

A Method to Exploit the Structure of Genetic Ancestry Space to Enhance Case-Control Studies

Corneliu A. Bodea,^{1,3,4} Benjamin M. Neale,^{2,3,4} Stephan Ripke,^{3,4,7} The International IBD Genetics Consortium, Mark J. Daly,^{2,3,4} Bernie Devlin,⁵ and Kathryn Roeder^{1,6,*}

One goal of human genetics is to understand the genetic basis of disease, a challenge for diseases of complex inheritance because risk alleles are few relative to the vast set of benign variants. Risk variants are often sought by association studies in which allele frequencies in case subjects are contrasted with those from population-based samples used as control subjects. In an ideal world we would know population-level allele frequencies, releasing researchers to focus on case subjects. We argue this ideal is possible, at least theoretically, and we outline a path to achieving it in reality. If such a resource were to exist, it would yield ample savings and would facilitate the effective use of data repositories by removing administrative and technical barriers. We call this concept the Universal Control Repository Network (UNICORN), a means to perform association analyses without necessitating direct access to individual-level control data. Our approach to UNICORN uses existing genetic resources and various statistical tools to analyze these data, including hierarchical clustering with spectral analysis of ancestry; and empirical Bayesian analysis along with Gaussian spatial processes to estimate ancestry-specific allele frequencies. We demonstrate our approach using tens of thousands of control subjects from studies of Crohn disease, showing how it controls false positives, provides power similar to that achieved when all control data are directly accessible, and enhances power when control data are limiting or even imperfectly matched ancestrally. These results highlight how UNICORN can enable reliable, powerful, and convenient genetic association analyses without access to the individual-level data.

Introduction

To detect genetic variants affecting risk for complex disease, the ideal association study would contrast a large number of affected subjects to an even larger set of population-based samples used as control subjects. Ideally these control subjects would be so numerous and so well-matched to case subjects, ancestrally, that the power to detect risk variants would be limited solely by the size of the case sample. This article outlines an approach to turn this ideal into reality.

The challenges in accruing a large control sample are numerous. It requires a substantial portion of the research budget; although data repositories, such as dbGaP,^{1,2} contain genetic data from tens of thousands of potential control samples, using these data requires considerable and independent effort from each research team; and issues such as population structure and genotyping platform require additional work before an adequately controlled association test can be performed. Family-based studies obviate concerns about ancestry,^{3,4} but they have other drawbacks.^{5–9}

Instead we show here that it is theoretically possible to build a web resource that enables research teams to focus on maximizing the value of their case sample by providing control allele frequency information that is optimally matched to the available case subjects. Additionally, infor-

mation can be exchanged via a web server similar to the existing Exome Aggregation server, without revealing individual genetic information. We call such a resource the Universal Control Repository Network (UNICORN), because it provides matched control data for a variety of ancestries. In our vision, and to ensure the confidentiality of both case and control subjects, no case genotype information is passed to UNICORN, nor will the control subjects' data processed to produce UNICORN be accessible to this resource.

Our approach to building UNICORN employs the spectral graph approach,¹⁰ which has similarities to principal-component analysis,^{11,12} to obtain a hierarchical representation of ancestry, where individuals are clustered into increasingly finer ancestry spaces. Using a Bayesian model, we infer allele frequencies over all such clusters, always borrowing strength across the entire hierarchy to maximize power. We then perform a second layer of inference within clusters to model spatial variation. This step picks up fine-grained ancestry structure that the hierarchical clustering did not by assuming that deviations from a cluster-wide average follow a Gaussian process with a covariance structure that is inferred from the ancestry space. This model is appropriate because it is flexible enough to accommodate smooth allele frequency fluctuations with varying degrees of spatial correlation.

¹Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213, USA; ²Program in Medical and Population Genetics, Broad Institute of MIT and Harvard, Cambridge, MA 02142, USA; ³Analytical and Translational Genetics Unit, Department of Medicine, Massachusetts General Hospital and Harvard Medical School, Boston, MA 02114, USA; ⁴Stanley Center for Psychiatric Research, Broad Institute of MIT and Harvard, Cambridge, MA 02142, USA; ⁵Department of Psychiatry, University of Pittsburgh School of Medicine, Pittsburgh, PA 15213, USA; ⁶Computational Biology Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA; ⁷Department of Psychiatry and Psychotherapy, Charité, Campus Mitte, 10117 Berlin, Germany

*Correspondence: roeder@andrew.cmu.edu
<http://dx.doi.org/10.1016/j.ajhg.2016.02.025>

©2016 by The American Society of Human Genetics. All rights reserved.

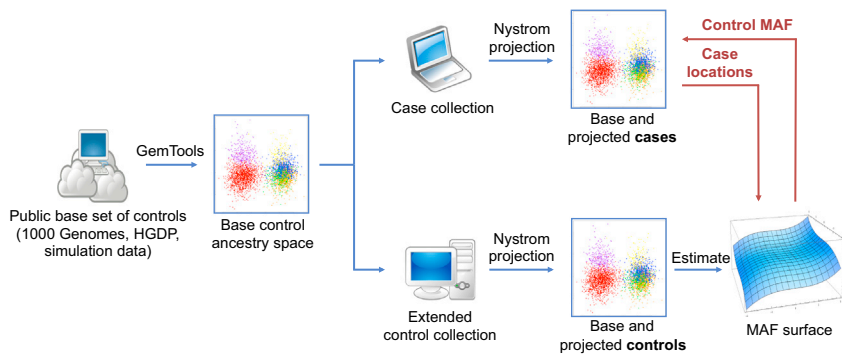


Figure 1. Overview of the UNICORN Model
The UNICORN pipeline starts with a public base set of control subjects and constructs the corresponding base control ancestry space. All subsequent case and control subjects can be projected independently via GemTools onto this space. This approach ensures that, having only knowledge of the base set, new individuals can be compared to existing ancestries. An extended set of control subjects is then projected onto the base control ancestry space, which is used to estimate the minor allele frequency distribution (MAFD) over the ancestry space. To query the repository, researchers project their case subjects onto the base control ancestry

space and submit the resulting coordinates to the UNICORN server. Users then receive control allele frequencies as well as the degree of uncertainty associated with these estimates for all relevant locations, based on the pre-computed MAFD. Users can then proceed with an association test. Users need to submit only ancestry coordinates and the system returns only frequency inferences for the corresponding locations (red arrows). No other information is exchanged.

Our results on both simulated data and imputation-based genotype-level data from seven studies of Crohn disease show that UNICORN has the potential to greatly improve power in genetic association tests. First we show that UNICORN not only controls false positives but also that it makes efficient use of the control data, providing power similar to a setting in which all control data are directly accessible to the researcher. We then show that UNICORN can improve power relative to a carefully matched case-control study simply by using all available control information, even though the additional control subjects are not perfectly matched to case subjects.

Subjects and Methods

Overview of UNICORN

The steps involved in building our version of UNICORN (henceforth simply UNICORN) and performing an association study are now outlined (Figure 1). Existing publicly available collections of control data determine a common genetic ancestry space onto which case and control subjects can be projected independently. GemTools^{10,13,14} constructs ancestry spaces and performs such projections. The projected controls are then used to estimate the control minor allele frequency distribution (MAFD) over the ancestry space. For efficiency of computations, the MAFD would be precomputed and stored for application whenever users request control information. To query the repository, researchers project their case subjects onto the public control ancestry space and submit the locations to UNICORN. Based on the pre-computed surface, the system will infer allele frequencies as well as the degree of uncertainty associated with the estimates at all relevant locations and return the results to the users, who can then proceed with an association test, such as the one we describe in the next section.

To estimate the MAFD, UNICORN employs a combination of empirical Bayesian analysis across a hierarchical clustering of the control subjects and, for localized ancestry regions, a Gaussian process model of the minor allele frequency (Figure 2). To visualize the algorithm in action, we utilize the Europeans in the Population Reference Sample (POPRES)¹⁵ (dbGaP: phs000145.v4.p2), which yields an ancestry map that approximates the geographic

map of Europe.^{16,17} Two SNPs in LCT (lactase persistence) and OCA2 (hair, skin, and eye color) provide examples of UNICORN's MAFD for SNPs under selection and provide an illustration of clines in allele frequency across Europe. Intensity of color displays allele frequency estimates that vary smoothly across the map (Figure 3).

Conceptually, UNICORN aims to use as many control samples as justifiable, based on ancestry, to estimate the MAFD associated with each case sample. To motivate this model, consider two different matched case-control studies: one with equal numbers of case and control subjects and the other with ten control subjects for each case subject. In the first instance, the statistical power is driven equally by case and control subjects; for the latter, the number of case subjects is the key determinant for power. For UNICORN, the matching of control to case subjects is determined by how many control subjects are located near each case subject in ancestry space. Regardless of the number of case and control subjects, if there were very few control subjects similar in ancestry to case subjects, any test will have a large variance and little power. Alternatively, if there are many control subjects that are close in ancestry space to each case subject, then the variance of the test will be dominated by the case sample size. UNICORN seeks to achieve power by using information on allele frequencies from a very large sample of control subjects.

Ancestry Mapping via GemTools

Dimension reduction techniques such as principal-component analysis (PCA) are traditionally used to model complex genetic structure and to control for population stratification.^{11,12,17–20} These approaches often require many dimensions to describe the ancestry space, and this is not ideal for downstream steps of UNICORN. Instead, our algorithm first discovers clusters of subjects with relatively homogeneous ancestry, which then require fewer eigenvectors to represent ancestry within a cluster. To achieve this purpose, we use GemTools,¹³ a software tool based on a spectral graph approach¹⁰ quite similar to PCA. We note, however, that many popular ancestry mapping techniques could be successfully paired with UNICORN in place of GemTools.

A first step in the UNICORN algorithm involves plotting both case and control samples onto a common ancestry map without data exchanging hands (Figure 1). This is achieved by generating an ancestry map using a publicly available repository, called the “base sample,” and then projecting case and control subjects

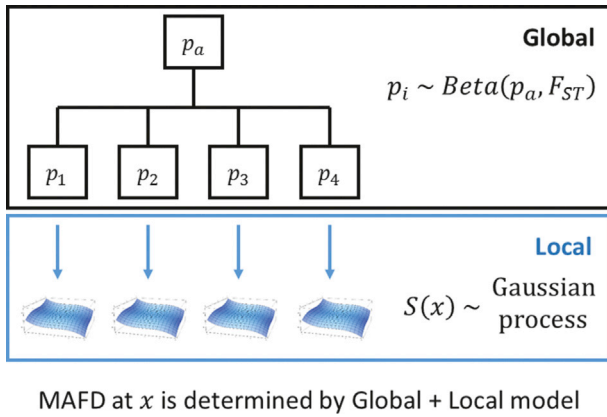


Figure 2. Overview of the Inference Levels

The Global step operates on a cluster-wide resolution, providing estimates for entire clusters based on a beta-binomial model of allele frequencies. The Local step operates within clusters, providing localized estimates across the ancestry space spanned by the individuals in each cluster. This step models allele frequencies as spatial processes operating within clusters. The Global and Local inference modules complement each other, the former picking up larger fluctuations in allele frequencies, and the latter generating a fine map that would otherwise have been hidden by the strong signal at the Global level.

onto this map via the Nyström approximation.^{21–24} When samples are projected onto a given ancestry map, it accurately reflects their ancestry only if the base sample spans the full range of ancestries included in the new samples.²¹ Individuals with unrepresented ancestry will be projected into the available range and they will be falsely represented as more similar to the base sample. Thus, as with any genetic association study, the case collection should be restricted to samples with ancestry similar to the available control samples.

The aim of the spectral graph approach is to obtain a useful eigenmap of the genetic ancestry present in a sample. The population is represented as a weighted graph with vertices denoting individuals and weights denoting genetic similarity. Define the matrix Y such that y_{ik} is the minor allele count for the i^{th} subject at the k^{th} SNP. Center and scale the columns of Y . Instead of proceeding with computing eigenvectors and eigenvalues of YY^t , define the weight matrix W as $w_{ij} = \sqrt{y_i^t y_j}$ if $y_i^t y_j \geq 0$ and 0 otherwise for similarity between the i^{th} and j^{th} subjects. Setting a threshold on YY^t to guarantee non-negative weights creates a skewed distribution of weights, so the choice of a square-root transformation leads to more symmetric distributions. This transformation also increases the robustness to outliers. Let the degree of vertex i be $d_i = \sum_{j=1}^n w_{ij}$ and define the diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$. The normalized graph Laplacian matrix for W is defined as $1 - L$, where $L = D^{-1/2} W D^{-1/2}$. Let v_i and u_i be the eigenvalues and eigenvectors of $1 - L$ and let $\lambda_i = \max\{0, 1 - v_i\}$. We can then map the i^{th} subject onto an s -dimensional ancestry space according to: $[\lambda_i^{1/2} u_1(i), \dots, \lambda_s^{1/2} u_s(i)]$. See Lee et al.¹⁰ for further details.

GemTools builds on this spectral graph approach to construct eigenmaps and provide a hierarchical clustering of individuals based on ancestry.¹³ To speed up computation, it is useful to avoid the cost of calculating the inner product matrix YY^t and then performing a spectral decomposition on a large matrix. GemTools uses a divide-and-conquer approach that clusters individuals of

similar ancestry and then finds eigenmaps for each cluster. Homogeneous clusters of individuals are derived via Ward's k-means algorithm. In addition to reducing computation time, this approach focuses on fine-scale structure across clusters, leading to more informative maps than those resulting from a brute force computation of a single eigenmap of the entire dataset.^{10,14}

New subjects are mapped onto an existing map via Nyström projection. Let Y represent the scaled and centered allele count vectors for the initial n subjects. Let z be the scaled allele count vector of a new individual we wish to project. We define the edge weights between the new subject and an existing individual as $w_{ij} = \sqrt{z^t y}$ if $z^t y \geq 0$ and 0 otherwise. The vertex degree of z is $d(z) = w(z, z) + \sum_{i=1}^n w(z, y_i)$. Then the eigenvector coordinates of z for dimensions $k = 1, \dots, s$ are $u_k(z) = \lambda k^{-1} \sum_{i=1}^n L(z, y_i) u_k(y_i)$, where $L(z, y_i) = [d(z)d(y_i)]^{-1/2} w(z, y_i)$. Nyström projection plays a critical role in UNICORN because it allows two datasets to be mapped to the same ancestry space without the need for data sharing.

To highlight the importance of choosing a representative base sample, we estimate the eigenvectors using two different base samples derived from POPRES¹⁵ and HGDP²⁵ European samples (Figure 4).²¹ When using the HGDP populations as a base (Figure 4A), the axes do not differentiate the POPRES sample. Rather, the points clump together in the center of the eigenspace because their differences are dwarfed by the differences in the more diverse HGDP sample. Likewise, we found that when using the POPRES sample as a base (Figure 4B), the axes do not capture the strong differences in the highly diverse HGDP data.

Cluster-wide Inference

UNICORN estimates ancestry-specific allele frequencies via an efficient, flexible semi-parametric model. Frequencies are modeled in two stages to account for global and local structure. In the first stage, the data are partitioned into approximately homogeneous ancestry clusters based on eigenanalysis.¹⁴ Next, each of these clusters is subsequently described by a secondary eigenanalysis that models local ancestry within a cluster. In stage two, local variability is modeled over the ancestry space via a Gaussian spatial process. The key to modeling local variation in allele frequency is to obtain a parsimonious representation of the ancestry not unlike a geographic map. GemTools recursively partitions the subjects until the clusters are approximately homogeneous, as judged by the leading eigenvalues.¹⁴ Consequently, two eigenvectors are sufficient to describe the residual ancestry differences within clusters at the final stage.

Each stage of the model is amenable to a simple statistical model that accounts for allele frequency variation over the ancestry space and records variability in the allele frequency estimate. In the first stage, the allele counts are modeled via a beta-binomial model with variance a function of the well-known genetic parameter F_{ST} . Assume we have a dataset in which GemTools detects n subpopulations. At this stage we want to find good estimates of the true cluster-wide allele frequencies p_i . We model each of these frequencies as

$$p_i \sim \text{Beta}\left(\frac{p_a(1 - F_{ST})}{F_{ST}}, \frac{(1 - p_a)(1 - F_{ST})}{F_{ST}}\right). \quad (\text{Equation 1})$$

Following an empirical Bayesian setting, we use Equation 1 as a prior for the cluster-wide allele frequency and use the data to guide us in selecting appropriate values for the two hyperparameters p_a and F_{ST} . Let \hat{p}_i be the average allele count in cluster i . Although

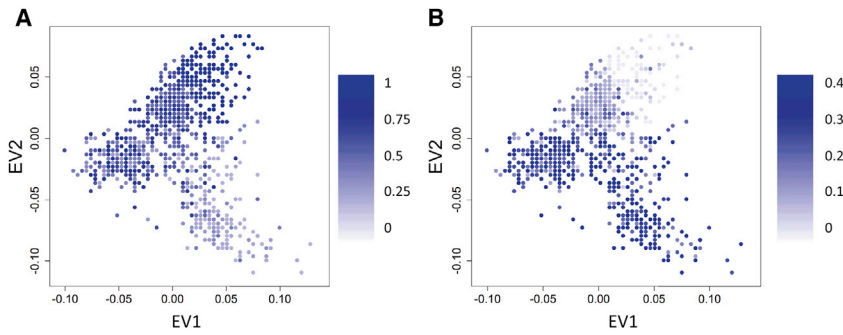


Figure 3. Clines Detected by UNICORN in the POPRES Data for Two SNPs under Strong Selection

Intensity of color displays allele frequency estimates that vary smoothly across the map. (A) Cline of a SNP within the LCT region (lactase persistence). (B) Cline of a SNP within the OCA2 region (hair, skin, and eye color).

\hat{p}_i is an unbiased estimator of p_i , it can have a large variance if few individuals reside in the cluster. From Equation 1 we have

$$p_a = E(p_i) \approx \frac{1}{n} \sum_{i=1}^n \hat{p}_i \stackrel{\text{def}}{=} \hat{p}_a \quad (\text{Equation 2})$$

and

the posterior variance is

$$\text{mean}(\hat{p}_i | \mathbf{y}) = \frac{\hat{p}_a(1 - \hat{F}_{ST}) + \frac{\sum_{j=1}^{n_i} \gamma_j}{2n_i}}{\frac{1 - \hat{F}_{ST}}{2n_i \hat{F}_{ST}} + 1}, \quad (\text{Equation 6})$$

$$\text{var}(\hat{p}_i | \mathbf{y}) = \frac{\left[\frac{\hat{p}_a(1 - \hat{F}_{ST}) + \frac{\sum_{j=1}^{n_i} \gamma_j}{2n_i}}{2n_i \hat{F}_{ST}} + 1 \right] \left[\frac{(1 - \hat{p}_a)(1 - \hat{F}_{ST})}{2n_i \hat{F}_{ST}} + 1 - \frac{\sum_{j=1}^{n_i} \gamma_j}{2n_i} \right]}{\left[\frac{1 - \hat{F}_{ST}}{2n_i \hat{F}_{ST}} + 1 \right]^2 \left[\frac{1}{\hat{F}_{ST}} + 2n_i \right]}, \quad (\text{Equation 7})$$

$$F_{ST} = \frac{\text{var}(p_i)}{p_a(1 - p_a)} \approx \frac{\text{var}(\hat{p}_i)}{\hat{p}_a(1 - \hat{p}_a)} \stackrel{\text{def}}{=} \tilde{F}_{ST}. \quad (\text{Equation 3})$$

This estimator can be improved by taking into account linkage disequilibrium, the tendency of nearby alleles to descend from the same ancestral chromosome. The F_{ST} of nearby alleles must thus be similar, creating a smooth F_{ST} function across the genome. However, \tilde{F}_{ST} can exhibit excessive variation that is alleviated by local smoothing through kernel regression based on genomic location.

Using Equations 1, 2, and 3, we estimate the prior for p_i through

$$\hat{p}_i \sim \text{Beta} \left(\frac{\hat{p}_a(1 - \hat{F}_{ST})}{\hat{F}_{ST}}, \frac{(1 - \hat{p}_a)(1 - \hat{F}_{ST})}{\hat{F}_{ST}} \right). \quad (\text{Equation 4})$$

Assume we observe the genotype vector \mathbf{y} for the n_i individuals located in cluster i . Then the posterior distribution of \hat{p}_i is

$$\hat{p}_i | \mathbf{y} \sim \text{Beta} \left(\frac{\hat{p}_a(1 - \hat{F}_{ST})}{\hat{F}_{ST}} + \sum_{j=1}^{n_i} \gamma_j, \frac{(1 - \hat{p}_a)(1 - \hat{F}_{ST})}{\hat{F}_{ST}} + 2n_i - \sum_{j=1}^{n_i} \gamma_j \right). \quad (\text{Equation 5})$$

This is the distribution for the cluster-wide allele frequency that we will proceed to use for local inference.

From Equation 5 it follows that the posterior mean of \hat{p}_i is

and the posterior $\hat{p}_i | \mathbf{y}$ is a consistent estimator of the true minor allele frequency p_i .

Within-Cluster Inference

In the second stage, local structure is quantified using models made popular in the geostatistics/kriging literature.²⁶ To describe the model, we require the following notation: $Y(x)$ = minor allele count at location x in the eigenspace; $P(x)$ = minor allele frequency at location x ; $S(x)$ = deviation from cluster-wide average allele frequency at location x (spatial structure); β = cluster-wide log odds of minor allele frequency; σ^2 = variance of the stationary Gaussian process (SGP); and ϕ = rate at which the correlation ρ between values of S at different locations decays with increasing distance u . Kriging methods consider a stochastic process $S = \{S(x) : x \in \mathbb{R}^p\}$, called the signal, whose realized values are not directly observed. We do observe Y , the vector of allele counts, which are located in the eigenspace indexed by x . We assume that the distribution of $Y(x)$ depends on $S(x)$ and that the allele counts are a noisy version of S for a given set of locations $x_i, i \in 1, \dots, n$. The goal is to predict $S(x)$ at new locations where the case subjects have been sampled. To model local structure in an ancestry space, we assume that deviations from a cluster-wide average follow a stationary Gaussian process with mean 0 and a covariance structure that will be inferred from the data. Consider the Bayesian kriging setup:

$$\begin{aligned} Y(x) | S(x) &\sim \text{Bin}[2, P(x)] \\ \log \left[\frac{P(x)}{1 - P(x)} \right] &= \beta + S(x) \\ S &\sim \text{SGP}[0, \sigma^2, \rho(u) = e^{-u\phi}]. \end{aligned} \quad (\text{Equation 8})$$




Figure 4. Importance of the Choice of Base Sample for Ancestry Maps

When projecting new samples onto an existing ancestry map, it is crucial that the base sample spans the full range of ancestries present in the new samples. If the projected samples contain unrepresented ancestries, they will still be mapped onto the ancestry range of the base set, thus distorting their true background and leading to strongly heterogeneous clusters that do not accurately reflect the allele frequencies of the new samples.

(A) Base = HGDP (black), projected = POPRES (turquoise). In this scenario we get poor resolution of ancestries in the POPRES sample. This set projects as a clump, because it looks very homogeneous relative to the more diverse HGDP base set. (B) Base = POPRES, projected = HGDP. In this scenario, the HGDP ancestries not present in the POPRES base set are still projected within the POPRES ancestry range.

This model is appropriate because it is flexible enough to accommodate smooth allele frequency fluctuations with varying degrees of spatial correlation. With this two-stage model, we can make use of our hierarchical clustering and at the same time adapt local inference to the variability present in the data, all in a Bayesian framework. Inferences are performed via Metropolis-Hastings. For additional details, see [Appendix A](#). Ultimately, the distribution of the MAF is well approximated by a function that captures the mean and variance of the estimate.

The variance parameters of UNICORN, σ^2 and ϕ , determine how fast allele frequencies fluctuate over the ancestry space. We can use the available ancestry space to make an informed choice of priors for these parameters. To extract the necessary variability information from the data, we use a well-established method from the kriging literature: the variogram.²⁶ The theoretical variogram $\gamma(x, y)$ describes the spatial dependence in a random field:

$$\gamma(x, y) = \frac{\text{var}[Y(x) - Y(y)]}{2}. \quad (\text{Equation 9})$$

If the random field is stationary and isotropic, which is assumed here, then the theoretical variogram can be rewritten as:

$$\begin{aligned} \gamma(u) &= \frac{1}{2} \{\text{var}[Y(x+u) - Y(x)]\} \\ &= \frac{1}{2} E\{[Y(x+u) - Y(x)]^2\} - \frac{1}{2} E[Y(x+u) - Y(x)]^2. \end{aligned} \quad (\text{Equation 10})$$

If $E[Y(x+u)] = E[Y(x)]$, thus under the assumption that there exists no spatial structure, the theoretical variogram is routinely estimated via the empirical variogram:

$$\hat{\gamma}(u) = \frac{1}{2|N(u)|} \sum [Y(x_i) - Y(x_j)]^2, \quad (\text{Equation 11})$$

where the sum is over $N(u) = \{(i, j): x_i - x_j = u\}$ and $|N(u)|$ is the number of distinct elements of $N(u)$. But because we expect spatial structure to be present, we cannot compute the empirical variogram via Equation 11 directly. Instead, we estimate the spatial structure in Y first through linear regression in the ancestry space, and then we use the residuals and Equation 11 to compute a residual empirical variogram. The next step uses both the theoretical

and empirical variogram to derive values for the variance parameters. Because an algebraic expression of the theoretical variogram is complicated, we use simulation to find appropriate estimates of the variance parameters. Priors for σ^2 and ϕ are then chosen so that their mean equals the value derived from the variogram analysis.

Association Test

Each case sample is mapped to a cluster in the hierarchical tree and an ancestry position x within the cluster. Combining the results from our cluster-wide inference and within-cluster inference ([Figure 2](#)), we can obtain the MAFD for this case

$$P(x) = \frac{e^{\beta+S(x)}}{1 + e^{\beta+S(x)}}, \quad (\text{Equation 12})$$

where $S(x)$ is the spatial structure determined via the Gaussian process model (within-cluster inference) and β is based on the beta-binomial model (cluster-wide inference). Specifically, this expression determines the mean, $E[P(x)]$, and the variance, $\text{var}[P(x)]$, of the MAFD (x) which is required to perform an association test.

For an association study, we sample minor allele counts $[Y(x_1), \dots, Y(x_n)]$ for a sample of n cases. Under the null hypothesis (no association), we assume that $Y(x) \sim \text{Bin}(2, P(x))$. It follows that $E[Y(x)] = E[E[Y(x)|P(x)]] = 2E[P(x)]$ and

$$\text{var}[Y(x)] = E[\text{var}(Y(x) | P(x))] + \text{var}[E(Y(x) | P(x))] \quad (\text{Equation 13})$$

$$= 2E[P(x)(1 - P(x))] + 4\text{var}[P(x)]. \quad (\text{Equation 14})$$

The null distribution of \bar{Y} follows from the central limit theorem:

$$\bar{Y} \approx N \left[\frac{2}{n} \sum_{i=1}^n E[P(x_i)], \frac{1}{n^2} \sum_{i=1}^n 2E[P(x_i)(1 - P(x_i))] + 4\text{var}[P(x_i)] \right]. \quad (\text{Equation 15})$$

Z-scores and subsequently p values can be computed for association tests based on Equation 15. This result shows that if many control subjects become available for each case subject (decreasing

$\text{var}[P(x)]$), then the variance of the null distribution will be dominated by the binomial sampling variance in the case subjects. In this setting, the statistic reduces approximately to a test comparing the minor allele frequency in case subjects to a known population quantity and the power of the test is largely determined by the number of case subjects sampled. At the other extreme, if only one matched control is available for each case, then the statistic is equivalent to the usual 2-sample test and has twice the variance attainable by UNICORN with a large sample of control subjects. Provided the case subjects are well matched to a large UNICORN control sample, the power can be approximated using a genetic power calculator with control:case ratio set suitably high, say ten.

Results

Analysis of POPRES Data

To illustrate UNICORN, we use data from POPRES,¹⁵ from which we selected 160,000 high-quality SNPs (MAF > 1% and less than 1% missing genotypes) and 1,000 individuals of European ancestry (each subject must have no more than 1% missing genotypes). The hierarchical ancestry structure was determined via GemTools, yielding an ancestry map that approximates the geographic map of Europe.^{16,27}

For any particular study we expect the UNICORN repository will include 10–20 times as many control samples as case subjects. Moreover, it is likely that only a fraction of these control subjects will be suitably matched in ancestry to the case subjects. Thus to mimic the realistic performance of UNICORN using the POPRES data, we needed to select a small case sample with a particular regional distribution. Specifically, we randomly selected 60 POPRES samples of French and Swiss ancestry to serve as case subjects (6% of the total). For this constructed case-control sample, we simulated causal variants of varying allele frequencies and odds ratios. We first performed a matched-control association test, where the selected case individuals were matched to the nearest control subjects in the ancestry space. We then analyzed the simulated variants via UNICORN and found that it delivered more powerful results even when compared to a standard case-control association test comparing 60 case subjects to 600 ancestry-matched control subjects (Figure S1).

Application to IBD Data

The large meta-analysis study of Crohn disease (CD), including 5,956 case subjects and 14,927 control subjects,²⁸ provides a realistic test of the validity and power of the UNICORN approach. This study is perfect for detailed investigation for two reasons: first, it provides a very large sample of data that include the challenges of genotypes imputed across multiple arrays; and second, all SNPs with moderately promising signals were genotyped for 75,000 individuals in a validation study to reveal the true risk status of many SNPs.

To assess the performance of UNICORN, we performed two experiments. (1) A direct comparison between UNICORN and an analysis of the full set of case and control subjects. In this experiment we learn whether UNICORN efficiently utilizes all the data in the control sample by comparing the power of the two tests. We do not expect UNICORN to have greater power, but we can determine whether it loses power compared to a direct analysis of the data. To determine whether UNICORN produces false positives, we permute case and control labels and look for deviations from the expected null distribution. (2) We mimic a realistic application of UNICORN by focusing on a particular study within the larger sample, complete with ancestry-matched case and control subjects. In this experiment we compare performance of a direct analysis of the matched case-control study to UNICORN applied to the same case subjects but with the full unselected sample of control subjects, excluding the matched controls.

Experiment 1

To obtain a baseline for power in the CD dataset, we performed a traditional logistic regression analysis on the full sample of case and control subjects, adjusting for ancestry via principal components (LRegr). For comparison, UNICORN used the full sample of controls to construct the MAFD for each case subject and then performed an association test using all case subjects. The results for the two methods were extremely similar (Figure 5A); notably, all SNPs that yielded significant results for LRegr ($p < 5 \times 10^{-8}$) also yielded significant results for UNICORN. This shows that in spite of the fact that UNICORN handled the control data only indirectly via the MAFD, it maintains full power to detect association signals. Moreover, each of the significant SNPs was also significant in the validation study.²⁸

To examine the overall validity of the tests, we computed the λ_{1000} genomic control factors^{29,30} and found both tests performed well: $\lambda = 1.03$ for UNICORN and $\lambda = 1.02$ for LRegr. To further evaluate the validity of UNICORN in the absence of polygenic effects, we permuted the case and control labels to remove association.³¹ The distribution of p values produced by UNICORN is well calibrated to meet null expectations (Figure 5B) and the genomic control factor for this distribution is $\lambda = 1.01$. In total this experiment shows that UNICORN makes efficient use of the full data without inducing false positives.

Finally, to illustrate the impact of each level of population structure, we analyzed these data three ways: (1) ignoring the effect of ancestry altogether; (2) modeling only the global structure using the first level of UNICORN; and (3) modeling the global and local structure with UNICORN. As expected, not accounting for ancestry leads to a P-P plot with strong evidence of overdispersion; incorporating the global level of UNICORN leads to a marked improvement in the distribution of test statistics; and finally, modeling additional structure at the local level leads to even greater reduction of false positives (Figure S2).

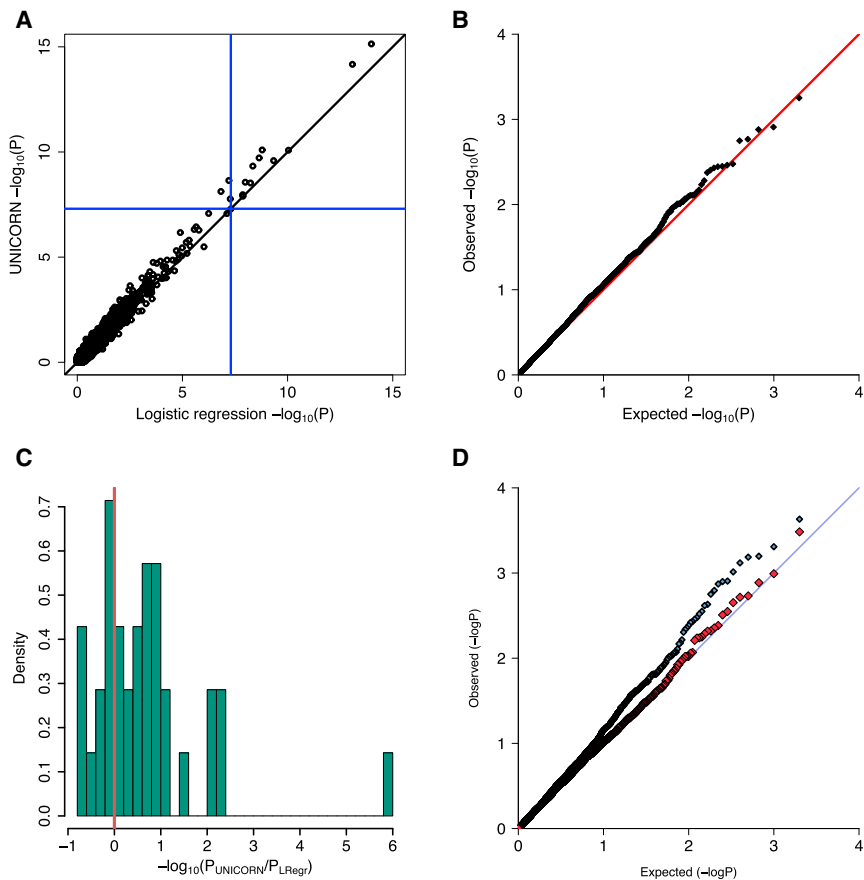


Figure 5. IBD Analysis via UNICORN versus Logistic Regression Controlling for Ancestry

(A) Comparison between UNICORN and LRegr on the full 7-study CD dataset. All significant SNPs detected by LRegr were also significant in UNICORN, and each of these SNPs was significant in the validation study as well.

(B) UNICORN null distribution obtained by permuting affection status in the full case-control dataset. The resulting distribution of p values produced by UNICORN is well calibrated, indicating a good control of false positives.

(C and D) UNICORN applied only to case subjects from the Belgian study using all control subjects excluding that study.

(C) Difference in p value magnitude between UNICORN and LRegr applied only to Belgian case-controls. Results are shown only for SNPs that were found significant in the validation study.²⁸ All SNPs showing a substantial difference favored UNICORN, particularly the SNP that had the highest signal in Jostins et al.²⁸

(D) P-P plot for UNICORN (blue) compared to the null distribution with permuted phenotype labels (red). The blue P-P plot shows some signal was detected and the red P-P plot shows that UNICORN yields an appropriate null distribution when there is no signal present.

Experiment 2

UNICORN is designed to permit analysis of a case-only sample by utilizing control subjects drawn from a repository. To evaluate performance in this setting, we extracted the IBD-CD Belgian study for further investigation. This study consists of a sample of 666 CD case subjects and 978 control subjects of similar ancestry. A case-only sample applying UNICORN would have access to all 14,927 control subjects minus the 978 Belgian control subjects. For comparison, we contrasted the results of analysis of this well-matched study using LRegr with UNICORN using all non-Belgian control subjects.

Not surprisingly, no SNP is genome-wide significant for this modest sample of case subjects for either analysis. To compare the power, we evaluated the behavior of the association tests for the 163 SNPs that showed genome-wide significance in the validation study.²⁸ Taking this as truth, we favor whichever method yields a smaller p value in the comparison (Figure 5C). For 70% of these loci, the evidence for association from the UNICORN analysis was stronger than the LRegr analysis, and for the 30% where it was not superior, the tests were nearly identical for both approaches with neither test showing a signal. One surprising result was that the most highly significant SNP in the validation study exhibited a p value six orders of magnitude smaller with UNICORN than LRegr (Figure 5C). This SNP has a relatively large F_{ST} .

Based on the P-P plot, UNICORN p values detect a modest signal for many SNPs. To assess the validity of the test, we permuted the case and control labels to remove association and found that the overall distribution of the UNICORN test was appropriate (Figure 5D, red). This experiment supports the great potential of UNICORN to increase power without incurring false positive findings.

Detection and Removal of False Positives

One of the major challenges in the analysis of genetic data is controlling for the technical variability across different SNP arrays, imputation pipelines, and genotyping approaches. This challenge is equally great when applying UNICORN; however, careful attention to process and quality control (QC) can greatly enhance the reliability of the analysis.

Ultimately, the UNICORN repository will consist of an assimilation of samples from tens, if not hundreds, of individual studies. Therefore it will certainly include multiple SNP arrays. To avoid exacerbating study-specific biases, all samples in the repository will be imputed via a common pipeline. As proof of concept, imputation was performed jointly for the CD control subjects used here, which stem from seven different studies and arrays. After first performing the QC procedures described below, no significant array bias was detected in the study.²⁸ The IBD study demonstrates that a homogeneous control collection can be

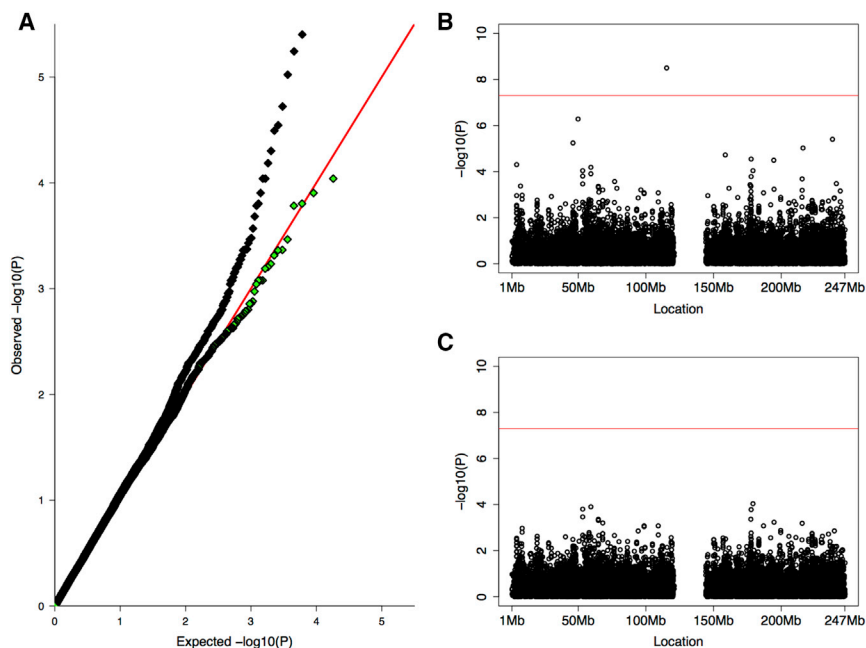


Figure 6. Detection and Removal of False Positives via Nonparametric Smoothing

We created a UNICORN study by using individuals selected for European ancestry from the HGDP dataset and comparing them to the CD controls. Any signals in this comparison are probably due to technical artifacts.

(A) P-P plot of UNICORN results before (black) and after (green) smoothing to reduce noise. Notice the strong presence of signal in the black P-P plot despite the expectation of no signal when comparing two control datasets.

(B) Manhattan plot of UNICORN p values before smoothing exhibits isolated signals without support in the immediate LD neighborhood.

(C) Isolated signals are removed after smoothing.

assembled from different sources, provided care is taken with the imputation and QC. Likewise, imputation on the case subjects should follow the same pipeline as implemented for UNICORN control subjects.

Based on our investigation of challenges due to imputation and array-based biases, we have identified a reliable approach that is quite similar to the technical QC and assay evaluation in use for routine genetic analysis. The objective is to identify SNPs with unusually small p values relative to their linkage disequilibrium (LD) neighbors. Such signals are almost always due to technical artifacts. SNPs with p values not supported by their LD neighbors can be identified using either nonparametric regression or a hidden Markov model via DIST.³² Both procedures successfully flag SNPs with outlier p values.

To illustrate this approach, we used individuals with European ancestry from the HGDP dataset as case subjects in comparison with the CD control subjects as part of a UNICORN association analysis. Any signal detected by such a test stems from technical artifacts and should be flagged as such. Prior to QC, UNICORN did indeed return some signals on chromosome 1 (Figures 6A and 6B). We ran a nonparametric kernel regression with binwidth of 2 Mbp across the series of $-\log_{10}$ p values and flagged results as noise if the smoothed value differed from the actual value by more than one order of magnitude. This procedure eliminated all the isolated signals (Figures 6A and 6C).

Another precaution can be employed to remove false positives due to differing SNP arrays. A comparison between UNICORN control subjects and control subjects measured on the same array as the case subjects should reveal SNPs that cannot be reliably compared across these arrays. Any SNPs exhibiting a signal in these experiments should be removed from further investigation. Moreover,

SNPs identified by internal comparisons across chips in development of the UNICORN repository will be noted on the UNICORN web site.

In conclusion, we note that similar to any large genetic association study, UNICORN can yield false positives due to technical artifacts. This challenge arises in part because UNICORN requires imputation in the control datasets to obtain a common set of SNPs across arrays for subsequent analyses. When the case sample is genotyped on an array that is not well represented in the control sample the challenge is greater; however, we have found that post analysis cleaning can remove false positives that arise.

Discussion

An essential feature of a genetic association study is a large control sample, chosen to represent the case sample in ancestry.^{33–35} Although suitable control samples can sometimes be obtained from public repositories, it typically requires substantial analytical effort to process control subjects along with the case sample. The goal of UNICORN is to obviate this need, at least partially, by automatically providing equivalent information from control subjects collected previously, for example controls deposited in dbGaP. For each case sample, the algorithm uses available control subjects to estimate the allele frequencies matched by ancestry. In this way UNICORN will facilitate case-control studies, even if the study has characterized only case samples, and thereby optimize the discovery of risk variants. By providing a control sample that is ancestrally matched to case subjects, without requiring resources or effort from the user, UNICORN provides advantages even for case-control studies for which a set of control samples has already been collected. In our proposed

implementation, the end user would not experience significant compute time because the MAFD can be pre-computed and easily queried based on the user's case subjects.

Not sampling controls as part of the study design precludes the direct inclusion of covariates in the analysis. In some settings a properly chosen covariate can greatly enhance power, whereas in other scenarios covariates can reduce power³⁶ or bias the analysis.³⁷ Even in the former setting when covariates are useful, UNICORN can provide a more powerful analysis due to the enhanced estimate of the population allele frequency derived from a much larger sample of controls (Figure S3). Moreover, if covariates are measured in the case subjects, it is possible to perform conditional genetic analysis on subsets of the data via UNICORN. Such an analysis contrasts the SNP allele frequencies in a subset of the case subjects with the estimated population allele frequency of ancestry-matched controls.

A population control sample by definition includes some individuals that should be classified as case subjects. This will reduce power, but it will not generate an excess of false positives, and the impact on power increases with the frequency of the disorder under investigation. When screened control subjects have not been collected, however, it is common practice to rely on population control subjects and UNICORN has no special weaknesses or strengths with regard to this issue.

The current analysis features samples of European ancestry, but the framework is applicable to other ancestries as well. Due to its multiple levels of inference, UNICORN can accommodate populations of quite complex structure, such as that found from African populations,³⁸ as well as the simpler structure of European populations. Our experiments suggest that UNICORN models ancestry as effectively as PCA, so we expect it will perform well in other ancestries. Analyses of recently admixed populations are more challenging, however, and will require new additions to the UNICORN methodology.

In addition to the potential gain in power, UNICORN also has the potential to strengthen subject privacy. The ability to identify an individual from their anonymous genetic information in a public database threatens the principle of subject confidentiality.³⁹ Knowledge of an individual's genotype at relatively few SNPs is sufficient to uniquely identify a person; indeed, this is the basis of DNA forensics. Protected repositories such as dbGaP exist so that genome-wide data can be shared among responsible parties without exposing subjects to a loss of privacy. But a second level of privacy loss is also of concern. Based on reported allele frequencies in case and control subjects, given a very large number of SNPs, it is possible to determine with high probability whether an individual is a case or control subject in the study, or not in the study at all.^{40–42} By restricting the exchange of genetic data to ancestry coordinates, UNICORN could overcome both of these challenges. Additionally, our grid-based approach, where we return frequency estimates from the

pre-computed grid point closest to the case subject instead of the actual case location, provides another layer of security for the control identities by adding a small degree of randomness to our predictions.

Results from the Exome Aggregation Consortium (ExAC) motivate this work. ExAC has made substantial progress toward the goal of assembling exome sequencing data from a variety of large-scale sequencing projects to make summary data widely available, with more than 60,000 individuals' worth of data available.⁴³ Currently, ExAC provides allele frequency information for these samples. Through UNICORN, we aim to enhance this concept by generating ancestry-matched MAFD estimates for additional subjects. Although there are more technical challenges involved with sequence data than genotyping arrays, the ExAC project provides support to the belief that these data can be successfully aggregated and harmonized for use in UNICORN. We are thus currently in the process of extending the UNICORN framework, which will require further development and refinement for rare variation.

The UNICORN database and web server are in preparation and a limited version focusing on populations of European descent is slated for release late in 2016. If successful, UNICORN will dramatically improve access to the control resources stored in repositories such as dbGaP and can also make use of control samples from the same study as well as from other studies. We predict that UNICORN will hasten the discovery of genetic variation conferring risk for disease in three ways: by providing ancestrally matched allele frequencies, by its careful integration of datasets, and by making genetic association analysis simpler.

Appendix A.

The covariance between two points of the Gaussian process at distance u is $\sigma^2\rho(u) = \sigma^2/e^{u\phi}$. Notice that ϕ is the characteristic length-scale of our process: it determines how far apart two individuals must be for the allele frequency to change significantly.

Inference in this model is performed via MCMC. Write $\mathbf{S} = [S(x_1), \dots, S(x_n)]$ for the vector of values of S at the observed locations x_i and $\mathbf{S}^* = [S(x_1^*), \dots, S(x_n^*)]$ for the vector of values of S at the target locations x_i^* for which predictions are requested. Define \mathbf{P} and \mathbf{P}^* similarly and let \mathbf{Y} be the genotype data at the observed locations.

A cycle of the MCMC algorithm involves first sampling from $(\sigma^2, \phi) | (\mathbf{Y}, \mathbf{S}, \beta)$, then from $S_i | (\mathbf{S}_{-i}, \mathbf{Y}, \sigma^2, \phi, \beta)$, and finally from $\beta | (\mathbf{Y}, \mathbf{S}, \sigma^2, \phi)$. Here, \mathbf{S}_{-i} denotes the vector \mathbf{S} without its i^{th} element. Note that because conditionally on S the random variables Y_i are mutually independent, we have:

$$p(\mathbf{Y} | \mathbf{S}, \beta) = \prod_{j=1}^n f(y_j | s_j, \beta). \quad (\text{Equation A1})$$

Using Equation A1 we have:

$$p((\sigma^2, \phi) | \mathbf{Y}, \mathbf{S}, \beta) = p((\sigma^2, \phi) | \mathbf{S}) \propto p(\mathbf{S} | \sigma^2, \phi) p(\sigma^2, \phi) \quad (\text{Equation A2})$$

$$p(S_i | \mathbf{S}_{-i}, \mathbf{Y}, \sigma^2, \phi, \beta) \propto p(\mathbf{Y} | \mathbf{S}, \beta) p(S_i | \mathbf{S}_{-i}, \sigma^2, \phi) \\ = p(S_i | \mathbf{S}_{-i}, \sigma^2, \phi) \prod_{j=1}^n f(y_j | s_j, \beta) \quad (\text{Equation A3})$$

$$p(\beta | (\mathbf{Y}, \mathbf{S}, \sigma^2, \phi)) = p(\beta | (\mathbf{Y}, \mathbf{S})) \propto p(\mathbf{Y} | \mathbf{S}, \beta) p(\beta) \\ = p(\beta) \prod_{j=1}^n f(y_j | s_j, \beta). \quad (\text{Equation A4})$$

$$\min \left\{ 1, \frac{\left[\prod_{j=1}^{i-1} f(y_j | s_j, \beta) \right] \cdot f(y_i | s'_i, \beta) \cdot \left[\prod_{j=i+1}^n f(y_j | s_j, \beta) \right] \cdot p(S'_i | \mathbf{S}_{-i}, \sigma^2, \phi) \cdot q(S_i | S'_i)}{\left[\prod_{j=1}^{i-1} f(y_j | s_j, \beta) \right] \cdot f(y_i | s_i, \beta) \cdot \left[\prod_{j=i+1}^n f(y_j | s_j, \beta) \right] \cdot p(S_i | \mathbf{S}_{-i}, \sigma^2, \phi) \cdot q(S'_i | S_i)} \right\} = \quad (\text{Equation A9})$$

From the Gaussian process assumption, it follows that $p(\mathbf{S} | \sigma^2, \phi)$ has a multivariate normal density (mean 0 and covariance matrix $\sigma^2 e^{-U\phi}$ where U is the Euclidean distance matrix for the locations referred to by \mathbf{S}) and $p(S_i | \mathbf{S}_{-i}, \sigma^2, \phi)$ has a univariate normal distribution. Also, we know that $f(y_j | s_j, \beta)$ follows a binomial distribution (with success probability $e^{\beta+s_j}/1 + e^{\beta+s_j}$) and $p(\beta)$ and $p(\sigma^2, \phi)$ are the priors. Being able to draw from all these distributions enables us to apply the following component-wise Metropolis-Hastings algorithm.

1. Set initial values of β , σ^2 , and ϕ by drawing from their respective priors. Set the starting value for each S_i to 0.
2. Update (σ^2, ϕ)

- choose a new value (σ'^2, ϕ') from some appropriate proposal distribution $q((\sigma'^2, \phi') | (\sigma^2, \phi))$
- using Equation A2, accept (σ'^2, ϕ') with probability

$$\min \left\{ 1, \frac{p[(\sigma'^2, \phi') | \mathbf{Y}, \mathbf{S}, \beta] \cdot q[(\sigma^2, \phi) | (\sigma'^2, \phi')]}{p[(\sigma^2, \phi) | \mathbf{Y}, \mathbf{S}, \beta] \cdot q[(\sigma'^2, \phi') | (\sigma^2, \phi)]} \right\} = \quad (\text{Equation A5})$$

$$\min \left\{ 1, \frac{p(\mathbf{S} | \sigma'^2, \phi') \cdot p(\sigma'^2, \phi') \cdot q[(\sigma^2, \phi) | (\sigma'^2, \phi')]}{p(\mathbf{S} | \sigma^2, \phi) \cdot p(\sigma^2, \phi) \cdot q[(\sigma'^2, \phi') | (\sigma^2, \phi)]} \right\} = \quad (\text{Equation A6})$$

$$\min \left\{ 1, \frac{p(\mathbf{S} | \sigma'^2, \phi') \cdot p(\sigma'^2) \cdot p(\phi') \cdot q_{\sigma^2}(\sigma^2 | \sigma'^2) \cdot q_{\phi}(\phi | \phi')}{p(\mathbf{S} | \sigma^2, \phi) \cdot p(\sigma^2) \cdot p(\phi) \cdot q_{\sigma^2}(\sigma^2 | \sigma^2) \cdot q_{\phi}(\phi' | \phi)} \right\}, \quad (\text{Equation A7})$$

where the last equality holds if σ^2 and ϕ are independent.

- the prior distributions for σ^2 and ϕ as well as the jumping distributions q_{σ^2} and q_{ϕ} can be gammas.

3. Update \mathbf{S}

- choose a new value S'_i for the i^{th} component of \mathbf{S} from the transition probability function $q(S'_i | S_i) = p(S'_i | \mathbf{S}_{-i}, \sigma^2, \phi)$
- using Equation A3, accept S'_i with probability

$$\min \left\{ 1, \frac{p(S'_i | \mathbf{S}_{-i}, \mathbf{Y}, \sigma^2, \phi, \beta) \cdot q(S_i | S'_i)}{p(S_i | \mathbf{S}_{-i}, \mathbf{Y}, \sigma^2, \phi, \beta) \cdot q(S'_i | S_i)} \right\} = \quad (\text{Equation A8})$$

$$\min \left\{ 1, \frac{f(y_i | s'_i, \beta)}{f(y_i | s_i, \beta)} \right\}. \quad (\text{Equation A10})$$

- repeat the previous two steps for all $i = 1, \dots, n$ to complete updating \mathbf{S}

4. Update β

- choose a new value β' from some appropriate proposal distribution $q(\beta' | \beta)$
- using Equation A4, accept β' with probability

$$\min \left\{ 1, \frac{p(\beta' | (\mathbf{Y}, \mathbf{S}, \sigma^2, \phi)) \cdot q(\beta | \beta')}{p(\beta | (\mathbf{Y}, \mathbf{S}, \sigma^2, \phi)) \cdot q(\beta' | \beta)} \right\} = \quad (\text{Equation A11})$$

$$\min \left\{ 1, \frac{\prod_{j=1}^n f(y_j | s_j, \beta') \cdot p(\beta') \cdot q(\beta | \beta')}{\prod_{j=1}^n f(y_j | s_j, \beta) \cdot p(\beta) \cdot q(\beta' | \beta)} \right\}. \quad (\text{Equation A12})$$

- the prior distribution of β is determined by the cluster-wide allele frequency inference step, and the jumping distribution q can be a normal distribution. Repeat steps 2–4 (with an optional burn-in period and thinning) to obtain draws from the equilibrium distributions. We are now able to draw from the posteriors of $\sigma^2, \phi, \beta, \mathbf{S}$. We proceed with:

5. Draw a sample from the multivariate Gaussian distribution of $\mathbf{S}^* | (\mathbf{S}, \mathbf{Y}, \sigma^2, \phi, \beta)$ where the values of $\mathbf{S}, \sigma^2, \phi, \beta$, are those generated in steps 2–4. Using the

conditional independence structure of our model, this step reduces to drawing from $\mathbf{S}^* | (\mathbf{S}, \sigma^2, \phi)$. The Gaussian process assumption implies that:

$$\mathbf{S}^* | (\mathbf{S}, \sigma^2, \phi) \sim MVN(\Sigma_{12}^T \Sigma_{11}^{-1} \mathbf{S}, \Sigma_{22} - \Sigma_{12}^T \Sigma_{11}^{-1} \Sigma_{12}),$$

(Equation A13)

where

$$\Sigma_{11} = \text{var}(\mathbf{S})$$

(Equation A14)

$$\Sigma_{12} = \text{cov}(\mathbf{S}, \mathbf{S}^*)$$

(Equation A15)

$$\Sigma_{22} = \text{var}(\mathbf{S}^*).$$

(Equation A16)

Each of these matrices can be computed based on the variance properties defined by σ^2 and ϕ .

6. Compute \mathbf{P}^* based on the current values of \mathbf{S}^* and β :

$$P(x_i^*) = \frac{e^{\beta + S(x_i^*)}}{1 + e^{\beta + S(x_i^*)}}.$$

(Equation A17)

Iterating steps 5 and 6 gives us the predictive distribution $P(x_i^*)$ for all points at which we want to infer allele frequencies.

Supplemental Data

Supplemental Data include three figures and a list of consortia authors and can be found with this article online at <http://dx.doi.org/10.1016/j.ajhg.2016.02.025>.

Acknowledgments

This work was supported by the National Institute of Mental Health grant R37MH057881 (B.D. and K.R.) and NIH grant MH101244-02 (B.M.N.). We thank the referees for insightful comments.

Received: October 7, 2015

Accepted: February 29, 2016

Published: April 14, 2016

Web Resources

The URLs for data presented herein are as follows:

dbGaP, <http://www.ncbi.nlm.nih.gov/gap>

ExAC Browser, <http://exac.broadinstitute.org/>

GemTools, <http://wpcr.wpic.pitt.edu/WPICCompGen/GemTools/GemTools.htm>

References

- Mailman, M.D., Feolo, M., Jin, Y., Kimura, M., Tryka, K., Bagoutdinov, R., Hao, L., Kiang, A., Paschall, J., Phan, L., et al. (2007). The NCBI dbGaP database of genotypes and phenotypes. *Nat. Genet.* *39*, 1181–1186.
- Koike, A., Nishida, N., Inoue, I., Tsuji, S., and Tokunaga, K. (2009). Genome-wide association database developed in the Japanese Integrated Database Project. *J. Hum. Genet.* *54*, 543–546.
- Spielman, R.S., McGinnis, R.E., and Ewens, W.J. (1993). Transmission test for linkage disequilibrium: the insulin gene region and insulin-dependent diabetes mellitus (IDDM). *Am. J. Hum. Genet.* *52*, 506–516.
- Lange, C., and Laird, N.M. (2002). Power calculations for a general class of family-based association tests: dichotomous traits. *Am. J. Hum. Genet.* *71*, 575–584.
- Bacanu, S.-A., Devlin, B., and Roeder, K. (2000). The power of genomic control. *Am. J. Hum. Genet.* *66*, 1933–1944.
- Risch, N. (2001). Implications of multilocus inheritance for gene-disease association studies. *Theor. Popul. Biol.* *60*, 215–220.
- Ferreira, M.A., Sham, P., Daly, M.J., and Purcell, S. (2007). Ascertainment through family history of disease often decreases the power of family-based association studies. *Behav. Genet.* *37*, 631–636.
- Klei, L., Sanders, S.J., Murtha, M.T., Hus, V., Lowe, J.K., Willsey, A.J., Moreno-De-Luca, D., Yu, T.W., Fombonne, E., Geschwind, D., et al. (2012). Common genetic variants, acting additively, are a major source of risk for autism. *Mol. Autism* *3*, 9.
- Gaugler, T., Klei, L., Sanders, S.J., Bodea, C.A., Goldberg, A.P., Lee, A.B., Mahajan, M., Manaa, D., Pawitan, Y., Reichert, J., et al. (2014). Most genetic risk for autism resides with common variation. *Nat. Genet.* *46*, 881–885.
- Lee, A.B., Luca, D., and Roeder, K. (2010). A spectral graph approach to discovering genetic ancestry. *Ann. Appl. Stat.* *4*, 179–202.
- Price, A.L., Patterson, N.J., Plenge, R.M., Weinblatt, M.E., Shadick, N.A., and Reich, D. (2006). Principal components analysis corrects for stratification in genome-wide association studies. *Nat. Genet.* *38*, 904–909.
- Patterson, N., Price, A.L., and Reich, D. (2006). Population structure and eigenanalysis. *PLoS Genet.* *2*, e190.
- Klei, L., Kent, B.P., Melhem, N., Devlin, B., and Roeder, K. (2011). Gemtools: a fast and efficient approach to estimating genetic ancestry. *arXiv:1104.1162*. <http://arxiv.org/abs/0901.0633v2>.
- Lee, A.B., Luca, D., Klei, L., Devlin, B., and Roeder, K. (2010). Discovering genetic ancestry using spectral graph theory. *Genet. Epidemiol.* *34*, 51–59.
- Nelson, M.R., Bryc, K., King, K.S., Indap, A., Boyko, A.R., Novembre, J., Briley, L.P., Maruyama, Y., Waterworth, D.M., Waiber, G., et al. (2008). The Population Reference Sample, POPRES: a resource for population, disease, and pharmacological genetics research. *Am. J. Hum. Genet.* *83*, 347–358.
- Novembre, J., Johnson, T., Bryc, K., Kutalik, Z., Boyko, A.R., Auton, A., Indap, A., King, K.S., Bergmann, S., Nelson, M.R., et al. (2008). Genes mirror geography within Europe. *Nature* *456*, 98–101.
- Yang, W.-Y., Novembre, J., Eskin, E., and Halperin, E. (2012). A model-based approach for analysis of spatial structure in genetic data. *Nat. Genet.* *44*, 725–731.
- Epstein, M.P., Allen, A.S., and Satten, G.A. (2007). A simple and improved correction for population stratification in case-control studies. *Am. J. Hum. Genet.* *80*, 921–930.
- Engelhardt, B.E., and Stephens, M. (2010). Analysis of population structure: a unifying framework and novel methods based on sparse factor analysis. *PLoS Genet.* *6*, e1001117.

20. Wang, C., Zhan, X., Liang, L., Abecasis, G.R., and Lin, X. (2015). Improved ancestry estimation for both genotyping and sequencing data using projection procrustes analysis and genotype imputation. *Am. J. Hum. Genet.* *96*, 926–937.
21. Crossett, A., Kent, B.P., Klei, L., Ringquist, S., Trucco, M., Roeder, K., and Devlin, B. (2010). Using ancestry matching to combine family-based and unrelated samples for genome-wide association studies. *Stat. Med.* *29*, 2932–2945.
22. Bengio, Y., Delalleau, O., Le Roux, N., Paiement, J.-F., Vincent, P., and Ouimet, M. (2004). Learning eigenfunctions links spectral embedding and kernel PCA. *Neural Comput.* *16*, 2197–2219.
23. Williams, C., and Seeger, M. (2001). Using the nyström method to speed up kernel machines. In *Proceedings of the 14th Annual Conference on Neural Information Processing Systems* number EPFL-CONF-161322. pp. 682–688.
24. Kumar, S., Mohri, M., and Talwalkar, A. (2012). Sampling methods for the nyström method. *J. Mach. Learn. Res.* *13*, 981–1006.
25. Li, J.Z., Absher, D.M., Tang, H., Southwick, A.M., Casto, A.M., Ramachandran, S., Cann, H.M., Barsh, G.S., Feldman, M., Cavalli-Sforza, L.L., and Myers, R.M. (2008). Worldwide human relationships inferred from genome-wide patterns of variation. *Science* *319*, 1100–1104.
26. Diggle, P.J., Tawn, J., and Moyeed, R. (1998). Model-based geostatistics. *J. R. Stat. Soc. Ser. C Appl. Stat.* *47*, 299–350.
27. Heath, S.C., Gut, I.G., Brennan, P., McKay, J.D., Bencko, V., Fabianova, E., Foretova, L., Georges, M., Janout, V., Kabisch, M., et al. (2008). Investigation of the fine structure of European populations with applications to disease association studies. *Eur. J. Hum. Genet.* *16*, 1413–1429.
28. Jostins, L., Ripke, S., Weersma, R.K., Duerr, R.H., McGovern, D.P., Hui, K.Y., Lee, J.C., Schumm, L.P., Sharma, Y., Anderson, C.A., et al.; International IBD Genetics Consortium (IIBDGC) (2012). Host-microbe interactions have shaped the genetic architecture of inflammatory bowel disease. *Nature* *491*, 119–124.
29. Devlin, B., and Roeder, K. (1999). Genomic control for association studies. *Biometrics* *55*, 997–1004.
30. Freedman, M.L., Reich, D., Penney, K.L., McDonald, G.J., Mignault, A.A., Patterson, N., Gabriel, S.B., Topol, E.J., Smoller, J.W., Pato, C.N., et al. (2004). Assessing the impact of population stratification on genetic association studies. *Nat. Genet.* *36*, 388–393.
31. Yang, J., Weedon, M.N., Purcell, S., Lettre, G., Estrada, K., Willet, C.J., Smith, A.V., Ingelsson, E., O’Connell, J.R., Mangino, M., et al.; GIANT Consortium (2011). Genomic inflation factors under polygenic inheritance. *Eur. J. Hum. Genet.* *19*, 807–812.
32. Lee, D., Bigdeli, T.B., Riley, B.P., Fanous, A.H., and Bacanu, S.-A. (2013). DIST: direct imputation of summary statistics for unmeasured SNPs. *Bioinformatics* *29*, 2925–2927.
33. Lander, E.S., and Schork, N.J. (1994). Genetic dissection of complex traits. *Science* *265*, 2037–2048.
34. Cardon, L.R., and Palmer, L.J. (2003). Population stratification and spurious allelic association. *Lancet* *361*, 598–604.
35. Luca, D., Ringquist, S., Klei, L., Lee, A.B., Gieger, C., Wichmann, H.E., Schreiber, S., Krawczak, M., Lu, Y., Styche, A., et al. (2008). On the use of general control samples for genome-wide association studies: genetic matching highlights causal variants. *Am. J. Hum. Genet.* *82*, 453–463.
36. Pirinen, M., Donnelly, P., and Spencer, C.C.A. (2012). Including known covariates can reduce power to detect genetic effects in case-control studies. *Nat. Genet.* *44*, 848–851.
37. Aschard, H., Vilhjálmsson, B.J., Joshi, A.D., Price, A.L., and Kraft, P. (2015). Adjusting for heritable covariates can bias effect estimates in genome-wide association studies. *Am. J. Hum. Genet.* *96*, 329–339.
38. Tishkoff, S.A., Reed, F.A., Friedlaender, F.R., Ehret, C., Ranciaro, A., Froment, A., Hirbo, J.B., Awomoyi, A.A., Bodo, J.-M., Doumbo, O., et al. (2009). The genetic structure and history of Africans and African Americans. *Science* *324*, 1035–1044.
39. Kaye, J., Boddington, P., de Vries, J., Hawkins, N., and Melham, K. (2010). Ethical implications of the use of whole genome methods in medical research. *Eur. J. Hum. Genet.* *18*, 398–403.
40. Homer, N., Szlinger, S., Redman, M., Duggan, D., Tembe, W., Muehling, J., Pearson, J.V., Stephan, D.A., Nelson, S.F., and Craig, D.W. (2008). Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genet.* *4*, e1000167.
41. Jacobs, K.B., Yeager, M., Wacholder, S., Craig, D., Kraft, P., Hunter, D.J., Paschal, J., Manolio, T.A., Tucker, M., Hoover, R.N., et al. (2009). A new statistic and its power to infer membership in a genome-wide association study using genotype frequencies. *Nat. Genet.* *41*, 1253–1257.
42. Im, H.K., Gamazon, E.R., Nicolae, D.L., and Cox, N.J. (2012). On sharing quantitative trait GWAS results in an era of multiple-omics data and the limits of genomic privacy. *Am. J. Hum. Genet.* *90*, 591–598.
43. Lek, M., Karczewski, K., Minikel, E., Samocha, K., Banks, E., Fennell, T., O’Donnell-Luria, A., Ware, J., Hill, A., Cummings, B., et al. (2015). Analysis of protein-coding genetic variation in 60,706 humans. *bioRxiv* 10.1101/030338.