multiplying normalized genotypes for all SNPs in GERA (\mathbf{Y}_{GERA}) by the PCs (\mathbf{V}) and scaling by the singular values ($\mathbf{\Sigma}$), the number of SNPs used to calculate the PCs (M) and the number of SNPs used for projection (M_{GERA}): $\mathbf{U} = \mathbf{Y}_{GERA} \mathbf{V} \mathbf{\Sigma}^{-1} \sqrt{M/M_{GERA}}$. Projected PCs were then calculated by multiplying the corresponding set of SNPs in POPRES by these singular vectors and scaling again by the singular values: $\mathbf{V}_{POPRES} = \mathbf{Y}_{POPRES}^T \mathbf{U} \mathbf{\Sigma}^{-1}$. The projected individuals were overlaid on the PCA plot of GERA individuals and colored according to population membership and consistently with population assignment from SNPweights⁵⁷.

Results

FastPCA Simulations

We used simulated data to compare the running time and memory usage of FastPCA to three previous algorithms: smartpca^{26,31}, PLINK2-pca⁵⁸, and flashpca⁶⁰ (see Web Resources). We simulated genotype data from six populations with a star-shaped phylogeny using 100k SNPs (typical for real data after LD-pruning) and up to 100k individuals (see Methods). For each run, running time was capped at 100 hours and memory usage was capped at 40GB. The running time and memory usage of FastPCA scaled linearly with simulated dataset size (Figure 1), compared with quadratically or cubically for other methods. The computation became intractable at 50k-70k individuals for smartpca, PLINK2-pca and flashpca. The largest dataset, with 100k SNPs and 100k individuals, required only 56 minutes and 3.2GB of memory with FastPCA (Table S1). (We also note that shellfish (see Web Resources), a parallel PCA implementation, requires $O(MN^2 + N^3)$ and is not computationally tractable on large data sets, as previously demonstrated⁶⁰). Thus, FastPCA—unlike other publicly available software packages for analyzing genetic data—enables rapid principal components analysis without specialized computing facilities.

We next assessed the accuracy of FastPCA, using PLINK2-pca⁵⁸ as a benchmark. We used the same simulation framework as before, with 10k individuals (1,667k individuals per population) and 50k SNPs. We varied the divergence between populations, as quantified by F_{ST} ⁶¹. We assessed accuracy using the Mean of Explained Variances (MEV) of the 5 population structure PCs (see Methods). We determined that the results of FastPCA and PLINK-pca were virtually identical (Figure 2). This indicates that FastPCA performs comparably to standard PCA algorithms while running much faster.

PC-based Selection Statistic Simulations

We evaluated the calibration and power of the PC-based selection statistic. To evaluate calibration, we simulated 60k SNPs undergoing random drift with up to $N = 50k$ individuals from two populations differentiated by $F_{ST} = \{0.1, 0.01, 0.001\}$. At all values of N and F_{ST} , the proportion of truly null SNPs reported as significant was well-calibrated at p -value thresholds ranging from 10^{-1} to 10^{-5} . Similar results indicating appropriate calibration were obtained for simulations with admixture (Table S2), as expected since the drift model still applies in the case of admixture²⁷. The median of the selection statistic was slightly inflated at $F_{ST} = 0.1$ due to a deficiency in the tail (Figure S1 and Table S2), but well-calibrated at the small values of F_{ST} that correspond to our analyses of real data. The selection statistic in the presence of a population bottleneck performed identically to populations differentiated by the same F_{ST} level (Table S2). We also simulated five populations with a phylogenetic structure (see Methods) that mimics the population structure found in the GERA data (see below) and found that the statistic remained well-calibrated here as well (Figure S1 and Table S2).

We evaluated power using the same number of SNPs and samples but at $F_{ST} = \{0.1, 0.05, 0.02, 0.01, 0.005, 0.002, 0.001, 0.0005\}$ and using a separate set of SNPs under selection where the allele frequency between the two populations was varied ($|D| = |p_1 - p_2|$). The significance threshold was set to 8.3×10^{-7} based on 60K SNPs tested. There was no power to detect selection at $F_{ST} = 0.1$. We observed a phase-change in the power simulations that was sharper for smaller F_{ST} , where there was no power to detect selection below a specified allele frequency difference threshold, but there was complete power to detect selection at a slightly higher threshold (Figure 3a). We examined this effect in more depth using a range of samples sizes, and determined that the transition from no-power to complete-power was more sample size dependent at $F_{ST} = 0.001$ (Figure 3b) than at $F_{ST} = 0.01$ (Figure 3c), indicating that sample size is more important when analyzing more closely related populations. The PC-based selection statistic performed very similarly to the discrete-population

test of selection¹⁹ in the case of data from discrete subpopulations (Figure S2). We also assessed effect of admixture on power by sampling ancestry for individuals between the two populations using a $Beta(a, a)$ distribution. We determined that increasing the admixture parameter a (which reduces the variation in ancestry across samples) had a similar effect to reducing sample size (Figure S3).

Application of FastPCA to a European American Cohort

We ran FastPCA on the GERA cohort (see Web Resources), a large European American dataset containing 54,734 individuals and 162,335 SNPs after QC filtering and LD-pruning (see Methods). This computation took 57 minutes and 2.6GB of RAM. PC1 and PC2 separated individuals along the canonical Northwest European (NW), Southeast European (SE) and Ashkenazi Jewish (AJ) axes¹⁵, as indicated by labeling the individuals by predicted fractional ancestry from SNPweights⁵⁷ (Figure 4). These results are consistent with Banda *et al.* 2015⁵⁶ which also examined this dataset. PC3 and PC4 detected additional population structure within the NW population.

To further investigate this subtle structure, we projected POPRES individuals from throughout Europe⁵⁹ onto these PCs³¹ (see Methods). This analysis recapitulated the position of SE populations via the placement of the Italian individuals, and determined that PC3 and PC4 separate the NW individuals into Irish (IR), Eastern European (EE) and Northern European (NE) populations (Figure 5). This visual subpopulation clustering was confirmed via k-means clustering on the top 4 PCs, which consistently grouped the AJ, SE, NE, IR and EE populations separately (Figure S4. k-Means clustering confirms visually-observed subpopulations. Figure S4). We note that, in general, K PCs can cluster samples into $K+1$ subpopulations.

Application of PC-based Selection Statistic to a European American Cohort

For each of the top PCs, we computed our PC-based selection statistic for 608,981 non-LD-pruned SNPs (see Methods). The resulting Manhattan plots for PCs 1-4 are displayed in Figure 6 (QQ plots are

displayed in Figure S5). Analyses of PCs 5-10 indicated that these PCs do not represent true population structure (Figure S6), but are either dominated by a small number of long-range LD loci^{34,62,63} or correlated with the missing data rate across individuals. Selection statistics for PCs 1-4 exhibited little or no inflation, particularly after removing Table 1 regions (Table S3).

Genome-wide significant signals (listed in Table 1) included several known selection regions^{9,64-67} and novel signals at ADH1B, IGFBP3 and IGH (see below). Suggestive signals were observed at additional known selection regions^{66,68} (Table S4). After removing the regions in Table 1, rerunning FastPCA and recalculating selection statistics, all of these regions remained significant except for a region on chromosome 8 with a known chromosomal inversion^{34,62} (Figure S7, Table S5). Thus, the remaining regions are not due to PC artifacts caused by SNPs inside these regions. We also found that a significantly greater proportion of SNPs under selection failed Hardy-Weinberg equilibrium, although the converse is not true, indicating that signals of selection are not a result of H-W artifacts (Figure S8).

Detecting subtle signals of selection benefited from the large sample size, as subsampling the GERA data set at smaller sample sizes and recomputing PCs and selection statistics generally led to less significant signals (Table 2). We note that several suggestive selection signals, including signals at the known selected loci TLR1⁶⁶ [MIM 601194] and SLC45A2⁶⁸ [MIM 606202], are on the cusp of being significant and further increases to sample size may increase power to detect selection at suggestive loci.

We identified a genome-wide significant signal of selection at rs1229984, a coding SNP (Arg47His) in the ADH1B alcohol dehydrogenase gene (Table 1). The allele rs1229984*T has been shown to have a protective effect on alcoholism risk⁴⁴⁻⁴⁷ and to produce an REHH signal in East Asians^{14,46,48,49}, but was not previously known to be under selection in Europeans. (Previous studies noted the higher frequency of the rs1229984*T allele in western Asia compared to Europe, but indicated that selection or random drift were both plausible explanations^{69,70}.) We examined the allele frequency of the rs1229984*T allele in the five subpopulations: AJ, SE, NE, IR and EE (Table S7). We observed allele frequencies of 0.21 in AJ,

0.10 in SE, and 0.05 or lower in other subpopulations, consistent with the higher frequency of the rs1229984*T in western Asia. A comparison of NE to the remaining subpopulations using the discrete subpopulation selection statistic¹⁹ also produced a genome-wide significant signal after correcting for all hypotheses tested (Table S7); this is not an independent experiment, but indicates that this finding is not due to assay artifacts affecting PCs.

To further understand the selection at this locus, we examined the allele frequency of rs1229984*T in 1000 Genomes project⁷¹ populations (see Web Resources), along with the allele frequency of the regulatory SNP rs3811801 that may also have been a target of selection in Asian populations⁴⁶. The haplotype carrying rs3811801*A (and corresponding haplotype H7) was absent in populations outside of East Asia (Table S8). This indicates that if natural selection acted on this SNP in Asian populations, selection acted independently at this locus in Europeans. One possible explanation for these findings is that rs1229984 is an older SNP under selection in Europeans, while rs3811801 is a newer SNP under strong selection in Asian populations leading to the common haplotype found in those populations.

The IGFBP3 insulin-like growth factor-binding protein gene had two SNPs reaching genome-wide significance. Genetic variation in IGFBP3 has been associated with breast cancer⁷², height⁷³, blood pressure⁷⁴ and hypertension⁷⁵, although the published associated SNPs are not in LD with the two SNPs we detected. The IGH immunoglobulin heavy locus had one genome-wide-significant SNP and two suggestive SNPs with p -value $< 10^{-6}$ (Table 1). Genetic variation in IGH has been associated with multiple sclerosis⁷⁶, although the published associated SNPs are not in LD with the three SNPs we detected. The IGFBP3 and IGH SNPs each had substantially higher minor allele frequencies in Eastern Europeans, but were not genome-wide significant under the discrete subpopulation selection statistic¹⁹(Table S9 and Table S10). The existence of multiple SNPs at each of these loci with $p < 10^{-6}$ for the PC-based selection statistic suggests that these findings are not the result of assay artifacts.

Discussion

We have detected new, genome-wide significant signals of selection by applying a PC-based selection statistic to top PCs computed using FastPCA, a computationally efficient (linear-time and linear-memory) algorithm. Although mixed model association methods are increasingly appealing for conducting genetic association studies^{55,77}, we anticipate that PCA will continue to prove useful in population genetic studies, in characterizing population stratification when present in association studies, in supplementing mixed model association methods by including PCs as fixed effects in studies with extreme stratification, and in correcting for stratification in analyses of components of heritability^{78,79}. Our PC-based selection statistic extends previous statistics developed for discrete populations¹⁹. In contrast to previous work on detecting selection using PCs^{63,80} or using the spatial ancestry analysis (SPA) method⁸¹, our statistic is able to detect signals at genome-wide significance, a key consideration in genome scans for selection⁸². Our work demonstrates the advantages of comparing closely related populations in very large sample sizes to detect novel loci, whereas very recent studies applying related methods to smaller sample sizes detected genome-wide significant signals only at previously known loci^{83,84}. In particular, we detected genome-wide significant evidence of selection in Europeans at the ADH1B locus, which was previously reported to be under selection in East Asian populations^{14,46,48,49} using REHH⁵³ (which can only detect relatively recent signals and does not work on standing variation³). We also detected genome-wide significant evidence of selection at the disease-associated IGFBP3 and IGH loci. While the SNPs under selection at these loci are not in LD with the disease-associated SNPs identified in previous association studies, these genes are biologically important and there may be other phenotypes associated with the selected SNPs. Although we emphasize the importance of genome-wide significance, loci with suggestive signals of selection that do not reach genome-wide significance could potentially be used to increase the power of disease mapping⁸⁵.

We note that our work has several limitations. First, top PCs do not always reflect population structure, but may instead reflect assay artifacts⁸⁶ or regions of long-range LD³⁴; however, PCs 1-4 in GERA data reflect true population structure and not assay artifacts, because the PCs (and the signals of selection they detect) remained nearly unchanged after removing regions with significant signals of selection (Table 1) and rerunning PCA. Second, common variation may not provide a complete description of population structure, which may be different for rare variants⁸⁷; we note that based on analysis of real sequencing data with known structure, we recommend that LD-pruning and removal of singletons (but not all rare variants) be applied in data sets with pervasive LD and large numbers of rare variants (see Appendix). Third, our selection statistic is only capable of detecting that selection occurred, but not when or where it occurred; indeed, top PCs may not perfectly represent the geographic regions in which selection occurred, underscoring that interpretation of results can be a fundamental limitation of model-free methods. Fourth, our selection statistic performs best when allele frequencies vary linearly along a PC; the SPA method⁸¹ (see above) models allele frequency as a logistic function and is not constrained by this limitation. Despite these limitations, we anticipate that FastPCA and our PC-based selection statistic will prove valuable in analyzing the very large data sets of the future.

Appendix

Inferring ancestry from genetic data is a common problem in both population and medical genetic studies, and many methods exist to address it^{31,32,88}. Principal components analysis (PCA)³¹ has been shown to be effective at elucidating geographic structure from genetic data⁸⁹ and correcting for confounding due to population stratification in association mapping²⁶. These uses of PCA depend critically on its ability to separate genetically disparate subpopulations when analyzing data from commercial genotyping arrays. However, as high-throughput sequence data becomes more common, enabling ancestry inference from this new class of data is becoming increasingly relevant.

As sequence data contains more variants, and many more population-specific variants⁹⁰, it may be reasonable to expect that PCA applied to high-throughput sequence data will be substantially more effective than the corresponding analysis on genotype data. However, our results suggest the opposite. Specifically, PCA makes assumptions about marker independence that are violated by the pervasive linkage disequilibrium in sequence data. In addition, assumptions about genetic drift that are reasonable for common SNPs on genotyping arrays are less so when applied to the numerous rare variants in sequence data⁸⁷.

Methods

Principal Components Analysis (PCA) is generally applied to a genetic relationship matrix (GRM) that is computed as:

$$g_s = \frac{x_s - 2p_s}{\sqrt{2p_s(1 - p_s)}}$$
$$G = \sum_{s \in SNPs} g_s g_s^T$$

where x_s is a vector of genotypes for SNP s and p_s is the minor allele frequency of SNP s . We propose modifications to standard PCA to deal with two challenges that are present in sequence data but absent from genotype data: pervasive linkage disequilibrium, and rare variants. Specifically, we recommend that LD pruning be applied to sequence data and singleton variants be removed. While we evaluated more sophisticated approaches to handling these issues, they did not improve our results beyond these simpler approaches. Importantly, we recommend against a commonly used strategy of removing all low frequency of rare variants as these variants contain significant information for detecting population structure.

Linkage Disequilibrium

It is well known that application of PCA to regions of the genome containing long-range LD blocks can confound PCA's ability to separate disparate populations^{31,63}. As a result, these LD blocks are often simply excluded from analysis. However, in sequence data, many regions of the genome outside of previously identified long-range LD blocks contain sufficient LD to bias results. As a result, we examine three methods to deal with LD: (1) LD Pruning (2) LD Shrinkage⁶³ and (3) LD Regression^{31,912,9}.

LD Pruning is a commonly applied approach to removing correlated SNPs from a dataset. To produce a data set pruned for LD above a threshold T , one SNP of any pair of SNPs in LD ($r^2 > T$) is removed from the data.

LD Shrinkage is a more sophisticated method of correcting for LD proposed by (Zou et al. 2012)⁶³. In LD shrinkage, each SNP s is weighted by its LD to surrounding SNPs before inclusion in the genetic relationship matrix.

$$GLDg_s = \frac{x_s - 2p_s}{\sqrt{2p_s(1-p_s)}}$$

$$w_s = \frac{1}{1 + \sum_{t \in \text{window}(s)} r_{s,t}^2}$$

$$G = \sum_{s \in \text{SNPs}} g_s g_s^T$$

We note that $t \in \text{window}(s)$ refers to SNPs t that are within some region of the genome surrounding SNP s . Intuitively, this is a heuristic to correct for the over representation in the GRM of some SNPs that are redundant with respect to nearby SNPs.

LD Regression was originally proposed in (Patterson et al. 2006)³¹ and utilized extensively in (Gusev et al. 2013)⁹¹. Only the residual of a SNP—after regressing out other SNPs in LD—in the GRM. :

$$g_s = \frac{x_s - 2p_s}{\sqrt{2p_s(1-p_s)}}$$

$$g_s \sim \sum_{t \in \text{window}(s)} g_t + \varepsilon_s$$

$$G = \sum_{s \in \text{SNPs}} \varepsilon_s \varepsilon_s^T$$

Rare Variants

In considering how to optimally include rare variants in the genome, we examined three strategies. The first strategy was to include all rare variants as described in the computations above without any modifications. The second strategy was to exclude all variants below a threshold, which is a standard strategy used in several recent papers. We compared these simple strategies to a strategy based on reweighting rare variants to optimize the separation between populations.

We considered a particular scenario to optimize. Specifically, we imagine that two populations that split from one another t generations ago are equally represented in our GRM. We would like to optimize the proportion of variance in our GRM that is explained by the true population labels. That is, our figure of merit is:

$$\frac{\frac{1}{n(n-1)} \sum_i \sum_{j \in \text{pop}(i)} g_{i,j} - \frac{1}{n^2} \sum_i \sum_{j \in \text{pop}(i)} g_{i,j}}{\sqrt{\text{var}(g_{i,j})}}$$

where $\text{pop}(i)$ refers to the subpopulation from which individual i came.

Now, considering the population split, our data contains two classes of variants: those variants that are result of mutations predating the population split (pre-split SNPs), and those variants arising after the population split (post-split SNPs). For pre-split SNPs we invoke the normal approximation to genetic drift

described. That is, the difference between allele frequencies p_1, p_2 (for populations 1 and 2, respectively) is:

$$(p_1 - p_2) \sim N(0, 2F_{ST}p(1-p))$$

where p is the allele frequency in the ancestral population prior to the split and F_{ST} quantifies the genetic drift that has occurred since the split. We note that this approximation is reasonable for common SNPs and for small values F_{ST} . If we assume that our data contains only pre-split SNPs then our figure of merit is optimized by the standard computation of the GRM given above. On the other hand, rare, post-split SNPs have the property that

$$|p_1 - p_2| = 2\hat{p}$$

where \hat{p} is the allele frequency estimated from the sample. This difference implies that the optimal weighting for pre-split SNPs is $\frac{1}{\sqrt{p_s(1-p_s)}}$ identically:

$$g_i^s = \frac{x_i - 2p_s}{\sqrt{p_s(1-p_s)}}, \text{ for pre-split SNP } s$$

but the optimal weighting for post-split SNPs is $\sqrt{\frac{F_{ST}^2 + 2F_{ST} + 2}{F_{ST}(1-2p_s)}}$.

$$g_i^s = (x_i - 2p_s) \sqrt{\frac{F_{ST}^2 + 2F_{ST} + 2}{F_{ST}(1-2p_s)}}, \text{ for post-split SNP } s$$

However, this modification requires knowledge of the F_{ST} between studied subpopulations and, more dauntingly, which SNPs are post-split. We believe it is reasonable to iterate over several values of F_{ST} (and find that in real data results are relatively robust to choice of F_{ST}). In order to deal with uncertainty over the set of post-split SNPs, we propose that a SNP be considered post-split if

$$\frac{1}{\sqrt{p_s(1-p_s)}} > \sqrt{\frac{F_{ST}^2 + 2F_{ST} + 2}{F_{ST}(1-2p_s)}}$$

We examine the effect of both of these modifications on the effectiveness of PCA to separate genetically disparate subpopulations.

Analysis of Northern vs. Southern Europe in POPRES Targeted Sequencing Data

We analyzed 531 individuals from the UK referred to as Northern European and 146 Italian, 134 Portuguese, 100 Spaniards, and 7 Swiss Italian individuals collectively referred to as Southern European¹⁰. We excluded 25.9 kb of sequence data from genes on the X chromosome, focusing solely on the autosomes. In total, 8,469 SNPs were polymorphic in either of the Northern or Southern European Samples. These variants were overwhelmingly rare, with 81.5% of variants having a MAF < 1% in the combined sample.

We tested various methods to correct for LD and better handle rare variants (see Methods). The results are summarized in Table S11. These results indicate that handling of both rare variants and LD is critical to maximizing the performance of PCA on this class of data. Applying standard PCA, the top 5 PCs explained only 2.3% of the variance ($r^2=0.023$) of the true population labels. This was improved substantially by removing or reweighting rare variants with ($r^2=0.287, 0.341, 0.352$) for removing variants with MAF < 0.02, removing singletons and reweighting, respectively. This indicates that rare variants, particularly singletons, may be problematic when analyzed using PCA. However, the difference between removing variants with MAF < 0.02 and reweighting ($r^2=0.287$ vs 0.352) suggests that these variants do contain useful information for ancestry inference and should not be universally excluded.

Additionally, application of a method to correct for LD significantly improved performance of PCA when performed in conjunction with singleton exclusion or rare variant reweighting. With rare variant reweighting, LD shrinkage⁸ ($r^2=0.563$) performing slightly better than LD regression ($r^2=0.528$)² and LD

pruning ($r^2=0.534$). While LD Pruning performed well, this may be due to the fact that LD is broken up because the dataset contains sequence data from separated chunks of genome.

Recommendations

In data sets that do not include pervasive LD or large numbers of rare variants (i.e. genotyping data), standard techniques are likely to be successful in detecting population structure. However, in data sets that have pervasive LD and large numbers of rare variants, we recommend that LD pruning and singleton removal be applied. While more sophisticated methods for dealing with these issues were assessed, we did not observe significant improvements above and beyond these simpler approaches. Importantly, we do not recommend that all low frequency and rare variants ($MAF < 0.02$) be removed as these variants do significantly improve detection of population structure.

Acknowledgements

We are grateful to D. Reich for helpful discussions and S. Pollack for assistance with FastPCA software. This research was funded by NIH grant R01 HG006399. SM is funded by NSF grants DMS-1209155 and DMS-1418261.

Web Resources

EIGENSOFT version 6.1, including open-source implementation of FastPCA and smartpca and the PC-based selection statistic: <https://data.broadinstitute.org/alkesgroup/EIG6.1/>

PLINK2: <https://www.cog-genomics.org/plink2/>

flashpca: <https://github.com/gabraham/flashpca>

GERA cohort: http://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000674.v1.p1

1000 Genomes: <http://www.1000genomes.org/>

Shellfish: <http://www.stats.ox.ac.uk/~davison/software/shellfish/shellfish.ph>

Description of Supplemental Data

Supplemental Data include eight figures and eleven tables.

References

1. Sabeti, P. C. *et al.* Positive Natural Selection in the Human Lineage. *Science* **312**, 1614–1620 (2006).
2. Nielsen, R., Hellmann, I., Hubisz, M., Bustamante, C. & Clark, A. G. Recent and ongoing selection in the human genome. *Nat. Rev. Genet.* **8**, 857–868 (2007).
3. Novembre, J. & Di Rienzo, A. Spatial patterns of variation due to natural selection in humans. *Nat. Rev. Genet.* **10**, 745–755 (2009).
4. Scheinfeldt, L. B. & Tishkoff, S. A. Recent human adaptation: genomic approaches, interpretation and insights. *Nat. Rev. Genet.* **14**, 692–702 (2013).
5. Karlsson, E. K., Kwiatkowski, D. P. & Sabeti, P. C. Natural selection and infectious disease in human populations. *Nat. Rev. Genet.* **15**, 379–393 (2014).
6. Shriver, M. D. *et al.* The genomic distribution of population substructure in four populations using 8,525 autosomal SNPs. *Hum. Genomics* **1**, 274 (2004).
7. Hamblin, M. T. & Di Rienzo, A. Detection of the Signature of Natural Selection in Humans: Evidence from the Duffy Blood Group Locus. *Am. J. Hum. Genet.* **66**, 1669–1679 (2000).
8. Akey, J. M., Zhang, G., Zhang, K., Jin, L. & Shriver, M. D. Interrogating a High-Density SNP Map for Signatures of Natural Selection. *Genome Res.* **12**, 1805–1814 (2002).
9. Bersaglieri, T. *et al.* Genetic Signatures of Strong Recent Positive Selection at the Lactase Gene. *Am. J. Hum. Genet.* **74**, 1111–1120 (2004).
10. Lamason, R. L. *et al.* SLC24A5, a Putative Cation Exchanger, Affects Pigmentation in Zebrafish and Humans. *Science* **310**, 1782–1786 (2005).
11. Tishkoff, S. A. *et al.* Convergent adaptation of human lactase persistence in Africa and Europe. *Nat. Genet.* **39**, 31–40 (2007).
12. Perry, G. H. *et al.* Diet and the evolution of human amylase gene copy number variation. *Nat. Genet.* **39**, 1256–1260 (2007).
13. Ayodo, G. *et al.* Combining Evidence of Natural Selection with Association Analysis Increases Power to Detect Malaria-Resistance Variants. *Am. J. Hum. Genet.* **81**, 234–242 (2007).
14. Han, Y. *et al.* Evidence of positive selection on a class I ADH locus. *Am. J. Hum. Genet.* **80**, 441–456 (2007).
15. Price, A. L. *et al.* The Impact of Divergence Time on the Nature of Population Structure: An Example from Iceland. *PLoS Genet* **5**, e1000505 (2009).

16. Xu, S. *et al.* Genomic Dissection of Population Substructure of Han Chinese and Its Implication in Association Studies. *Am. J. Hum. Genet.* **85**, 762–774 (2009).
17. Yi, X. *et al.* Sequencing of 50 Human Exomes Reveals Adaptation to High Altitude. *Science* **329**, 75–78 (2010).
18. Bigham, A. *et al.* Identifying Signatures of Natural Selection in Tibetan and Andean Populations Using Dense Genome Scan Data. *PLoS Genet* **6**, e1001116 (2010).
19. Bhatia, G. *et al.* Genome-wide Comparison of African-Ancestry Populations from CARE and Other Cohorts Reveals Signals of Natural Selection. *Am. J. Hum. Genet.* **89**, 368–381 (2011).
20. Hancock, A. M. *et al.* Adaptations to Climate-Mediated Selective Pressures in Humans. *PLoS Genet* **7**, e1001375 (2011).
21. Ko, W.-Y. *et al.* Identifying Darwinian Selection Acting on Different Human APOL1 Variants among Diverse African Populations. *Am. J. Hum. Genet.* **93**, 54–66 (2013).
22. Engelken, J. *et al.* Extreme Population Differences in the Human Zinc Transporter ZIP4 (SLC39A4) Are Explained by Positive Selection in Sub-Saharan Africa. *PLoS Genet* **10**, e1004128 (2014).
23. Gurdasani, D. *et al.* The African Genome Variation Project shapes medical genetics in Africa. *Nature* **517**, 327–332 (2015).
24. Fumagalli, M. *et al.* Greenlandic Inuit show genetic signatures of diet and climate adaptation. *Science* **349**, 1343–1347 (2015).
25. Kwiatkowski, D. P. How Malaria Has Affected the Human Genome and What Human Genetics Can Teach Us about Malaria. *Am. J. Hum. Genet.* **77**, 171–192 (2005).
26. Price, A. L. *et al.* Principal components analysis corrects for stratification in genome-wide association studies. *Nat. Genet.* **38**, 904–909 (2006).
27. Price, A. L. *et al.* Discerning the Ancestry of European Americans in Genetic Association Studies. *PLoS Genet* **4**, e236 (2008).
28. Novembre, J. & Stephens, M. Interpreting principal component analyses of spatial population genetic variation. *Nat. Genet.* **40**, 646–649 (2008).
29. Tishkoff, S. A. *et al.* The Genetic Structure and History of Africans and African Americans. *Science* **324**, 1035–1044 (2009).
30. Leslie, S. *et al.* The fine-scale genetic structure of the British population. *Nature* **519**, 309–314 (2015).
31. Patterson, N., Price, A. L. & Reich, D. Population Structure and Eigenanalysis. *PLoS Genet* **2**, e190 (2006).

32. Pritchard, J. K., Stephens, M. & Donnelly, P. Inference of Population Structure Using Multilocus Genotype Data. *Genetics* **155**, 945–959 (2000).
33. Alexander, D. H., Novembre, J. & Lange, K. Fast model-based estimation of ancestry in unrelated individuals. *Genome Res.* **19**, 1655–1664 (2009).
34. Tian, C. *et al.* Analysis and Application of European Genetic Substructure Using 300 K SNP Information. *PLoS Genet* **4**, e4 (2008).
35. Seldin, M. F. & Price, A. L. Application of Ancestry Informative Markers to Association Studies in European Americans. *PLoS Genet* **4**, e5 (2008).
36. Li, J. Z. *et al.* Worldwide Human Relationships Inferred from Genome-Wide Patterns of Variation. *Science* **319**, 1100–1104 (2008).
37. Jakobsson, M. *et al.* Genotype, haplotype and copy-number variation in worldwide human populations. *Nature* **451**, 998–1003 (2008).
38. Schlebusch, C. M. *et al.* Genomic Variation in Seven Khoe-San Groups Reveals Adaptation and Complex African History. *Science* **338**, 374–379 (2012).
39. Lazaridis, I. *et al.* Ancient human genomes suggest three ancestral populations for present-day Europeans. *Nature* **513**, 409–413 (2014).
40. Moreno-Estrada, A. *et al.* The genetics of Mexico recapitulates Native American substructure and affects biomedical traits. *Science* **344**, 1280–1285 (2014).
41. Rokhlin, V., Szlam, A. & Tygert, M. A Randomized Algorithm for Principal Component Analysis. *SIAM J. Matrix Anal. Appl.* **31**, 1100–1124 (2009).
42. Halko, N., Martinsson, P., Shkolnisky, Y. & Tygert, M. An Algorithm for the Principal Component Analysis of Large Data Sets. *SIAM J. Sci. Comput.* **33**, 2580–2594 (2011).
43. Halko, N., Martinsson, P. & Tropp, J. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Rev.* **53**, 217–288 (2011).
44. Edenberg, H. J. & Foroud, T. Genetics and alcoholism. *Nat. Rev. Gastroenterol. Hepatol.* **10**, 487–494 (2013).
45. Whitfield, J. B. Alcohol Dehydrogenase and Alcohol Dependence: Variation in Genotype-Associated Risk between Populations. *Am. J. Hum. Genet.* **71**, 1247–1250 (2002).
46. Li, H. *et al.* Diversification of the ADH1B Gene during Expansion of Modern Humans. *Ann. Hum. Genet.* **75**, 497–507 (2011).
47. Gelernter, J. *et al.* Genome-wide association study of alcohol dependence: significant findings in African- and European-Americans including novel risk loci. *Mol. Psychiatry* **19**, 41–49 (2014).

48. Osier, M. V. *et al.* A Global Perspective on Genetic Variation at the ADH Genes Reveals Unusual Patterns of Linkage Disequilibrium and Diversity. *Am. J. Hum. Genet.* **71**, 84–99 (2002).
49. Peter, B. M., Huerta-Sanchez, E. & Nielsen, R. Distinguishing between Selective Sweeps from Standing Variation and from a De Novo Mutation. *PLoS Genet* **8**, e1003011 (2012).
50. Golub, G. H. & Van Loan, C. F. *Matrix Computations*. (Johns Hopkins University Press, 1996).
51. Nicholson, G. *et al.* Assessing population differentiation and isolation from single-nucleotide polymorphism data. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64**, 695–715 (2002).
52. Billingsley, P. *Probability and Measure*. (Wiley-Interscience, 1995).
53. Sabeti, P. C. *et al.* Detecting recent positive selection in the human genome from haplotype structure. *Nature* **419**, 832–837 (2002).
54. Galassi, M. *et al.* *GNU Scientific Library Reference Manual*. (Network Theory Limited, 2009).
55. Yang, J., Zaitlen, N. A., Goddard, M. E., Visscher, P. M. & Price, A. L. Advantages and pitfalls in the application of mixed-model association methods. *Nat. Genet.* **46**, 100–106 (2014).
56. Banda, Y. *et al.* Characterizing Race/Ethnicity and Genetic Ancestry for 100,000 Subjects in the Genetic Epidemiology Research on Adult Health and Aging (GERA) Cohort. *Genetics* genetics.115.178616 (2015). doi:10.1534/genetics.115.178616
57. Chen, C.-Y. *et al.* Improved ancestry inference using weights from external reference panels. *Bioinformatics* **29**, 1399–1406 (2013).
58. Purcell, S. *et al.* PLINK: A Tool Set for Whole-Genome Association and Population-Based Linkage Analyses. *Am. J. Hum. Genet.* **81**, 559–575 (2007).
59. Nelson, M. R. *et al.* The Population Reference Sample, POPRES: A Resource for Population, Disease, and Pharmacological Genetics Research. *Am. J. Hum. Genet.* **83**, 347–358 (2008).
60. Abraham, G. & Inouye, M. Fast Principal Component Analysis of Large-Scale Genome-Wide Data. *PLoS ONE* **9**, e93766 (2014).
61. Bhatia, G., Patterson, N., Sankararaman, S. & Price, A. L. Estimating and interpreting F_{ST}: The impact of rare variants. *Genome Res.* **23**, 1514–1521 (2013).
62. Fellay, J. *et al.* A Whole-Genome Association Study of Major Determinants for Host Control of HIV-1. *Science* **317**, 944–947 (2007).
63. Zou, F., Lee, S., Knowles, M. R. & Wright, F. A. Quantification of Population Structure Using Correlated SNPs by Shrinkage Principal Components. *Hum. Hered.* **70**, 9–22 (2010).
64. de Bakker, P. I. W. *et al.* A high-resolution HLA and SNP haplotype map for disease association studies in the extended human MHC. *Nat. Genet.* **38**, 1166–1172 (2006).

65. Voight, B. F., Kudaravalli, S., Wen, X. & Pritchard, J. K. A Map of Recent Positive Selection in the Human Genome. *PLoS Biol* **4**, e72 (2006).
66. Burton, P. R. *et al.* Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature* **447**, 661–678 (2007).
67. Pickrell, J. K. *et al.* Signals of recent positive selection in a worldwide sample of human populations. *Genome Res.* **19**, 826–837 (2009).
68. Sabeti, P. C. *et al.* Genome-wide detection and characterization of positive selection in human populations. *Nature* **449**, 913–918 (2007).
69. Li, H. *et al.* Geographically Separate Increases in the Frequency of the Derived ADH1B*47His Allele in Eastern and Western Asia. *Am. J. Hum. Genet.* **81**, 842–846 (2007).
70. Treutlein, J., Frank, J., Kiefer, F. & Rietschel, M. ADH1B Arg48His Allele Frequency Map: Filling in the Gap for Central Europe. *Biol. Psychiatry* **75**, e15 (2014).
71. The 1000 Genomes Project Consortium. An integrated map of genetic variation from 1,092 human genomes. *Nature* **491**, 56–65 (2012).
72. The Endogenous Hormones and Breast Cancer Collaborative Group. Insulin-like growth factor 1 (IGF1), IGF binding protein 3 (IGFBP3), and breast cancer risk: pooled individual data analysis of 17 prospective studies. *Lancet Oncol.* **11**, 530–542 (2010).
73. Wood, A. R. *et al.* Defining the role of common variation in the genomic and biological architecture of adult human height. *Nat. Genet.* **46**, 1173–1186 (2014).
74. Ganesh, S. K. *et al.* Effects of Long-Term Averaging of Quantitative Blood Pressure Traits on the Detection of Genetic Associations. *Am. J. Hum. Genet.* **95**, 49–65 (2014).
75. Zhu, X. *et al.* Meta-analysis of Correlated Traits via Summary Statistics from GWASs with an Application in Hypertension. *Am. J. Hum. Genet.* **96**, 21–36 (2015).
76. Buck, D. *et al.* Genetic variants in the immunoglobulin heavy chain locus are associated with the IgG index in multiple sclerosis. *Ann. Neurol.* **73**, 86–94 (2013).
77. Loh, P.-R. *et al.* Efficient Bayesian mixed-model analysis increases association power in large cohorts. *Nat. Genet.* **47**, 284–290 (2015).
78. Yang, J. *et al.* Common SNPs explain a large proportion of the heritability for human height. *Nat. Genet.* **42**, 565–569 (2010).
79. Yang, J. *et al.* Genome partitioning of genetic variation for complex traits using common SNPs. *Nat. Genet.* **43**, 519–525 (2011).
80. Suo, C. *et al.* Natural positive selection and north–south genetic diversity in East Asia. *Eur. J. Hum. Genet.* **20**, 102–110 (2012).

81. Yang, W.-Y., Novembre, J., Eskin, E. & Halperin, E. A model-based approach for analysis of spatial structure in genetic data. *Nat. Genet.* **44**, 725–731 (2012).
82. Bhatia, G. *et al.* Genome-wide Scan of 29,141 African Americans Finds No Evidence of Directional Selection since Admixture. *Am. J. Hum. Genet.* **95**, 437–444 (2014).
83. He, Y. *et al.* A probabilistic method for testing and estimating selection differences between populations. *Genome Res.* gr.192336.115 (2015). doi:10.1101/gr.192336.115
84. Chen, G.-B., Lee, S. H., Zhu, Z.-X., Benyamin, B. & Robinson, M. R. EigenGWAS: finding loci under selection through genome-wide association studies of eigenvectors in structured populations. *bioRxiv* 023457 (2015). doi:10.1101/023457
85. Ko, A. *et al.* Amerindian-specific regions under positive selection harbour new lipid variants in Latinos. *Nat. Commun.* **5**, (2014).
86. Clayton, D. G. *et al.* Population structure, differential bias and genomic control in a large-scale, case-control association study. *Nat. Genet.* **37**, 1243–1246 (2005).
87. Mathieson, I. & McVean, G. Differential confounding of rare and common variants in spatially structured populations. *Nat. Genet.* **44**, 243–246 (2012).
88. Lawson, D. J., Hellenthal, G., Myers, S. & Falush, D. Inference of Population Structure using Dense Haplotype Data. *PLoS Genet* **8**, e1002453 (2012).
89. Novembre, J. *et al.* Genes mirror geography within Europe. *Nature* **456**, 98–101 (2008).
90. The 1000 Genomes Project Consortium (last). A map of human genome variation from population-scale sequencing. *Nature* **467**, 1061–1073 (2010).
91. Gusev, A. *et al.* Quantifying Missing Heritability at Known GWAS Loci. *PLoS Genet* **9**, e1003993 (2013).

Figures

Figure 1. Running time and memory requirements of FastPCA and other algorithms.

The CPU time and memory usage of FastPCA scale linearly with the number of individuals. On the other hand, smartpca and PLINK2-pca scale between quadratically and cubically, depending on whether computing the GRM (quadratic) or the eigendecomposition (cubic) is the rate-limiting step. The running time of flashpca scales quadratically (because it computes the GRM), but its memory usage scales linearly because it stores the normalized genotype matrix in memory. With 50k individuals, smartpca exceeded the time constraint (100 hours) and flashpca exceeded the memory constraint (40GB). With 70k individuals, PLINK2-pca exceeded the memory constraint (40GB). Run times are based on one core of a 2.26-GHz Intel Xeon L5640 processor; we caution that run time comparisons may vary by a small constant factor as a function of the computing environment. Numerical data are provided in Table S1.




Figure 2. Accuracy of FastPCA and PLINK2-pca.

FastPCA and PLINK2-pca were run on simulated populations of varying divergence. The simulated data comprised 50k SNPs and 10k total individuals from six subpopulations derived from a single ancestral population. PCs computed by PLINK2-pca and FastPCA were compared to the true population PCs and to each other using the Mean of Explained Variances (MEV) metric (see text). FastPCA explained the same amount of true population variance as PLINK2-pca in all experiments, and the methods output nearly identical PCs (MEV>0.999).




Figure 3. Power of PC-based selection statistic.

The allele frequency difference at selected SNPs was varied between two populations separated by varying F_{ST} . The significance threshold was set to 8.3×10^{-7} based on 60K SNPs tested. (a) With 50k samples, the power curves for $F_{ST} = \{0.05, 0.02, 0.01, 0.005, 0.002, 0.001, 0.0005\}$ showed a phase change. (b) Varying the number of samples for $F_{ST} = 0.001$ demonstrated that this phase change was more gradual at smaller sample sizes. (c) Varying the number of samples at $F_{ST} = 0.01$ showed that the impact of sample size was less pronounced than at $F_{ST} = 0.001$.






Figure 4. FastPCA results on GERA data set.

FastPCA and SNPweights⁵⁷ were run on the GERA cohort and the principal components from FastPCA were plotted. Individuals were colored by mapping Northwest European (NW), Southeast European (SE) and Ashkenazi Jewish (AJ) ancestry estimated by SNPweights to the red/green/blue color axes (see Methods). PC1 and PC2 separate the GERA cohort into northwest (NW), southeast (SE) and Ashkenazi Jewish (AJ) subpopulations. PC3 separates the AJ and SE individuals, while PC3 and PC4 further separates the NW European individuals.




Figure 5. Separation of Irish, Eastern European and Northern European individuals in GERA data.

We report results of projecting POPRES⁵⁹ individuals onto top PCs. The plot of PC3 vs PC4 shows that the Northwest European (NW) individuals are further separated into Irish and Eastern European and Northern European populations. Projected populations were colored based on correspondence to the ancestry assignment from SNPweights⁵⁷, except that Irish and Eastern European individuals were colored purple and orange, respectively, to indicate additional population structure.


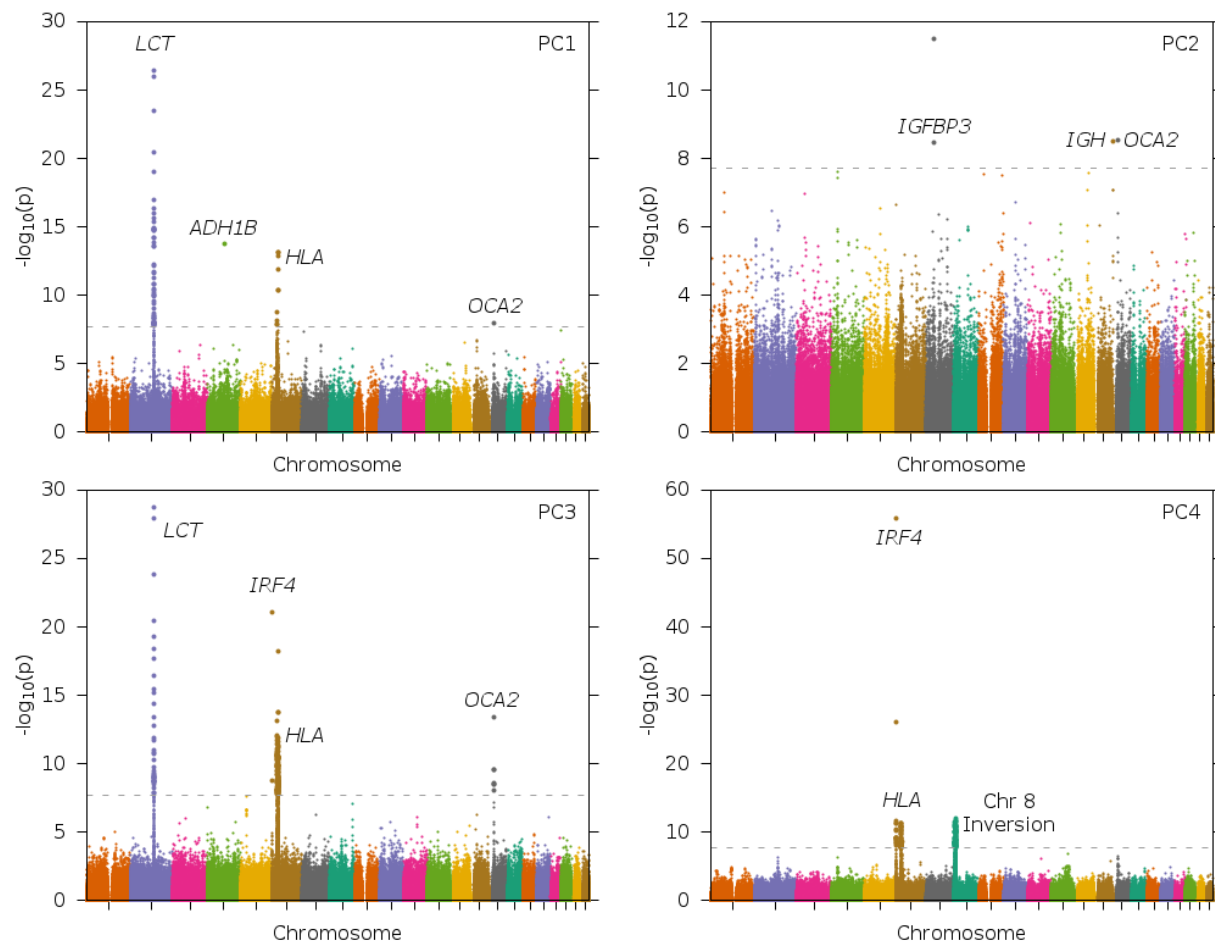


Figure 6. Signals of selection in the top PCs of GERA data.

We display Manhattan plots for selection statistics computed using each of the top 4 PCs. The grey line indicates the genome-wide significance threshold of 2.05×10^{-8} based on 2,435,924 hypotheses tested ($\alpha = 0.05$, 608,981 SNPs x 4 PCs).



Tables

Table 1. Genome-wide significant signals of selection in GERA data.

We list regions with genome-wide significant ($\alpha = 0.05$, Bonferroni correction with 608,981 SNPs x 4 PCs = 2,435,924 hypotheses tested, $p < 2.05 \times 10^{-8}$) evidence of selection in the top 4 PCs. Loci that were not previously known to be under selection in Europeans are indicated in bold font. The chromosome 8 inversion signal is due to a PC artifact (see main text). Regions with suggestive evidence of selection ($10^{-6} < p < 2.05 \times 10^{-8}$) are listed in Table S3.

Locus	Chromosome	Region (Mb)	PC	Best Hit	<i>p</i> -value
LCT ⁹	2	134.8 – 137.6	1	rs6754311	4.15×10^{-27}
			3	rs4988235	1.83×10^{-29}
ADH1B	4	100.5	1	rs1229984	1.67×10^{-14}
IRF4 ^{66,67}	6	0.3 – 0.5	3	rs12203592	8.69×10^{-22}
			4		1.83×10^{-56}
HLA ⁶⁴	6	30.8 – 33.3	1	rs382259	7.95×10^{-14}
			3	rs9268628	6.52×10^{-19}
			4	rs34707463	4.76×10^{-12}
IGFBP3	7	45.3-45.9	2	rs150353309	3.14×10^{-12}
Chr8 Inversion ³⁴	8	8.2 – 11.9	4	rs6984496	9.21×10^{-13}
IGH	14	106.0-106.1	2	rs34614900	3.34×10^{-9}
OCA2 ^{65,67}	15	25.9 – 26.2	1	rs12916300	1.12×10^{-8}
			2		3.07×10^{-9}
			3		4.29×10^{-14}

Table 2. Performance of natural selection statistic in subsampled data.

The selection statistic was computed in random subsets of individuals of specified size for each SNP in Table 1 (except for the chromosome 8 inversion region) and the known selection regions TLR1⁶⁶ and SLC45A2⁶⁸ in Table S4. We report the median selection statistic P-value across 100 random subsets.

Locus	SNP	Sample size						
		Full data set	1k	2k	5k	10k	20k	50k
LCT	rs6754311	2.15e-25	4.91e-17	2.97e-20	1.53e-23	1.17e-24	2.63e-25	1.02e-26
	rs4988235	1.15e-27	7.44e-17	9.80e-20	4.64e-23	3.11e-24	2.69e-25	1.62e-27
	rs17346504	8.41e-7	2.86e-2	1.25e-2	9.49e-4	6.03e-5	8.12e-6	9.80e-7
ADH1B	rs1229984	1.26e-13	3.91e-9	3.51e-11	1.97e-12	5.54e-13	1.50e-13	1.31e-13
IRF4	rs12203592	5.52e-55	3.15e-6	9.18e-12	7.47e-25	7.21e-36	7.02e-45	2.19e-54
HLA	rs382259	5.38e-13	8.68e-9	1.23e-10	7.07e-12	1.85e-12	7.51e-13	5.77e-13
	rs9268628	8.66e-18	3.62e-5	3.41e-7	5.97e-12	2.10e-14	2.68e-16	1.00e-17
	rs4394275	9.36e-12	8.40e-2	1.94e-3	1.44e-5	4.00e-8	7.86e-10	1.24e-11
IGFBP3	rs150353309	5.82e-12	5.90e-4	1.49e-5	2.72e-8	3.61e-10	3.34e-11	6.61e-12
IGH	rs34614900	5.23e-9	6.33e-3	2.24e-4	2.26e-6	2.01e-7	3.32e-8	5.32e-9
OCA2	rs12916300	2.80e-13	6.29e-6	1.07e-7	3.67e-9	1.94e-11	5.29e-12	3.11e-13
	rs2703951	5.11e-7	1.12e-1	2.45e-2	7.96e-4	7.17e-5	4.52e-6	5.74e-7
TLR1	rs5743611	5.42e-8	8.05e-3	4.27e-4	9.41e-6	1.19e-6	2.17e-7	5.60e-8
	rs4833095	6.52e-7	6.07e-4	3.37e-4	7.35e-5	3.64e-5	6.03e-6	7.10e-7
SLC45A2	rs16891982	6.89e-8	8.25e-4	2.17e-4	1.93e-5	4.55e-6	2.46e-7	7.31e-8