



Ultrafast Comparison of Personal Genomes via Precomputed Genome Fingerprints

Gustavo Glusman*, Denise E. Mauldin, Leroy E. Hood and Max Robinson

Institute for Systems Biology, Seattle, WA, United States

OPEN ACCESS

Edited by:

Max A. Alekseyev,
George Washington University,
United States

Reviewed by:

Filippo Geraci,
Consiglio Nazionale delle Ricerche
(CNR), Italy
Pavel Loskot,
Swansea University, United Kingdom
Son Pham,
Salk Institute for Biological Studies,
United States
Kui Lin,
Beijing Normal University, China

*Correspondence:

Gustavo Glusman
gustavo@systemsbiology.org

Specialty section:

This article was submitted to
Bioinformatics and Computational
Biology,
a section of the journal
Frontiers in Genetics

Received: 14 July 2017

Accepted: 12 September 2017

Published: 26 September 2017

Citation:

Glusman G, Mauldin DE, Hood LE
and Robinson M (2017) Ultrafast
Comparison of Personal Genomes via
Precomputed Genome Fingerprints.
Front. Genet. 8:136.
doi: 10.3389/fgene.2017.00136

We present an ultrafast method for comparing personal genomes. We transform the standard genome representation (lists of variants relative to a reference) into “genome fingerprints” via locality sensitive hashing. The resulting genome fingerprints can be meaningfully compared even when the input data were obtained using different sequencing technologies, processed using different pipelines, represented in different data formats and relative to different reference versions. Furthermore, genome fingerprints are robust to up to 30% missing data. Because of their reduced size, computation on the genome fingerprints is fast and requires little memory. For example, we could compute all-against-all pairwise comparisons among the 2504 genomes in the 1000 Genomes data set in 67 s at high quality (21 μ s per comparison, on a single processor), and achieved a lower quality approximation in just 11 s. Efficient computation enables scaling up a variety of important genome analyses, including quantifying relatedness, recognizing duplicative sequenced genomes in a set, population reconstruction, and many others. The original genome representation cannot be reconstructed from its fingerprint, effectively decoupling genome comparison from genome interpretation; the method thus has significant implications for privacy-preserving genome analytics.

Keywords: computational genomics, algorithms, genome comparison, study design, population genetics, privacy

INTRODUCTION

Personal genome sequences contain the information required for assessing genetic risks, matching genetic backgrounds between cases and controls in medical research, and detecting duplicate individuals or close relatives for medical, legal, or historical reasons. Research purposes served by personal genome sequencing include classifying individuals by population, reconstructing human history, assessing and controlling the quality of the sequence information itself, computing kinship matrices to support genome-wide association studies, and combining data sets for meta-analysis.

Many of these applications involve comparison of two or more personal genomes. However, the size, complexity, and diversity of representations in which they are stored makes comparison of personal genomes in their existing forms error-prone and slow, and therefore challenging to scale from pairs to the hundreds, thousands, or millions of individuals we will soon wish to compare in order to provide improved, personalized medical care.

Comparison of personal genomes requires cross-referencing on a very large number (potentially millions) of variants, most of which are single-nucleotide variants (SNVs); this process can be slow enough to make the computation of all pairwise comparisons take a prohibitive amount

of time. Current approaches to rapid genome comparison rely on severely limiting the number of single nucleotide polymorphisms (SNPs, defined as SNVs with population frequency above 1%) to compare. A carefully curated set of 24 autosomal exome SNPs (Pengelly et al., 2013) captures up to 38 bits of information per genome and thus limits utility to establishing identity and is susceptible to failure when data are partially missing. A larger set of 1500 pre-selected SNPs has been used to match genomes directly from BAM files (Wang et al., 2016), with the requirement that these be mapped to the same reference genome. Similarly, focusing on tens of thousands of pre-selected SNVs enables multiple fast quality control computations in pedigrees (Pedersen and Quinlan, 2017). Another powerful strategy involves deep compression of the data, coupled with advanced algorithms for fast querying and retrieval (Layer et al., 2016; Tatwawadi et al., 2016). The recent UNICORN method enables independent mapping of individuals onto ancestry spaces based on detailed maps of minor allele frequency distributions (Bodea et al., 2016). These strategies offer better speed at the expense of limited applicability, severely reduced accuracy, or strong reliance on detailed prior knowledge about the population being studied.

Perhaps more importantly, the diversity of representations of personal genomes remains a problem for each of these methods, which must translate variants represented in different ways to a common representation in order to assemble the data to compare. This diversity of representations arises from the variety of technologies used to determine personal genomes, their processing with diverse software pipelines and their representation relative to multiple versions of the reference genome. We explore each of these below.

The billions of base pairs in an individual's genome can be assayed using a variety of techniques, the most complete of which is whole-genome sequencing (WGS). Several WGS technologies have been used to determine individual genomic sequences, differing in experimental protocols and in specific parameters such as read length and error tolerances; the resulting data are processed using diverse computational pipelines. Each technology brings a different set of observational biases and error modalities; the resulting observations are expressed in a diversity of "native" vendor-specific file formats, and translated to one or more standard formats [e.g., variant call format (VCF)] in a variety of ways, leading to a bewildering diversity of flavors of semi-standard formats.

The resulting sequence is compactly represented as a list of differences (variants) from a reference genome. Due to the ongoing progress in improving the quality of the human reference genome, there are now several different versions (also known as "freezes") in use. To date, many thousands of personal genomes have been ascertained, most of them expressed relative to the GRCh37 (also known as "hg19") and GRCh38 ("hg38") freezes, though much legacy data exist using GRCh36 ("hg18") coordinates and even earlier versions. The position of a genomic variant depends on the reference version used: from one genome freeze to another, typically only the position of each variant along the chromosome changes, though in a small fraction of the cases a variant may be located in different chromosomes. Any two genomes of interest may be expressed relative to different

reference genomes, necessitating conversion prior to comparison and analysis.

Even relative to the same reference version, a variant may have more than one representation, necessitating additional conversion steps potentially leading to errors. Each chromosome has multiple names (e.g., "NC_000003.12," "chr3" or just "3" are all used refer to the third largest human chromosome), and different representations count the first nucleotide of each chromosome as position 0 or 1. Even with the same naming and numbering conventions, some variants can be expressed in more than one way, requiring normalization (Tan et al., 2015).

Sequencing the same genome using different technologies can yield differing results, as each technology has its own biases. Even when using the same technology, reference and encoding, sequencing the same genome repeatedly can give somewhat different results due to the stochastic nature of genome sequencing, to batch effects, or to differences in the computational pipelines used.

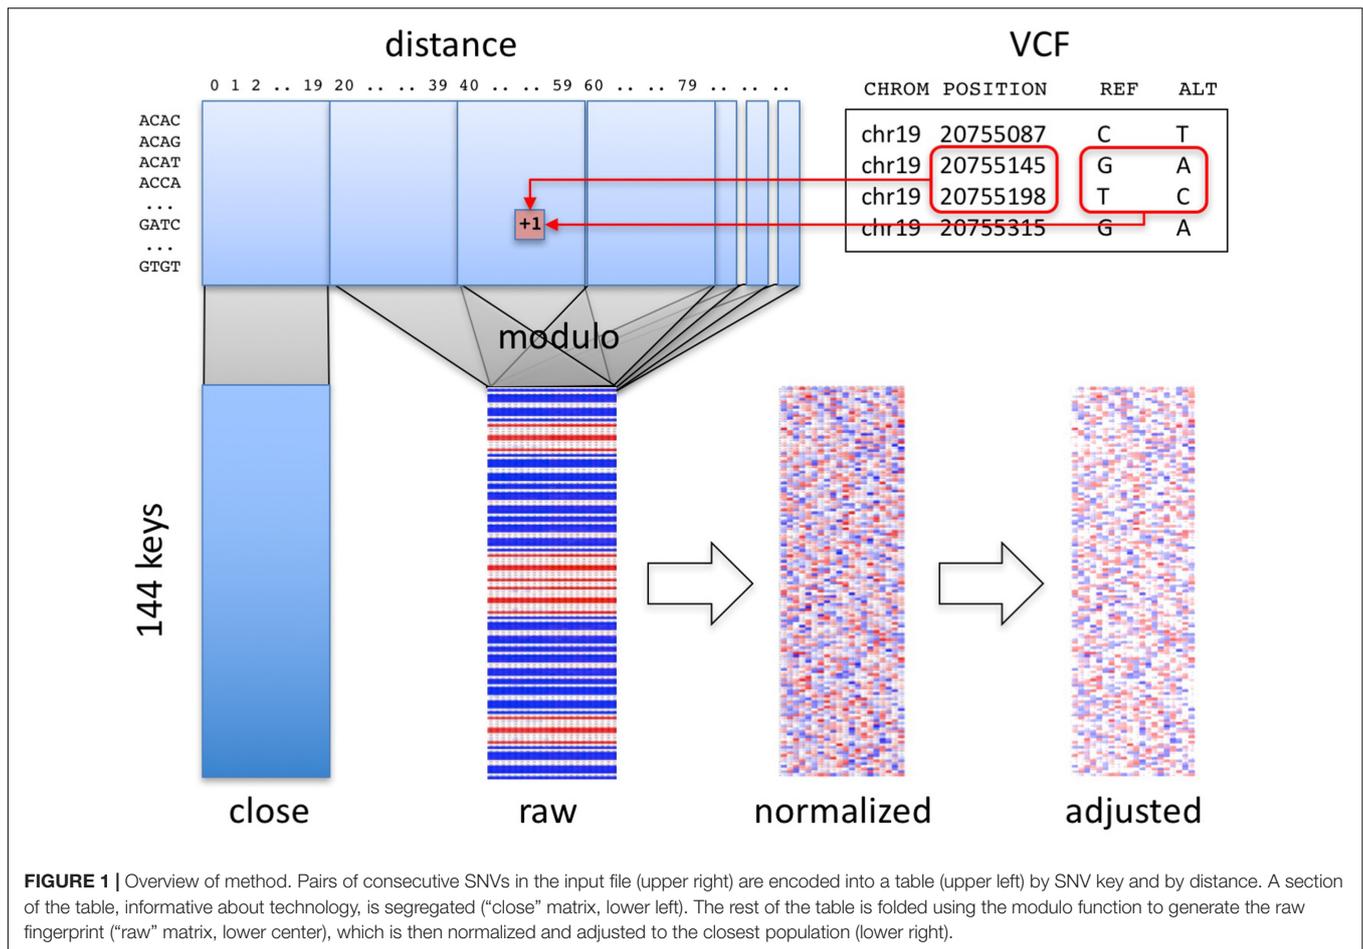
Looking to the future, additional (long-read) technologies will enable *de novo* genome assembly to become commonplace; the reference genome representation will change to a graph format, further breaking the concept of absolute coordinates; and the number of genomes available will soon be in the millions. These trends will deepen, not lessen the complexity of genome comparison.

We present here a new, rapid method for summarizing personal genomes that does not require knowledge of the technology, reference and encoding used, and yields "genome fingerprints" that can be used to facilitate many standard problems in genomics. Thanks to their reduced size, computation on the genome fingerprints is orders of magnitude faster and requires little memory, enabling comparison of much larger sets of genomes. No individual variants or other detailed features of the personal genome can be reconstructed from the fingerprint, thereby allowing private information to be more closely guarded and protected by *decoupling genome comparison from genome interpretation*. Fingerprints of different sizes allow balancing the speed and accuracy of the comparisons, and due to the high value of estimating relatedness, the potential applications of genome fingerprinting range from basic science (study design, population studies) to personalized medicine, forensics and genealogy research.

METHODS

Overview

Our algorithm summarizes a personal genome as a "fingerprint" (Figure 1). Conceptually, any variant-oriented representation of a personal genome (e.g., a VCF file or vendor-specific format, regardless of encoding or reference genome used) includes a list of variants, including their position and reference and alternative alleles, sorted by position. A "raw" fingerprint is a tally of consecutive biallelic SNV pairs grouped on a combination of these two attributes. We then normalize the raw fingerprint to account for systematic differences in frequency between groups by allele and by position. The resulting "normalized"



fingerprint preserves differences at the species level, e.g., between individuals from different populations. Averaging the normalized fingerprints of the individuals in a population yields a “population” fingerprint, which can be subtracted from an individual’s normalized fingerprint to produce a “population-adjusted” fingerprint suitable for more sensitive detection of related genomes.

Computation of Raw Fingerprints

The first stage in computing genome fingerprints yields a “raw” fingerprint, a $144 \times L$ table of SNV pair counts (L , the fingerprint length, is the main parameter of the method, defaulting to 20). We classify each pair of consecutive SNVs by the combination of reference and alternate alleles at both SNVs: this information determines the row in the table. We also consider the distance between the SNVs: this information determines the column in the table. When studying genomes produced using Complete Genomics and Illumina technologies, we observed differences in the distribution of distances between consecutive SNVs, for short distances (Supplementary Figure 1). These differences largely reflect each vendor pipeline’s encoding of multi-nucleotide variants; some pipelines represent these as individual SNVs, potentially leading to interpretation errors (Wei et al., 2015). We thus exclude SNV pairs that are separated by fewer than C

base pairs (defaulting to 20) from the final fingerprint. To avoid technology-specific artifacts, we recommend C at least 5 (see below).

To compute a raw fingerprint:

- (1) Parse the autosomes in the personal genome (e.g., from a VCF file) to identify biallelic SNVs as variants with reference allele of length 1 and single alternate allele of length 1, both in the case-insensitive alphabet [A C G T]. Ignore all other types of variants: two biallelic SNVs separated by other types of variants are still considered consecutive.
- (2) For each SNV, join its reference and alternative alleles into one of 12 possible SNV keys. For example, reference allele “G” and alternative allele “A” form key “GA.”
- (3) Join the keys of consecutive SNVs on the same chromosome to form one of 144 possible SNV pair keys; for example, consecutive SNVs with keys “GA” and “TC” form the SNV pair key “GATC.”
- (4) Compute the number of base pairs between the two SNVs (the “distance” between the SNVs along the reference genome). If this distance is smaller than C , increment by one the corresponding value in the $[144 \times C]$ “close” matrix (Figure 1), and skip the next steps.

- (5) Compute a reduced distance as the distance between the two SNVs, modulo the L .
- (6) Increment by one the corresponding value in the $[144 \times L]$ “raw” matrix.

Normalization of Fingerprints

The second stage in computing genome fingerprints performs a normalization of the raw fingerprint to account for the different frequencies of transitions and transversions, distance and parameter effects (Supplementary Figure 2).

The normalization is performed in two steps:

- (1) Normalization by distance: subtract the mean and divide by the standard deviation of each column.
- (2) Normalization by SNV pair key: subtract the mean and divide by the standard deviation of each row.

The normalized fingerprints do not show any remaining structure (Supplementary Figure 3). Normalizing fingerprints in the reverse order (first by SNV pair key, then by distance) fails to correct for the internal structure of raw fingerprints (Supplementary Figure 4).

Population Fingerprints

We compute a *population fingerprint* as the average of the normalized fingerprints from the individuals in the population. These fingerprints must have been computed using the same parameters (C and L).

Adjustment to Population

We compute a *population-adjusted fingerprint* for an individual by subtracting a population fingerprint from the normalized fingerprint of that individual. The individual fingerprint and the population fingerprint must have been computed using the same parameters (C and L). Similar to normalized fingerprints, population-adjusted fingerprints show no internal structure (Supplementary Figure 5).

Fingerprint Comparison

To compare two fingerprints, concatenate the rows of each fingerprint matrix into a vector and compute the Spearman correlation between the two vectors. The fingerprints to be compared thus have to be computed using the same parameter L . This same procedure is appropriate for comparing two normalized fingerprints or two population adjusted fingerprints, whether adjusted to the same or different populations.

Binary Fingerprints

We also evaluated a minimal version of the genome fingerprints, in the shape of a binary string of length 144, as follows (Supplementary Figure 6). We (a) compute a raw fingerprint matrix with $L = 2$, and (b) for each row in the matrix, corresponding to a SNV pair key, add to the binary fingerprint a 1 if the value of the second column is larger than the value in the first column. There is no need to normalize binary fingerprints. To compare two binary fingerprints, we count how many bits are identical between the two, divide by the number of bits (144) and square the resulting fraction.

Evaluation of the Robustness of Genome Fingerprints

Different Reference Versions

To evaluate the robustness of the fingerprinting method to different versions of the reference genome, we computed fingerprints ($L = 20$ and binary) for a set of 69 genomes sequenced by Complete Genomics, Inc. [a]; these genomes were mapped to both GRCh36 (hg18) and GRCh37 (hg19). We then computed all pairwise correlations, which include (a) each genome on both references, (b) combinations of different genomes on the same reference, and (c) combinations of different genomes with different references.

Format and Normalization

To evaluate the robustness of the fingerprinting method to transformation from vendor-specific formats to standard representations, we studied 2436 CGI genomes delivered in “var” format and 1618 CGI genomes delivered in “masterVar” format. The var format reports each allelic observation in a separate line, while the masterVar format represents each locus in one line (in similarity with the VCF format). We used custom parsers to transform these vendor-specific genome representations to VCF format, and normalized them using vt normalize (Tan et al., 2015). We then compared genome fingerprints computed from the original (var or masterVar) representations and from the normalized VCF representations. We performed these computations with $L = 20$ and with binary fingerprints.

Significant Filtering and Post-processing

To evaluate the robustness of the fingerprinting method to complex post-processing procedures, we studied 154 genomes from the 1000 Genomes Project for which individual WGS data are available from the International Genome Sample Resource [b]. The genomes were analyzed using Complete Genomics’ analysis pipeline versions 2.2.0.19 through 2.2.0.26 and reported in masterVar format. We compared for each genome fingerprints computed from the masterVar representation and from the version extracted from the multi-sample VCFs in release 20130502 (filenames: ALL.chrNN.phase3_shapeit2_mvncall_integrated_v5.20130502.genotypes.vcf.gz). We performed these computations with $L = 20$, $L = 200$, and binary fingerprints.

Sequencing Technology Effects

To evaluate the robustness of the fingerprinting method to differences between technologies, we computed fingerprints for several versions of the “platinum” NA12878 genome, some of them sourced from the Genome In A Bottle consortium (downloaded from [c]). These assemblies were sequenced using multiple versions of CGI’s technology and Illumina, processed using various pipeline versions, mapped to GRCh36, GRCh37, and GRCh38 (Figure 2). We performed these comparisons with $L = 5$, $L = 20$, $L = 120$, and binary fingerprints.

Version	CGI old hg18	CGI old hg19	CGI pipeline 2	CGI modern 1	CGI modern 2	CGI LFR	ILM Platinum	1000g	Technology	Reference	Identifier
CGI old hg18	1.000	0.990	0.902	0.887	0.869	0.857	0.827	0.834	CGI 1.10.0.26	GRCh36	GS00362-DNA_C04
CGI old hg19	0.959	1.000	0.907	0.890	0.872	0.860	0.830	0.838	CGI 1.10.0.26	GRCh37	GS00362-DNA_C04
CGI pipeline 2	0.892	0.905	1.000	0.900	0.871	0.854	0.829	0.836	CGI 2.0.0.26	GRCh37	GS00362-DNA_C04
CGI modern 1	0.778	0.790	0.778	1.000	0.887	0.879	0.844	0.847	CGI 2.5.0.19	GRCh37	GS02256-DNA_A01
CGI modern 2	0.815	0.778	0.766	0.706	1.000	0.843	0.806	0.812	CGI 2.5.0.20	GRCh37	GS02295-DNA_D01
CGI LFR	0.766	0.754	0.790	0.683	0.718	1.000	0.808	0.813	CGI LFR	GRCh37	GS000039392-ASM
ILM Platinum	0.828	0.815	0.853	0.766	0.706	0.778	1.000	0.879	illumina	GRCh38	NA12878
1000g	0.754	0.742	0.778	0.742	0.660	0.683	0.790	1.000	illumina **	GRCh37	NA12878

FIGURE 2 | Comparison of several versions of the same genome. We compared six versions of the genome of the same individual (NA12878), one of them (GS00363-DNA_C04) processed in three different ways. The values above the diagonal of the matrix represent pairwise Spearman correlations between fingerprints with $L = 20$, while numbers under the diagonal represent comparisons of binary fingerprints. Color scale from 0.5 (red) through 0.75 (white) to 1.0 (blue). CGI, Complete Genomics, Inc.; ILM, illumina; LFR, long fragment read; 1000g, 1000 Genomes Project. The asterisks “**” highlights the post-processing of the illumina-based genome sequence as done by the 1000 Genomes Project.

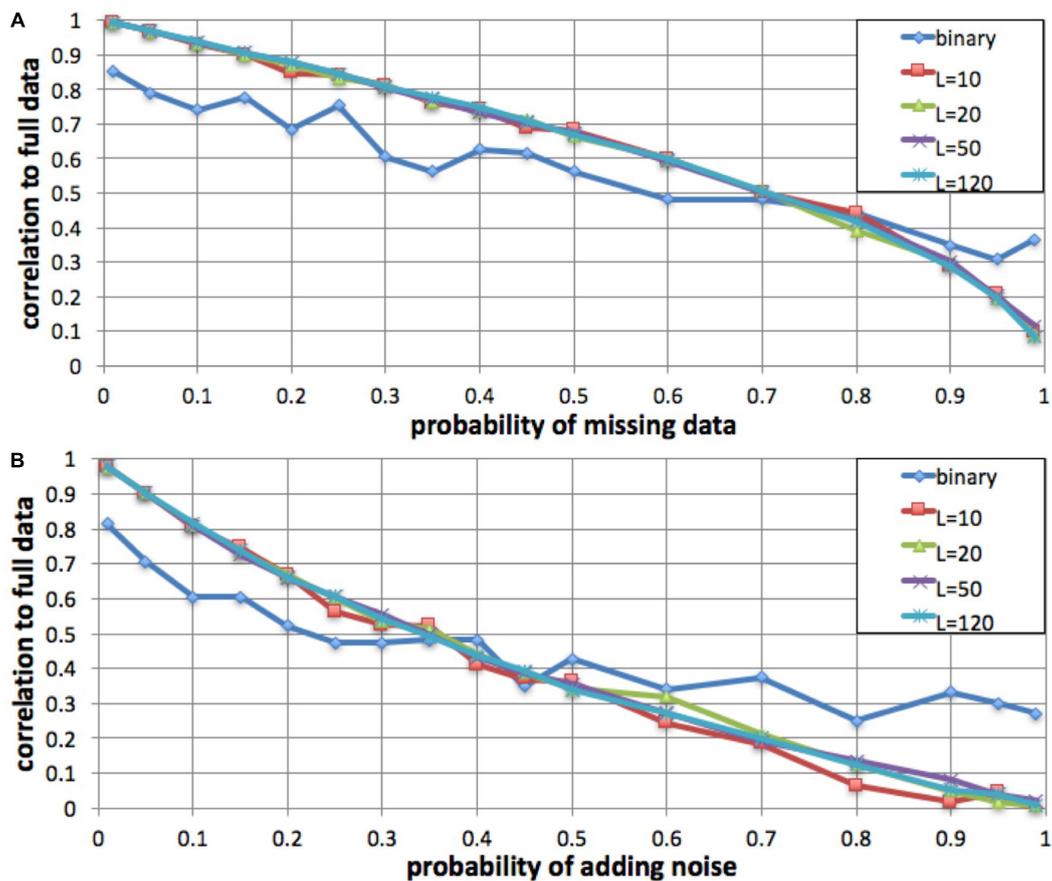


FIGURE 3 | Robustness to missing data and to added noise. We simulated increasing fractions of missing data (A) and to increasing probability of added noise (B) and computed the correlation between the resulting fingerprint and the fingerprint derived from the full data. For fingerprints with $L = 10$, $L = 20$, $L = 50$, and $L = 120$, the method was robust to up to 35% data loss and up to 15% added noise. The binary fingerprint was robust to up to 10% data loss and less than 5% added noise.

Missing Data and Added Noise

To evaluate the robustness of the fingerprinting method to missing data, we performed a simulation in which we degraded a genome to increasing degrees and compared the resulting

fingerprints to the original. We computed a series of fingerprints for the same genome by varying the probability of not observing each individual variant, excluding 1, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50, 60, 70, 80, 90, 95, and 99% of the variants at random

(**Figure 3**). We performed this simulation with $L = 10$, $L = 20$, $L = 50$, $L = 120$, and binary fingerprints.

Similarly, we evaluated robustness to added noise. We simulated noise from sequencing errors (false positives) by adding an additional, intervening spurious variant between each observed variant pair, with probabilities as in the missing data simulation.

Pedigree Study

We studied WGS data from 35 individuals in a large family (**Figure 4**). All genomes were sequenced by Complete Genomics, Inc. from blood ($n = 25$) or saliva ($n = 10$) samples, and processed using pipeline version 2.5.0.20. The genome data and a description of the family pedigree were donated by the private family. We categorized all pairwise relationships within the family as sibling, parent/child, half-sibling, aunt/uncle, grandparent, cousin, second cousin, second cousin once removed, unrelated, and other. The last category included more complex relations, e.g., child and grandchild of half-siblings. We computed for each individual a series of genome fingerprints using L in the range 2–200. For each L we computed all pairwise correlations, then computed their average and standard deviation stratified by the relationship categories listed above.

Population Reconstruction

We computed fingerprints ($L = 20$, $L = 120$, and binary) for each of the 2504 genomes in the 1000 Genomes Project data set from the version extracted from the multi-sample VCFs in release 20130502 (filenames: ALL.chrNN.phase3_shapeit2_mvncall_integrated_v5.20130502.genotypes.vcf.gz). For comparison, we analyzed the same data by principal components analysis (PCA) as follows. We identified SNPs with a minor allele frequency of 5% or more, removed SNPs in complete linkage disequilibrium with a SNP to the left (i.e., a smaller chromosomal position), retained 5% at random (298,454 SNPs) and counted occurrences of the minor allele (0, 1, or 2) in each genome to form a $2504 \times 298,454$ genotype matrix M . We performed PCA using the R function call `prcomp(M, center = TRUE, scale = TRUE)` and t-distributed Stochastic Neighbor Embedding (t-SNE) using the R package `Rtsne` with `perplexity = 25` and `max_iter = 500` (**Figure 5**).

We obtained data of the Sheep Genome Project Run 1 (International Sheep Genomics Consortium) from [d] as multi-sample VCFs (similar to the 1000 Genomes Project) with 453 samples. We computed fingerprints ($L = 120$) and analyzed using PCA as above (**Figure 6A**).

We obtained data of the 3000 Rice Genomes Project (International Rice Research Institute) from [e] as single-sample genome VCFs for 3000 rice varieties. We computed fingerprints ($L = 120$) and analyzed using PCA as above (**Figure 6B**).

Evaluation of Fingerprints for Population Classification

We computed a population fingerprint for each of the annotated populations studied by the 1000 Genomes Project. To identify the population closest to each individual, we compared each

individual's normalized fingerprint to the population fingerprints (using Spearman correlation, as described for comparing individual fingerprints). Each individual was considered classified as belonging to the closest population. To avoid bias, we excluded each individual from the computation of their own population fingerprint.

Population-Adjusted Fingerprints

We computed a population-adjusted fingerprint for each individual in the 1000 Genomes data set by subtracting the population fingerprint (excluding the individual) from the individual's normalized fingerprint. We computed all pairwise Spearman correlations and identified outliers using the absolute deviation around the median (Wilcox, 2016) with false discovery rate of 5%.

To compute kinship coefficients, we identified SNPs with a minor allele frequency of 5% or more, removed SNPs in complete linkage disequilibrium with a SNP to the left (i.e., a smaller chromosomal position), retained 5% at random (324,313 SNPs), reformatted using `vcftools` (Danecek et al., 2011) and `plink` (Purcell et al., 2007), and computed kinship coefficients using `KING` (Manichaikul et al., 2010). We obtained previously reported close relationships in the 1000 Genomes data set (Gazal et al., 2015) as computed using `RELPAIR` (Epstein et al., 2000).

RESULTS

A Novel Genome Encoding Method

We developed a novel algorithm for computing “fingerprints” from genome data; the algorithm is akin to locality-sensitive hashing (Indyk and Motwani, 1998). These fingerprints are rapidly computed, only need to be computed once per genome (not once per comparison), and can be rapidly compared to determine whether two genome sequences are derived from the same individual, closely related individuals, or unrelated individuals. Moreover, the original genome representation cannot be reconstructed from the fingerprint, and fingerprints can be shared when privacy concerns prevent sharing the genome itself.

Importantly, the fingerprints can be computed starting from any of several file formats, with a variety of encodings and relative to any reference version, and the resulting fingerprints are directly comparable without further conversion.

Based on characteristics of consecutive pairs of SNVs (see section “Methods”), our algorithm computes two small matrices of numbers for each genome (**Figure 1**), one of which (“close”) holds potentially useful information about the technology used to generate the genome, and the other (“raw”) represents the genome itself. The latter is then normalized and potentially adjusted to the relevant population fingerprint. Each of these matrix versions has its own utility.

The main parameter of our algorithm, L , determines the size of the fingerprint. Smaller fingerprints (e.g., $L = 20$) are useful for fast genome comparisons to determine identity, while larger fingerprints (e.g., $L = 120$, $L = 200$) retain more

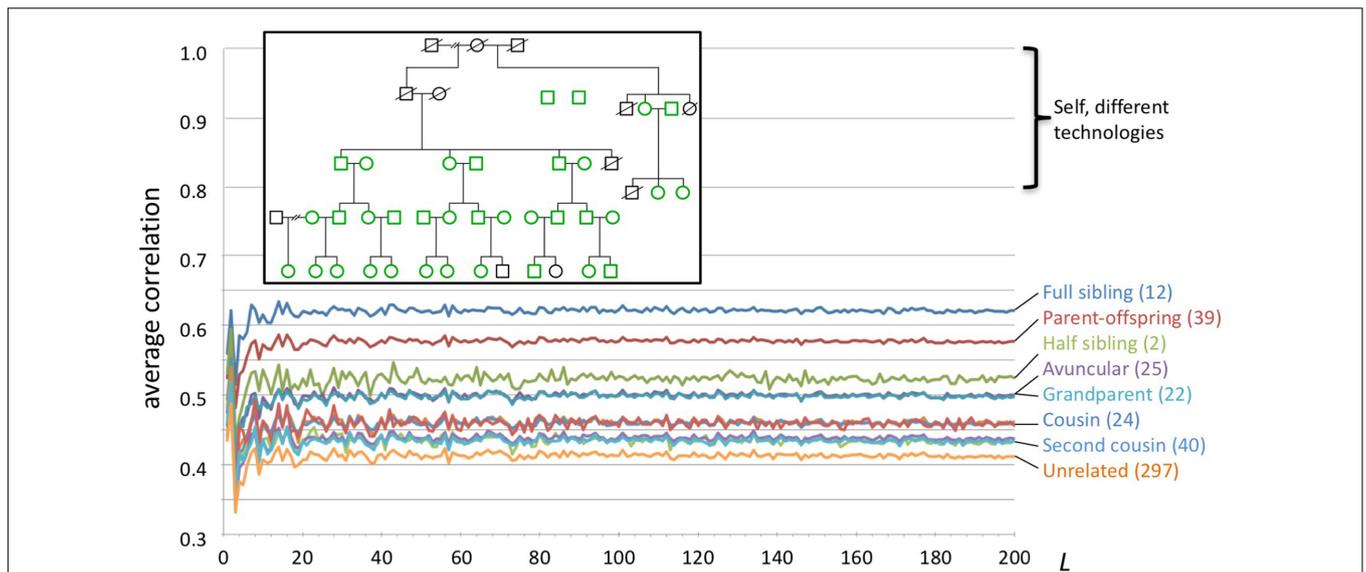


FIGURE 4 | Fingerprint correlation reflects degree of relatedness. Each trace represents the average correlation between fingerprints of individuals in each relatedness group, as a function of L . The number of pairs in each class is shown between parentheses. Inset: family structure; green icons represent individuals whose genomes were sequenced. Not shown: 169 additional pairs with complex relationships.

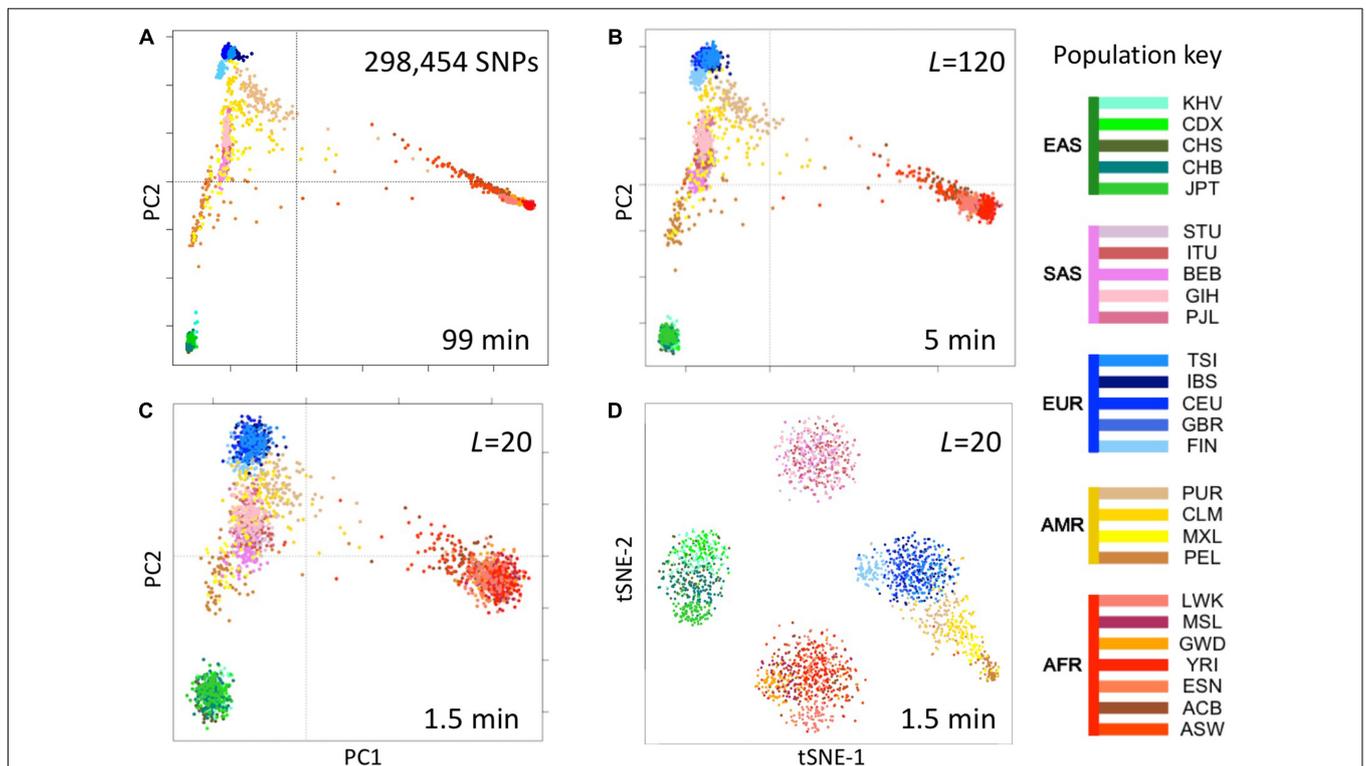
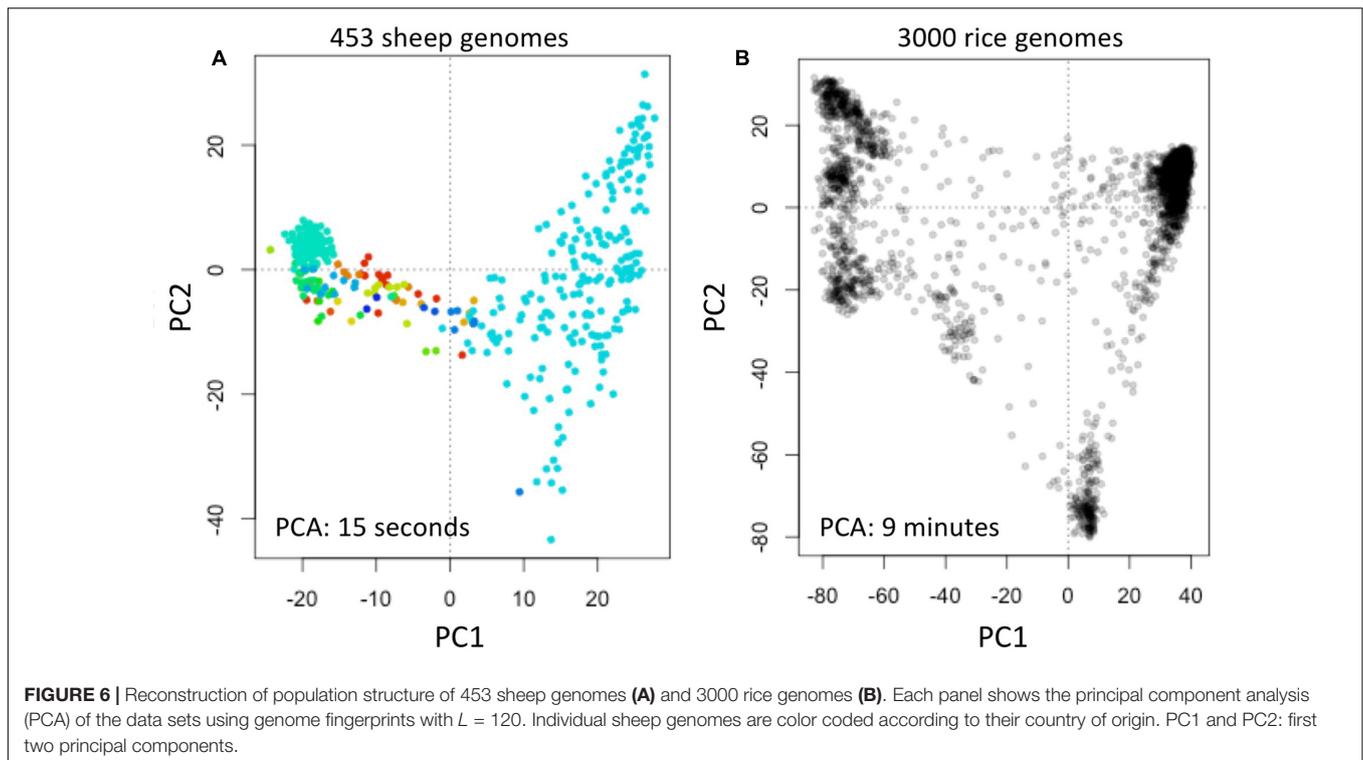


FIGURE 5 | Reconstruction of population structure of the 1000 Genomes Project data set. **(A)** Principal component analysis (PCA) of the 2504 individuals using ~300,000 SNPs. **(B)** PCA on genome fingerprints with $L = 120$, **(C)** PCA on genome fingerprints with $L = 20$, **(D)** t-SNE on genome fingerprints with $L = 20$. Individuals are color coded according to their population as per the key to the right. EAS, SAS, EUR, AMR, and AFR: East Asian, South Asian, European, Admixed American, and African, respectively. PC1 and PC2: first two principal components. tSNE-1 and tSNE-2: first two t-SNE components.



information and better support detailed analyses like population reconstruction.

Comparison of Genome Fingerprints

Two fingerprint matrices can be easily compared by first serializing each fingerprint into a vector of numbers (see section “Methods”) and then computing a correlation value between the two vectors. The correlation reflects the overall similarity between the two genomes being compared: identical input genomes yield perfect correlation. Based on our experiments, detailed below, comparison of two versions of the same genome typically display correlations above 0.75, and lower correlation values reflect various degrees of similarity that inform about family relatedness and population of origin. Various cutoff values can be used to determine the degree of relatedness represented by each correlation level; these cutoffs may depend on the parameters used for constructing fingerprints, and can be determined empirically through family studies as described below.

A Minimal Fingerprint

We further sought to create a minimal genome fingerprint. Using the smallest meaningful table ($L = 2$), we created a fingerprint of just 144 bits for each genome. Each bit represents, for each of the 144 possible SNV pair keys (e.g., “GAGT”), whether the distance (the number of intervening base pairs) between the consecutive SNVs is more frequently an odd or an even number. This simple function yields a “binary barcode” of a genome (Supplementary Figure 6) with negligible storage and memory footprints, trivially comparable to similarly computed “barcodes” by counting binary matches.

Computation on Genome Fingerprints Is Fast

Computation of genome fingerprints is very fast, typically requiring 15–45 s per genome, depending on the file format. The computation requires a single-pass read of the genome and depends principally on the time it takes to read the file (i.e., it is “I/O bound”). It is also trivially parallelizable: computation of one genome fingerprint does not depend on the results of similar computations for other genomes. It is also possible to compute partial raw fingerprints from sections of the genome, e.g., individual chromosomes; such partial fingerprints can then be combined by simple summation, yielding the raw fingerprint for the whole genome. This is particularly useful when the data are represented in separate multi-sample files per chromosome, as is the case for the 1000 Genomes data set.

Thanks to the small size of the fingerprint matrices, fingerprint comparisons are extremely fast. We performed all-against-all comparisons in the set of 2,504 genomes from the 1000 Genomes Project. The 3.1 million comparisons required 67 s (21.3 μ s per comparison) with $L = 120$, and 11 s (3.53 μ s per comparison) with $L = 20$. Fingerprint comparisons are also independent and trivially parallelizable.

Genome Fingerprints Are Robust Fingerprints Are Robust to Reference Versions

Mapping raw sequence data to different reference versions can yield disparate variant calls. The absolute coordinates for most observed variants will differ between references; a fraction of the variants will differ also in other aspects including reference

allele, variant allele and even chromosome assignment and strand. We studied a set of 69 genomes that were mapped to two reference versions and evaluated the effect of reference version on fingerprint comparison. With $L = 20$, the Spearman correlation between fingerprints derived from genomes mapped to GRCh36 and GRCh37 ranged from 0.989 to 0.991, while fingerprints derived from different genomes in this set ranged from 0.222 to 0.685 (see example in **Figure 2**). These differences are much larger than the standard deviations observed for each relatedness level (Supplementary Figure 8 and sections below). Binary fingerprints yielded lower correlations that nevertheless clearly separated self-comparisons (0.853–0.986) from comparisons of different individuals (0.243–0.683). Genomes mapped to different references therefore yield highly comparable fingerprints.

Fingerprints Are Robust to Format Differences

Genome sequencing vendors may deliver results in non-standard file formats that best convey technology-specific characteristics of the data; the variants themselves may be represented in inconsistent ways. Considering two vendor-specific file formats (Complete Genomics' var and masterVar), we evaluated the effect on fingerprinting of standardizing the genome data representation by transforming to the standard VCF format and by normalization of variants (Tan et al., 2015). With $L = 20$, the correlation between fingerprints derived from var files and from their normalized VCF counterparts was 0.9911 ± 0.0008 . For masterVar files, the correlation was 0.9982 ± 0.0015 . Binary fingerprints yielded marginally lower correlations: 0.9262 ± 0.0313 for var files and 0.9688 ± 0.0211 for masterVar files. We found that this transformation yields nearly identical fingerprints.

Fingerprints Are Robust to Complex Post-processing

The 1000 Genomes Project applied a number of post-processing steps to the collection of 2504 genomes to normalize variant calls and exclude spurious or unreliable variants. Thus, for each genome in the data set there are at least two versions: the initial, unprocessed individual genome, and the harmonized version in the context of the large study. We evaluated whether fingerprint comparison is robust to this complex post-processing procedure (see example in **Figure 2**). With $L = 20$, the correlation between initial and post-processed versions was 0.798 ± 0.010 . We observed nearly identical results with $L = 200$: 0.798 ± 0.008 . The lowest observed correlation for an individual genome was 0.776, well above the correlations observed for different individuals, even those closely related. As expected, binary fingerprints were more strongly affected by the post-processing procedure, yielding correlations of 0.660 ± 0.057 (Supplementary Figure 7). Here, the lowest observed correlation for an individual was 0.482, significantly below the highest correlation observed between binary fingerprints of two related individuals (0.683). Therefore, binary fingerprint comparisons for the purpose of determining identity between two genomes may frequently lead to false negatives in the presence of complex post-processing of genome data. We observed no false positive identifications using the 0.75

cutoff, though this could conceivably happen with much larger data sets.

Fingerprints Are Robust to Technology Differences

We compared fingerprints computed from various versions of the same genome (NA12878) as ascertained using different technologies, mapped with different pipelines and using different references (**Figure 2**). Using fingerprints with $L \geq 20$, all comparisons yielded correlations higher than 0.8. Shorter fingerprints ($L = 5$) yielded correlations above 0.75 (not shown). As expected, binary fingerprints were more sensitive to technology differences.

Fingerprints Are Robust in the Presence of Sequencing Errors

We evaluated, by simulation, whether fingerprint comparison is robust to missing data and to added noise. For a variety of fingerprint sizes, we degraded a personal genome by excluding (prior to computing fingerprints) a fraction of the variants (false negatives), and alternatively by adding spurious variant observations (false positives). For most fingerprint sizes, we observed a monotonic and equivalent decrease in Spearman correlation between the degraded genome fingerprints and the original, undegraded genome fingerprint. Given the 0.75 lowest correlation observed for different versions of the same genome, we conclude that the fingerprinting method is robust to up to 35% missing data and up to 15% added noise (**Figure 3**). As expected, the less informative binary fingerprint is less robust (up to 10% missing data, and to less than 5% added noise). All these error levels are significantly higher than what is typically observed in genome sequencing; this is consistent with the high correlations observed between technical replicates (**Figure 2**).

Fingerprint Similarity Reflects Degree of Relationship

We compared pairwise genome fingerprint correlations with known family relationships in a large pedigree (**Figure 4**). The family relationships spanned the range from close (e.g., siblings) to distant (e.g., second cousins once removed) and also included unrelated individuals (joining the family by marriage). We observed that fingerprint correlations decrease with increasing degree of relationship, from 0.62 ± 0.031 for full siblings to 0.41 ± 0.007 for unrelated individuals. This demonstrates that our genome fingerprinting method is locality-sensitive. We also observed that $L \geq 20$ is sufficient for distinguishing most relationship levels on average. Larger fingerprints have lower standard deviations (Supplementary Figure 8), improving resolution and the confidence of close relationship. The genomes of all individuals in this pedigree were sequenced using the same technology and processed using similar pipeline versions and the same reference version; this setup is typical for research studies. Combining data obtained using multiple technologies and processing pipelines could reduce the ability to confidently distinguish between relationship levels.

Of note, fingerprints derived from siblings tend to be more highly correlated than fingerprints of parents and their offspring (0.62 ± 0.031 vs. 0.58 ± 0.007 , respectively). In both cases the

degree of relationship is the same, but the identity by descent (IBD) pattern differs: parents and offspring share 100% of the genome in IBD1 state, while siblings share (on average) 50% of the genome in IBD1 state and 25% in IBD2 state. This increases the probability that sibling pairs will share heterozygous variants, especially heterozygous rare variants, as compared with parent/child pairs. For similar reasons, fingerprints of half siblings are more correlated than fingerprints of relatives of the same average degree (grandparental and avuncular relationships).

Fast Reconstruction of Population Structure

We tested the utility of genome fingerprints for population studies. We computed fingerprints for the 2504 individuals from 1000 Genomes Project and used PCA and t-SNE to reconstruct the known population structure (Figure 5) in a fraction of the time required to perform the same task using standard methods and with much smaller memory requirements. The quality of the reconstruction depended on fingerprint size: fingerprints with $L = 120$ yielded excellent population structure reconstruction and fingerprints with $L = 20$ yielded good separation at the continental level. PCA applied to fingerprints of different values of L provided results highly correlated with results from PCA applied to variants, with convergence to the same principal component axes as either the number of variants or the L increased. For a sufficient amount of data in either form, correlation between corresponding principal components was >0.99 for the first 5 to the first 10 components. This strongly suggests PCA of either kind of data provides equivalent information regarding population structure.

Importantly, the population reconstruction workflow using fingerprints is simple and requires no prior knowledge. The standard workflow requires creating multi-sample VCFs (which in turn requires normalizing variants), filtering out unreliable variants (which requires prior knowledge), filtering by allele frequency (requiring previous knowledge of population frequencies), paying attention to linkage disequilibrium, making assumptions about variants not observed in some samples and, finally, selecting a suitably sized subset of variants to balance resolution with memory and computational requirements. In some cases, these technical requirements can be quite severe, e.g., limiting the set of variants to a few thousands for t-SNE analysis (Li et al., 2017). In contrast, using genome fingerprints it is possible to reconstruct population structure by computing fingerprints directly from the unprocessed individual genomes and combining the fingerprints into a matrix ready for analysis with PCA, t-SNE or any other method of choice.

To test the general applicability of the method, we similarly studied 453 sheep (*Ovis aries*) genomes [d] and 3000 rice (*Oryza sativa*) genomes (The 3000 Rice Genome Project, 2014). These two data sets are available in different formats: the sheep genome data as multi-sample VCFs, while the rice genomes are provided as individual sample VCFs. We computed genome fingerprints with $L = 120$ and reconstructed population structure for each of these model species (Figure 6). While it took many hours to download the sheep and rice genome data

(14 and 30 h, respectively), fingerprint computation took just 1.5 h for the multi-sample sheep genome VCF and 55 min for the single-sample rice VCFs. Once the genome fingerprints were computed, the complete population reconstruction and visualization workflow took minutes to run (~ 5 min for sheep, ~ 15 min for rice), including the PCA component (15 s for sheep, 9 min for rice). These efforts did not require any prior knowledge of ovine (or rice) genome structure, linkage maps, population frequencies of variants, or metadata about the samples.

Utility for Fast and Simple Population Assignment

Individuals from the same population share some evolutionary history, and therefore some of the SNV pairs counted in computing genome fingerprints. It is thus useful to summarize the fingerprints of a population, both to estimate the “center” of the population’s fingerprints and their variability around that center (population diversity). Such “population fingerprints” have a variety of uses, including population assignment for individuals. We implemented a simple and very rapid method in which we first compute fingerprints ($L = 120$) for each population in the 1000 Genomes data set by simply averaging the fingerprints of the genomes in each population. We then compute the correlation between the fingerprint of a query genome and the fingerprint of each population: the genome is assigned to the population with which it is most strongly correlated. We tested this method by “leave one out” cross-validation. The correct population was identified as the best match for 2047 of 2504 query genomes (82% of cases). Some of the populations in the 1000 Genomes data set are very closely related and difficult to separate; for many practical purposes (e.g., selecting appropriate allele frequencies) using such closely related populations yields equivalent results. Accepting the second or third best population matches increased the success rate to 96 and 98%, respectively. At the continental resolution (five regions: AFR, AMR, EAS, EUR, and SAS), the best match was correct for all but 42 admixed AMR genomes.

Utility of Population-Adjusted Fingerprints

Fingerprints can be manipulated in a familiar way to perform different distance-based analyses. Population fingerprints, for example, are simple averages of a group that capture the features common to the group with respect to distance, but without corresponding interpretations that could be used for racial profiling or stereotyping. Population adjusted fingerprints allow analysis of relationships within a population, beyond the features shared among the population. This allows close relationships to be distinguished from membership in the same population and may help increase resolution of population structure reconstruction. We adjusted the $L = 120$ fingerprints for the 2504 individuals in the 1000 Genomes data set relative to their stated population of origin, performed all pairwise comparisons, compared with kinship coefficients computed using KING (Manichaikul et al., 2010) and with previously reported relationships (Gazal et al., 2015) computed using

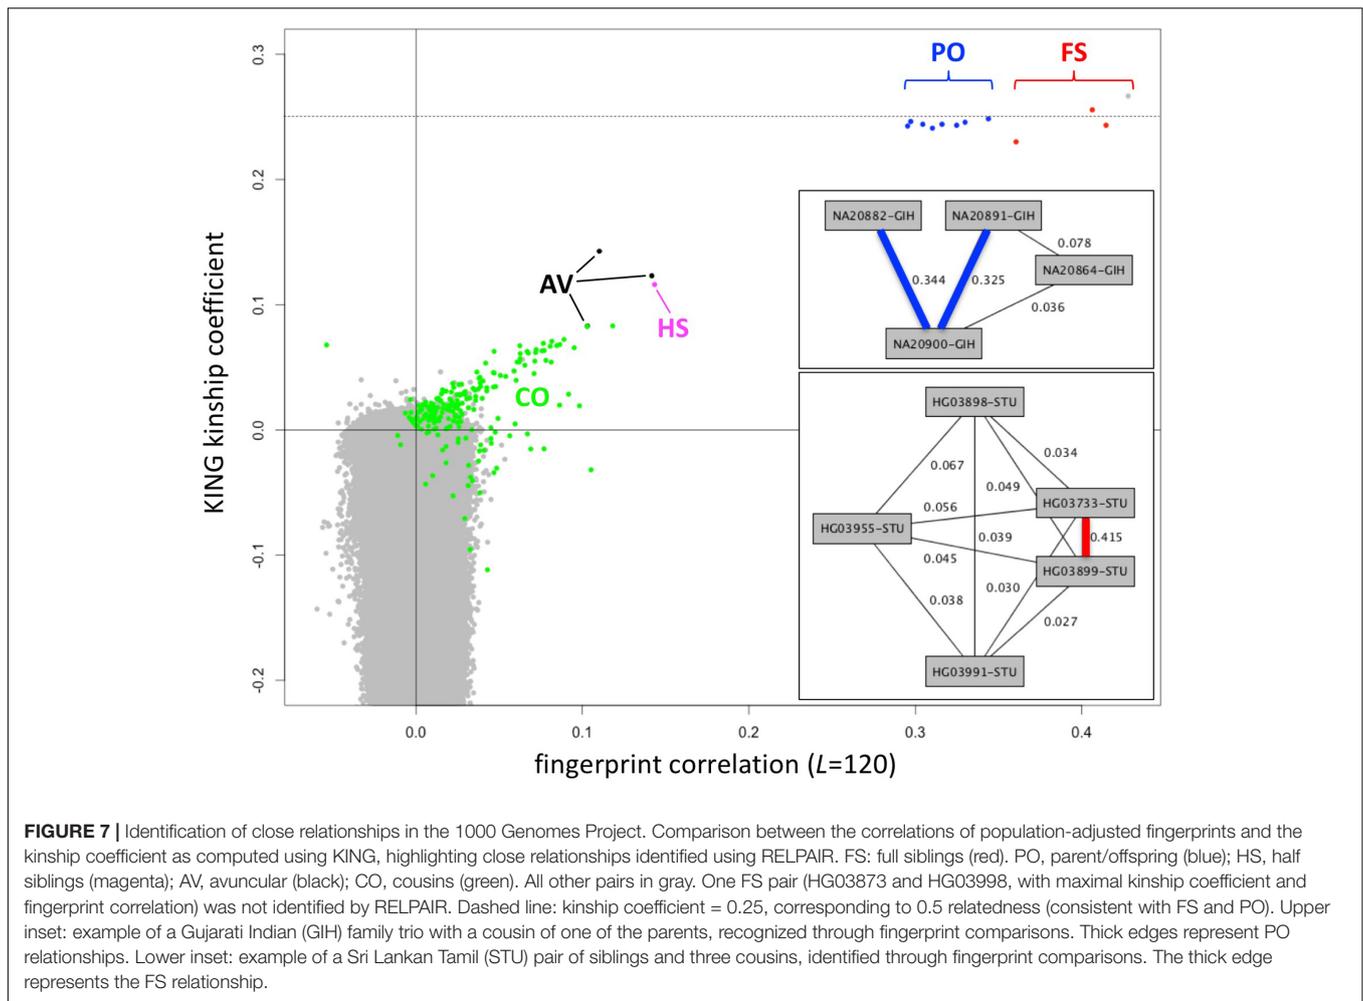


FIGURE 7 | Identification of close relationships in the 1000 Genomes Project. Comparison between the correlations of population-adjusted fingerprints and the kinship coefficient as computed using KING, highlighting close relationships identified using RELPAIR. FS: full siblings (red). PO, parent/offspring (blue); HS, half siblings (magenta); AV, avuncular (black); CO, cousins (green). All other pairs in gray. One FS pair (HG03873 and HG03998, with maximal kinship coefficient and fingerprint correlation) was not identified by RELPAIR. Dashed line: kinship coefficient = 0.25, corresponding to 0.5 relatedness (consistent with FS and PO). Upper inset: example of a Gujarati Indian (GIH) family trio with a cousin of one of the parents, recognized through fingerprint comparisons. Thick edges represent PO relationships. Lower inset: example of a Sri Lankan Tamil (STU) pair of siblings and three cousins, identified through fingerprint comparisons. The thick edge represents the FS relationship.

RELPAIR (Epstein et al., 2000) (Figure 7). Of note, these correlation values among population adjusted fingerprints are not comparable with those observed among normalized fingerprints in the family study (Figure 4). The similarity between unrelated individuals derived from the same population (e.g., Figure 4) is removed by adjustment to the population average. Thus, population-adjusted fingerprints for unrelated individuals show no significant correlation.

We observed 143 pairs of individuals from the same population with correlations higher than expected (at 5% false discovery rate) and two outlier pairs linking individuals of the ITU and STU populations (Supplementary Table 1), consistent with previous reports (Gazal et al., 2015). The highly correlated pairs correspond to a variety of degrees of relationship, from full siblings to cousins. In several cases they form networks of related individuals, e.g., a Gujarati Indian trio with a cousin of one of the parents (Figure 7, upper inset) and a Sri Lankan Tamil pair of full siblings and three cousins (Figure 7, lower inset). For related pairs, fingerprint correlation levels are correlated with the kinship coefficient as computed using KING (Figure 7). The fingerprint correlation method can further distinguish between full sibling and parent/offspring relationships. Many more cousin pairs are

clearly distinguished from unrelated pairs by their fingerprint correlations than by their kinship coefficient values. For unrelated pairs, fingerprint correlations are narrowly distributed around 0 (no correlation), while the kinship coefficients range from ~0.1 down to -0.4.

DISCUSSION

We have presented a novel algorithm for computing “fingerprints” of individual genomes, and various examples of their potential applications. Genome fingerprints are reduced representations that retain information about distances (between genomes) but not functional features, enabling ultrafast comparison of genomes. Importantly, the fingerprints need to be computed just once per genome, not once per comparison. They can also be computed directly from any representation of the genome, regardless of technology, reference version, formatting and encoding choices, and even significant levels of missing data and added noise. These features provide for the first time the ability to rapidly compute the level of similarity between genomes regardless of such representational differences.

We described an implementation of the method for whole-genome data and various aspects of comparing whole-genome fingerprints, pairwise and in large sets. Fingerprints can similarly be computed for exome data. The fingerprint computed from an individual's exome would not be directly comparable to the fingerprint derived from the same individual's whole genome; it is though possible and very efficient to approximate exome data from whole-genome data, by simple subsetting using bedtools (Quinlan and Hall, 2010).

Our current implementation of pairwise comparison of genomes via fingerprints takes microseconds per comparison; using multithreading or specialized hardware acceleration could potentially significantly improve performance. These speeds are already orders of magnitude faster than afforded by current methods, and enable orders of magnitude more comparisons to be routinely performed. As expected from the highly lossy data reduction involved in generating genome fingerprints, their comparison for various applications yields results that are somewhat inferior to direct computation on the entire dataset of all observed variants. The fast approximation afforded by genome fingerprints will be good enough for some applications, and will help prioritize more detailed computations where precise results are required. Widespread adoption of this methodology could revolutionize the field of comparative genomics by enabling comparisons on a scale not now attempted due to the time and effort required. For example, all-against-all pairwise comparisons of a million genomes could be performed with $L = 20$ in 20 CPU days (or under 1 day with minimal parallelization), followed by more refined comparison of the subset of genome pairs of interest.

We have shown that regardless of the technology, genome freeze, or encoding used for specific personal genomes, the corresponding fingerprints enable rapid testing of whether two genome representations are derived from the same person or from closely related individuals, rapid identification of the closest among a specified set of candidate populations, and acceleration of population structure studies. Genome fingerprints are also lightweight and restricted to distance information, and therefore suitable for databasing on a broader scale than the personal genomes themselves. Such databases would enable rapid testing of whether a new personal genome has already been observed, rapidly ensuring that a data set includes only unrelated individuals, and rapid testing for the presence of shared genomes in two or more studies, all of which are common steps in the planning and construction of cohorts of genomes for more sensitive, functional analysis. Distances between fingerprints also provide an unbiased scale for selecting pairs of individuals at comparable distances, for example, to select individuals from a control set matched to a set of cases at a consistent evolutionary distance.

Genome fingerprints are highly robust to a variety of important technical issues, including different reference versions, variant normalization and post-processing, different sequencing technologies, and different variant sampling efficiencies (missing data). We observed that significant post-processing of genome data in a large cohort, as done by the 1000 Genomes Project, can lead to reduced similarities with the unprocessed versions of the

same genomes. Several processing steps may contribute to this effect: exclusion of problematic “blacklisted” regions, imputation of missing values, changes in genotypes determined through “joint calling,” and the increased chance of observing multi-allelic variants—which are excluded from our analysis.

We evaluated the value of the most compact form of fingerprint, binary fingerprints. Despite their simplicity and compactness, binary fingerprints are typically sufficient to establish whether two genome representations are derived from the same individual, and even to distinguish between closely related and unrelated individuals; however, they do not retain enough information to make these determinations when comparing mixtures of unprocessed and highly post-processed versions of genomes.

Our fingerprinting algorithm is a fully deterministic form of locality-sensitive hashing (Indyk and Motwani, 1998), a class of algorithms which has been successfully applied to large-scale sequence comparisons (Buhler, 2001), protein classification (Caragea et al., 2012), metagenome clustering (Rasheed et al., 2013), and the detection of gene–gene interactions (Brinza et al., 2010). Locality sensitive hashing methods based on indexing substrings have been used to compare genomes, typically of different species, and other sequence data types. By encoding single-nucleotide differences from a reference sequence, our method focuses on representing distances among individuals of the same species, for whom the base (reference) sequence is identical and need not be indexed as k -mers. Accordingly, the main application of our method is to detect relationships at various levels—from identity to population structure.

We chose to focus on the most common and well-defined type of genome variation—SNVs. Other types of variation, including short insertions and deletions (indels), larger structural variants, and copy-number variants) are significantly less common than SNVs, and are frequently either represented differently in different genomes, not detected, or not reported: we therefore chose to exclude them for the sake of simplicity and consistency in analysis. Our results demonstrate that SNV-based fingerprints contain enough information to support many genome comparison tasks; the other types of variation are not necessary to estimate evolutionary distances within a species, and could be detrimental for the above reasons.

For each SNV, we only consider the nature of the reference allele and the alternate allele as observed in the individual. Our current implementation considers only positions in which an alternate allele is observed, regardless of zygosity: heterozygous SNVs and homozygous-alternate SNVs are treated equally. We also evaluated alternative versions of the method in which SNVs with different zygosity are weighted differently. Some of these alternative versions will be useful in particular contexts, e.g., for fingerprinting genomes in a fully reference-free fashion by considering only heterozygous SNVs. For each pair of consecutive SNVs, we consider the distance separating them along the reference genome. This is unaffected by any intervening indels or other complex variants in the individual.

We also chose to include only SNVs observed in autosomes. Inclusion of variants in the sex chromosomes leads to distorted similarity values. For example, due to shared variants on the X

chromosome, a female child would appear to be more closely related to her mother than to her father, or to be as closely related to her grandfather as to her grandfather's sister.

Pairs of consecutive SNPs within 200 bp (microhaplotype loci) have been recognized as a powerful tool for individual identification, ancestry inference, determining family/clan relationships, and other forensic applications (Kidd et al., 2013). Our method collects from each genome all SNV pairs regardless of distance and population frequency, thus extracting a phylogenetic signal in an unbiased way and without necessitating prior population knowledge. Most importantly, our algorithm successfully distributes this phylogenetic signal onto a matrix representation that preserves distances without preserving individual, functionally interpretable variants. While some algorithmic aspects of our method may seem somewhat arbitrary, it is simple and efficient to compute. Prior applications of locality sensitive hashing to genomics have demonstrated the utility of quite arbitrary-seeming algorithms. For example, the kernel of a locality sensitive hashing approach to genome-scale assembly of single-molecule sequencing data uses the XORShift pseudo-random number generator to transform k -mer hashes into comparable "sketches" (Berlin et al., 2015).

The focus on local patterns of SNVs offers an additional advantage of our method. To date, the reference genome has been represented as a collection of linear sequences (chromosomes) with absolute coordinates. One disadvantage of this method is that the absolute coordinates change from version to version; another is that the linear representation is not flexible enough to represent the rich diversity in structural variation observed within our species. The genomics community is now developing a new, graph-based reference genome format, which will further devalue the use of global sequence coordinates as the method for matching sequence variants. We expect that genome fingerprints can be computed from graph-based genome representations, which will be highly comparable to fingerprints we currently compute from linear-reference representations.

Public sharing of genome data has been limited by multiple personal privacy and confidentiality considerations. A central risk is the possibility of identifying genetic predispositions to certain diseases or other traits that could affect the individual's ability to obtain or maintain employment, insurance or financial services, or may carry social stigma, or could lead to other negative effects. Quantifying this risk is difficult since the ability to interpret genomic variants will expand over time with additional research. Our method enables sharing enough information about a genome to enable the comparison tasks, without concomitantly revealing the information needed for predicting phenotype. As such, our method *decouples genome comparison from interpretation*. This property has important implications for privacy-preserving genome analytics. Identifying genomes harboring a specific variant can be of interest both in the presence of associated phenotype, as facilitated through the Matchmaker Exchange (Philippakis et al., 2015), or in its absence, to ascertain novelty and frequency (Glusman et al., 2011 and the Beacon Network [f]). Genome fingerprints support the complementary task of matchmaking via identification of

closely related individuals, without exposing variant information, in similarity with the UNICORN method (Bodea et al., 2016) but with a much simpler algorithm that does not require extensive prior modeling of variant frequencies, nor samples to be expressed relative to the same reference. Genome fingerprints can furthermore be used to compare genomes from populations not previously studied.

Finally, as growing contingents of private individuals get access to their own genetic data, there is increasing public interest in efficient and private analysis of personal genomic variation, not just for interpreting variants and their combinations, but also for identifying related individuals. As such, our method has strong potential for empowering citizen science.

INTERNET RESOURCES

- (a) <http://www.completegenomics.com/public-data/69-genomes/>
- (b) <http://www.internationalgenome.org/announcements/complete-genomics-data-release-2013-07-26/>
- (c) <ftp://ftp-trace.ncbi.nlm.nih.gov/giab/ftp/data/NA12878/analysis/>
- (d) <http://www.ebi.ac.uk/ena/data/view/PRJEB14685>
- (e) <https://aws.amazon.com/public-datasets/3000-rice-genome/>
- (f) <http://ga4gh.org/#/beacon>
- (g) Documentation, code, sample datasets and more are available at:
 - http://db.systemsbiology.net/gestalt/genome_fingerprints
 - <https://github.com/gglusman/genome-fingerprints>

AUTHOR CONTRIBUTIONS

GG, LH, and MR designed the study. GG, DM, and MR performed analyses. All authors contributed to writing the manuscript and approved its final version.

FUNDING

This work was supported by NIH grant U54 EB020406.

ACKNOWLEDGMENTS

We wish to thank Chris Witwer, Ben Heavner, Irit Rubin, Jeremy Protas, and Terry Farrah for helpful discussions.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://journal.frontiersin.org/article/10.3389/fgene.2017.00136/full#supplementary-material>

REFERENCES

- Berlin, K., Koren, S., Chin, C.-S., Drake, J. P., Landolin, J. M., and Phillippy, A. M. (2015). Assembling large genomes with single-molecule sequencing and locality-sensitive hashing. *Nat. Biotechnol.* 33, 623–630. doi: 10.1038/nbt.3238
- Bodea, C. A., Neale, B. M., Ripke, S., Daly, M. J., Devlin, B., Roeder, K., et al. (2016). A method to exploit the structure of genetic ancestry space to enhance case-control studies. *Am. J. Hum. Genet.* 98, 857–868. doi: 10.1016/j.ajhg.2016.02.025
- Brinza, D., Schultz, M., Tesler, G., and Bafna, V. (2010). RAPID detection of gene-gene interactions in genome-wide association studies. *Bioinformatics* 26, 2856–2862. doi: 10.1093/bioinformatics/btq529
- Buhler, J. (2001). Efficient large-scale sequence comparison by locality-sensitive hashing. *Bioinformatics* 17, 419–428. doi: 10.1093/bioinformatics/17.5.419
- Caragea, C., Silvescu, A., and Mitra, P. (2012). Protein sequence classification using feature hashing. *Proteome Sci.* 10(Suppl. 1):S14. doi: 10.1186/1477-5956-10-S1-S14
- Danecek, P., Auton, A., Abecasis, G., Albers, C. A., Banks, E., DePristo, M. A., et al. (2011). The variant call format and VCFtools. *Bioinformatics* 27, 2156–2158. doi: 10.1093/bioinformatics/btr330
- Epstein, M. P., Duren, W. L., and Boehnke, M. (2000). Improved inference of relationship for pairs of individuals. *Am. J. Hum. Genet.* 67, 1219–1231. doi: 10.1016/S0002-9297(07)62952-8
- Gazal, S., Sahbatou, M., Babron, M.-C., Génin, E., and Leutenegger, A.-L. (2015). High level of inbreeding in final phase of 1000 Genomes Project. *Sci. Rep.* 5:17453. doi: 10.1038/srep17453
- Glusman, G., Caballero, J., Mauldin, D. E., Hood, L., and Roach, J. C. (2011). Kaviar: an accessible system for testing SNV novelty. *Bioinformatics* 27, 3216–3217. doi: 10.1093/bioinformatics/btr540
- Indyk, P., and Motwani, R. (1998). “Approximate nearest neighbors,” in *Proceedings of the Thirtieth Annual ACM Symposium on Theory of computing - STOC '98* (New York, NY: ACM Press), 604–613.
- Kidd, K. K., Pakstis, A. J., Speed, W. C., Lagace, R., Chang, J., Wootton, S., et al. (2013). Microhaplotype loci are a powerful new type of forensic marker. *Forensic Sci. Int. Genet. Suppl. Ser.* 4, e123–e124. doi: 10.1016/j.fsigss.2013.10.063
- Layer, R. M., Kindlon, N., Karczewski, K. J., Consortium, E. A., Quinlan, A. R., City, S. L., et al. (2016). Efficient genotype compression and analysis of large genetic variation datasets. *Nat. Methods* 13, 63–65. doi: 10.1038/nmeth.3654
- Li, W., Cerise, J. E., Yang, Y., and Han, H. (2017). Application of t-SNE to human genetic data. *J. Bioinform. Comput. Biol.* 15, 1750017. doi: 10.1142/S0219720017500172
- Manichaikul, A., Mychaleckyj, J. C., Rich, S. S., Daly, K., Sale, M., and Chen, W. M. (2010). Robust relationship inference in genome-wide association studies. *Bioinformatics* 26, 2867–2873. doi: 10.1093/bioinformatics/btq559
- Pedersen, B. S., and Quinlan, A. R. (2017). Who's Who? Detecting and resolving sample anomalies in human DNA sequencing studies with peddy. *Am. J. Hum. Genet.* 100, 406–413. doi: 10.1016/j.ajhg.2017.01.017
- Pengelly, R. J., Gibson, J., Andreoletti, G., Collins, A., Mattocks, C. J., and Ennis, S. (2013). A SNP profiling panel for sample tracking in whole-exome sequencing studies. *Genome Med.* 5:89. doi: 10.1186/gm492
- Philippakis, A. A., Azzariti, D. R., Beltran, S., Brookes, A. J., Brownstein, C. A., Brudno, M., et al. (2015). The matchmaker exchange: a platform for rare disease gene discovery. *Hum. Mutat.* 36, 915–921. doi: 10.1002/humu.22858
- Purcell, S., Neale, B., Todd-Brown, K., Thomas, L., Ferreira, M. A. R., Bender, D., et al. (2007). PLINK: a tool set for whole-genome association and population-based linkage analyses. *Am. J. Hum. Genet.* 81, 559–575. doi: 10.1086/519795
- Quinlan, A. R., and Hall, I. M. (2010). BEDTools: a flexible suite of utilities for comparing genomic features. *Bioinformatics* 26, 841–842. doi: 10.1093/bioinformatics/btq033
- Rasheed, Z., Rangwala, H., and Barbará, D. (2013). 16S rRNA metagenome clustering and diversity estimation using locality sensitive hashing. *BMC Syst. Biol.* 7(Suppl. 4):S11. doi: 10.1186/1752-0509-7-S4-S11
- Tan, A., Abecasis, G. R., and Kang, H. M. (2015). Unified representation of genetic variants. *Bioinformatics* 31, 2202–2204. doi: 10.1093/bioinformatics/btv112
- Tatwawadi, K., Hernaez, M., Ochoa, I., and Weissman, T. (2016). GTRAC: fast retrieval from compressed collections of genomic variants. *Bioinformatics* 32, i479–i486. doi: 10.1093/bioinformatics/btw437
- The 3000 Rice Genome Project (2014). The 3,000 rice genomes project. *Gigascience* 3:7. doi: 10.1186/2047-217X-3-7
- Wang, P. P. S., Parker, W. T., Branford, S., and Schreiber, A. W. (2016). BAM-matcher: a tool for rapid NGS sample matching. *Bioinformatics* 32, 2699–2701. doi: 10.1093/bioinformatics/btw239
- Wei, L., Liu, L. T., Conroy, J. R., Hu, Q., Conroy, J. M., Morrison, C. D., et al. (2015). MAC: identifying and correcting annotation for multi-nucleotide variations. *BMC Genomics* 16:569. doi: 10.1186/s12864-015-1779-7
- Wilcox, R. (2016). *Introduction to Robust Estimation and Hypothesis Testing*, 4th Edn. Cambridge, MA: Academic Press.

Conflict of Interest Statement: GG and MR hold a provisional patent application on the method described in this manuscript. GG and LH hold stock options in Arivale, Inc. Arivale, Inc. did not fund the study and was not involved in its design, implementation, or reporting.

The other author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Glusman, Mauldin, Hood and Robinson. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.