

METHODOLOGY

Open Access



# scAI-SNP: a method for inferring ancestry from single-cell data

Sung Chul Hong<sup>1</sup>, Francesc Muyas<sup>2</sup>, Isidro Cortés-Ciriano<sup>2</sup> and Sahand Hormoz<sup>1,3,4\*</sup>

## Abstract

**Background** Collaborative efforts, such as the Human Cell Atlas, are rapidly accumulating large amounts of single-cell data. To ensure that single-cell atlases are representative of human genetic diversity, we need to determine the ancestry of the donors from whom single-cell data are generated. Self-reporting of race and ethnicity, although important, can be biased and is not always available for the datasets already collected.

**Methods** Here, we introduce scAI-SNP, a tool to infer ancestry directly from single-cell genomics data. To train scAI-SNP, we identified 4.5 million ancestry-informative single-nucleotide polymorphisms (SNPs) in the 1000 Genomes Project dataset across 3201 individuals from 26 population groups. For a query single-cell dataset, scAI-SNP uses these ancestry-informative SNPs to compute the contribution of each of the 26 population groups to the ancestry of the donor from whom the cells were obtained.

**Results** Using diverse single-cell datasets with matched whole-genome sequencing data, we show that scAI-SNP is robust to the sparsity of single-cell data, can accurately and consistently infer ancestry from samples derived from diverse types of tissues and cancer cells, and can be applied to different modalities of single-cell profiling assays, such as single-cell RNA-seq and single-cell ATAC-seq.

**Discussion** Finally, we argue that ensuring that single-cell atlases represent diverse ancestry, ideally alongside race and ethnicity, is ultimately important for improved and equitable health outcomes by accounting for human diversity.

**Keywords** SNP, Single-cell transcriptomics, Ancestry, scAI-SNP, The 1000 Genomes Project

## Introduction

In recent years, single-cell technologies, especially single-cell RNA sequencing (scRNA-seq) and single-cell transposase-accessible chromatin sequencing (scATAC-seq), have become ubiquitous, revolutionizing our understanding of cellular heterogeneity in many biological

systems. Consequently, many datasets have emerged, epitomized by ambitious endeavors such as the Human Cell Atlas [1, 2] and the Human Tumor Atlas Network [3], which aim to create comprehensive reference maps of all human cells in healthy and diseased tissues. These datasets hold the potential for groundbreaking discoveries in personalized medicine, developmental biology, and a myriad of other fields.

The utility and impact of a reference atlas, such as the Human Cell Atlas, hinge crucially on its representativeness of human diversity. It is imperative to construct the reference atlases from a wide array of donors, encompassing all ancestries. This diversity is not merely a demographic criterion but a fundamental biological necessity. Genetic variation across different populations can significantly influence gene expression patterns,

\*Correspondence:

Sahand Hormoz  
sahand\_hormoz@hms.harvard.edu

<sup>1</sup> Department of Data Science, Dana-Farber Cancer Institute, Boston, MA 02215, USA

<sup>2</sup> European Molecular Biology Laboratory, European Bioinformatics Institute, Hinxton, Cambridge CB10 1SD, UK

<sup>3</sup> Department of Systems Biology, Harvard Medical School, Boston, MA 02115, USA

<sup>4</sup> Broad Institute of MIT and Harvard, Cambridge, MA 02142, USA



© The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

epigenetic landscapes, and cellular responses to environmental stimuli [4–6]. Without a diverse representation, reference atlases risk harboring biases, potentially leading to erroneous biological conclusions and limiting their applicability in global health contexts. For example, hematological characterization of individuals of European descent has shown significant differences in certain blood counts compared with those of African descent [7]. Consequently, reference datasets of blood counts constructed from predominantly European individuals have misled physicians into prescribing inappropriately low dosages of chemotherapy for early-stage breast cancer patients of African descent, who had different baseline ranges of neutrophils, resulting in poorer survival outcomes [8–10]. Therefore, single-cell atlases cannot truly represent the general population if the underlying data disproportionately represent or exclude specific ancestry groups.

One way to ensure that an atlas is representative of human genetic diversity is to require the donors to self-report their race or ethnicity. Unlike ancestry, which is determined by the genetic variants common to population groups that are inherited by an individual and can be reliably quantified through genetic analyses, race and ethnicity are socially constructed and ascribed. Race is socially assigned based on physical features, such as skin color, and ethnicity is socially assigned based on cultural similarities between individuals, such as a common language. However, using self-reported race or ethnicity to build a single-cell atlas has several drawbacks. People might not know or accurately report their complete racial or ethnic background, especially if they have a mixed heritage. Also, we cannot add self-reported race or ethnicity information retroactively to older samples that did not collect it. This limitation is particularly relevant given the vast repositories of historical single-cell data that hold potential insights. Furthermore, in genome-wide association studies (GWAS), using diverse groups based on ancestry in addition to race and ethnicity improves our ability to discover genetic predictors of complex traits [11]. Therefore, we need a method to infer ancestry directly from the single-cell data itself.

Single nucleotide polymorphisms (SNPs) that have significantly different frequencies across different population groups can be used to infer an individual's ancestry. Some methods have focused on identifying a small subset of such SNPs that are most informative for determining population structures [12–14]. Although for some applications, such as microarray measurements, having a selective set of SNPs is cost-effective, it is more strategic with sparse data like scRNA-seq data to cast a wider net; that is, to leverage numerous sites in the genome. More recently, whole genome sequencing (WGS) has

enabled the development of genome-wide approaches that account for all the SNPs detected in an individual's genome when assigning ancestry. There are two broad approaches for inferring ancestry from whole genome sequences: global and local. In global inference, the SNPs are used to assign an individual to an ancestral group or a mixture of the ancestral groups [15–17]. In local inference, each region of the genome of an individual is assigned to an ancestral group or a mixture of such groups, and the goal is to find the appropriate boundaries of these regions [18]. Here, we focus on global inference of ancestry from single-cell data.

Two general approaches have been proposed for inferring global ancestry from sequencing data, parametric and non-parametric. Parametric approaches, such as STRUCTURE [16], FRAPPE [19], fastSTRUCTURE [20], and ADMIXTURE [15], find the optimal choice of parameters for a statistical model where the parameters correspond to the fraction that each ancestry group contributes to each individual in the data and the frequency of the alleles in each population group. The posterior distribution of the parameters can be found using Bayesian methods such as Markov-Chain Monte Carlo or approximated efficiently using variational methods that rely on optimization. Non-parametric approaches, such as EIGENSTRAT [17], take a geometric perspective and identify the linear combinations of alleles that capture the directions of largest variation across the individuals in the data. The first few directions often capture the population structure in the data and can be used to assign the ancestral groups (or combinations thereof) to each individual [21]. Importantly, both approaches can work without any prior knowledge of the ancestral groups or their allele frequencies. The parametric methods can learn this information while simultaneously fitting a model to the data using expectation maximization, whereas non-parametric methods find the directions of maximum variation present in the data without relying on labeled individuals or population structure.

Our approach builds on the non-parametric methods for global inference of ancestry because these methods are more robust to missing data and are fast. Specifically, similar to EIGENSTRAT, we use principal component analysis (PCA) to identify the most informative directions in allele space for inferring ancestry (see also [22]). Unlike EIGENSTRAT, our goal is not to only capture variation in a dataset due to population structure but also to use these directions to assign out-of-sample individuals to a known ancestry group. Therefore, we used a dataset of individuals with known (self-reported) ancestry from the 1000 Genomes Project (1KGP) [23]. We identified 4.5 million ancestry-informative SNPs in the 1KGP WGS dataset and converted each individual's genotype

to a 4.5 million-dimensional vector where each element of the vector corresponds to the allele of the given SNP site detected in that individual. Next, we conducted PCA on these vectors to find a low-dimensional representation for them (Fig. 1b). Finally, we obtain a single vector representation of each population group by averaging the vector representation in PCA space of all individuals in the 1KGP data that belong to that population group.

In its final form, our method scAI-SNP (pronounced sky-snip) is remarkably easy to use. The user can take single-cell data of interest and extract the genotype of the 4.5 million ancestry-informative SNP sites. We genotype the SNP sites of scRNA-seq samples using a simple extension of the tool SComatic [24]. Other methods such as Monopogen [25], an elegant work that uses linkage disequilibrium to accurately detect germline variants, can also be used. Importantly, the inference is remarkably robust to missing data and sequencing errors and has little diminished accuracy when as many as 99% of the SNP sites are not detected.

To carry out the inference of ancestry, we simply multiply the genotype vector of the 4.5 million SNP sites (allele vector) by a 4.5 million by 600 dimensional PCA projection matrix to obtain a 600-dimensional vector representation of the input data. We then use convex optimization to find the linear combination of the 26 mean vectors of each population group that best approximates the input vector with the constraint that the coefficients of the linear combination are non-negative and sum to 1. The computed coefficients are the contributions of the 26 population groups to the ancestry of the individual from whom the single-cell data was obtained. A similar approach for linear deconvolution of ancestry has been applied to frequencies of variants [26] and principal components derived from these frequencies [27] but at the level of populations as opposed to individuals. We will show that scAI-SNP works for different modalities of single-cell profiling methods, across different tissues and cell types, and even when applied to data from cancer cells. The code for scAI-SNP is available on the Hormoz Lab GitHub page [28].

## Methods

scAI-SNP involves two stages: 1) model training and 2) model prediction. Briefly, for model training, we used the 1KGP dataset to generate a set of ancestry-informative SNPs and train an algorithm which takes in a vector corresponding to the alleles at these SNP locations and outputs a probability distribution over the 26 ancestral groups. For model prediction, we first used an extension of SComatic [24] to genotype the ancestry-informative SNP sites in single-cell data, thereby obtaining the input vector to the model and then applied scAI-SNP to infer

ancestry. Next, we describe each step of training and prediction in detail.

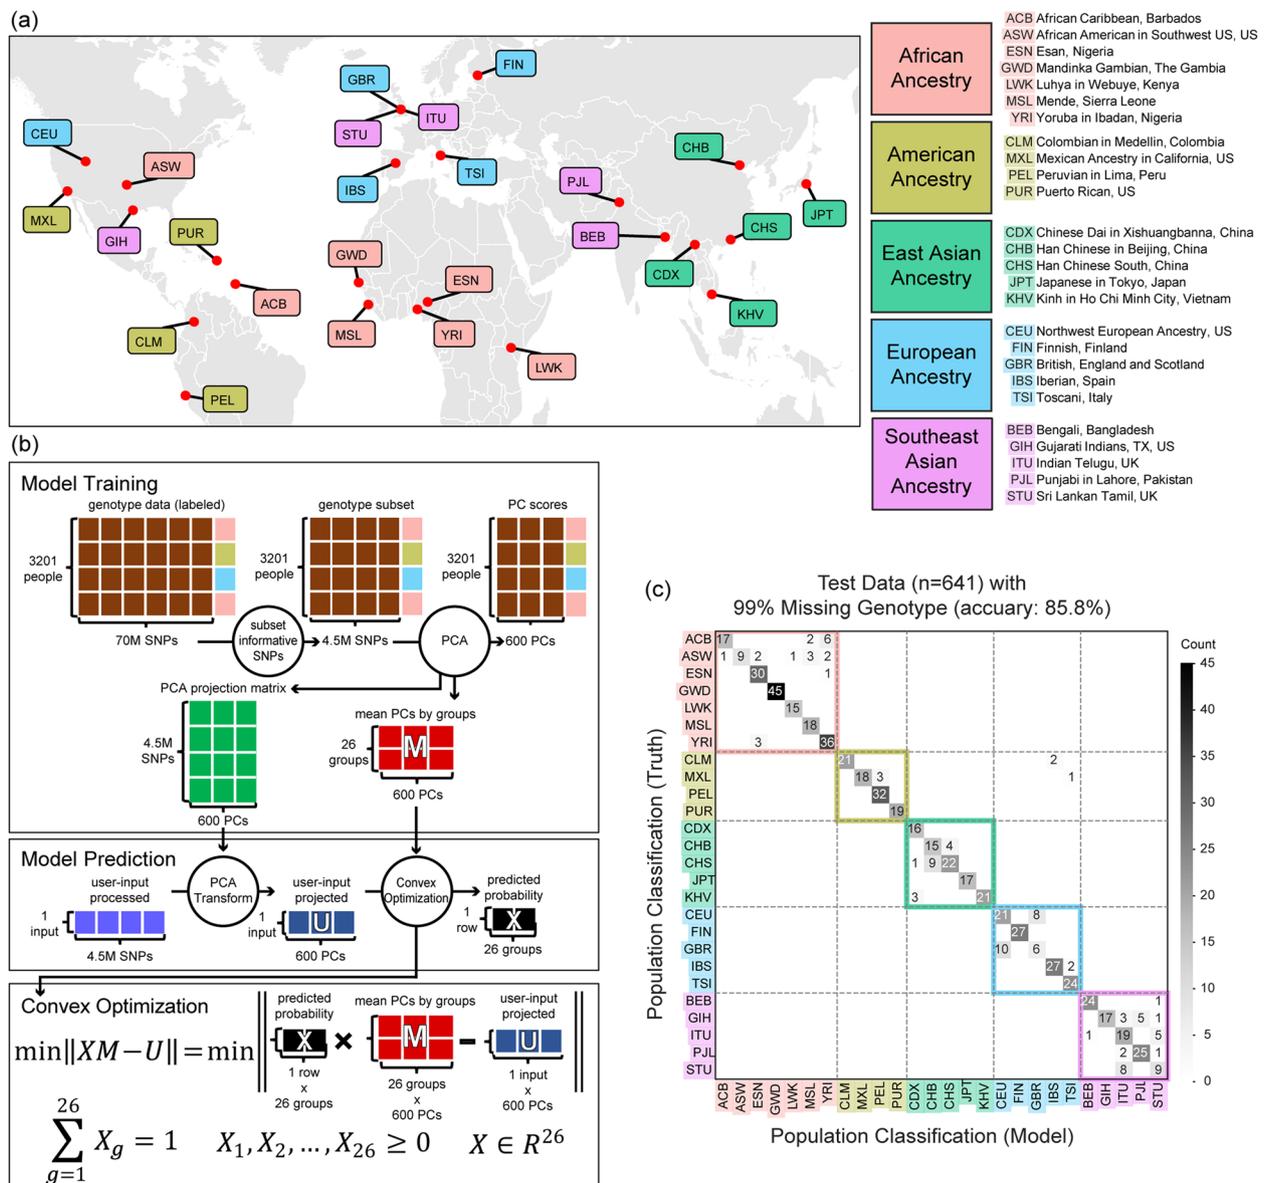
### Data source for training

The 1000 Genomes Project (1KGP), launched in 2008, has collected and published whole-genome sequence (WGS) data to enhance our understanding of the genetic contribution to human phenotypes [23]. The New York Genome Center (NYGC) used 3202 samples from this project to conduct whole-genome sequencing at 30X coverage, followed by variant calling, and has made the data publicly available [29]. We excluded one individual who was reported as admixed of two population groups, reducing the number of individuals to 3201, each of whom was assigned to a single population group from the 26 different population groups used in the study (Fig. 1a). From the 88 million genetic variations (SNPs, insertions and deletions, and structural variants) present in the data, we only used 70 million autosomal SNPs. 1KGP was chosen as the training data because it is the largest publicly available high quality WGS data with detailed racioethnic population group labels. We also examined the sensitivity of the model predictions to changes in the training data, as illustrated in Supplementary Fig. 2.

### Data pre-processing

To identify the ancestry-informative SNPs, we discarded SNPs where the most common genotype within a population was the same across all 26 population groups. This reduced the number of SNPs from approximately 70 million to 4.5 million; a set that we refer to as the ancestry-informative SNPs. Across the 3201 individuals, each of these SNPs was either the reference genotype (e.g. A) or a specific alteration (e.g. A to a C mutation). We did not see multiple alternative alleles in the dataset for a given SNP. Our procedure for reducing the number of SNPs may throw out SNPs that are unique to a particular group (and therefore informative about ancestry) but are not the most common genotype within the population. Nevertheless, we enforced this rule to minimize computation without comprising the accuracy of our inference. At least 74 individuals were included in each population group (geographic distribution shown in Fig. 1a and the numbers are listed in Supplementary Table 1).

As shown in Fig. 1a, the 26 populations can also be grouped by 5 different geographic regions: Africa, America, East Asia, Europe, and Southeast Asia, where “region” is defined as the regional ancestry, not where the individuals reside. For example, several populations in the United Kingdom, such as Indian Telugu and Sri Lankan Tamil, would be from Southeast Asian ancestry, not European.



**Fig. 1** Schematic of model training and inference. **a** The geographic locations of the 26 populations in the 1KGP dataset are shown in the truncated world map, with 5 color codes based on geographic regions. The right side shows the population description and three letter code that will be used throughout the paper. **b** Schematic of Model Training is shown on top. The training data is shown as a matrix of 3201 individuals by 70 million genotypes at SNP sites. We first subset the SNPs to keep 4.5 million ancestry-informative SNP sites as described in Sect. 2.2 of Methods. We then conducted PCA to reduce the dimensionality to 600 principal components (PCs). Model prediction is shown on the bottom. An example of a user input of 1 sample is shown although users can input an arbitrary number of samples. The input data is first centered using the mean vector from the training data, imputed for sites with missing genotypes, and scaled appropriately as described in Sect. 2.5 of Methods. Convex optimization is used to compute the contributions of the 26 population groups to the ancestry of the input data with the constraint that the contributions are non-negative and sum to 1. **c** The confusion matrix of classification results after splitting the 1KGP data into training-test (80–20%) is shown. Each test input is classified as the member of the population group with maximum contribution to that individual. Missing data was simulated by randomly removing 99% of the SNP sites from the data. Classification accuracy with 99% missing data is 86% (with no missing data, the accuracy is 89%), and all the misclassifications except for three cases are still within the same geographic region. Supplementary Fig. 1 shows the confusion matrix generated with different degrees of missing genotypes

Next, we converted the genotypes of the 4.5 million ancestry-informative SNPs for each individual from a vector of base-pairs to a vector of numbers. For a given individual and a given SNP, if both base pairs (maternal and paternal alleles) matched that of the reference genome GRCh38, we assigned a 0. If only one of the base pairs (either from the maternal or paternal allele) matched the reference, we assigned a 0.5. Otherwise, when both did not match the reference, we assigned a 1. This numerical conversion does not discard any information about the mutations, as described previously, we either see the reference or a single alternative allele for each SNP in the dataset. This is a reasonable numerical representation because the genotypes that an individual inherits are linear combinations of the genotypes of their parents.

Having obtained a numerical representation for the SNPs for every individual, we then centered the data. To do so, we thought of the data as a 3201-individual by 4.5-million SNP matrix. We computed the average of each column, corresponding to the average value of each SNP across the 3201 individuals. We then subtracted the average value of each column from all the entries of that column.

### Inference algorithm

In addition to the following description, the model training and inference algorithm are depicted visually in Fig. 1b. scAI-SNP leverages existing methods (principal component analysis and convex optimization) to infer ancestry.

#### Principal component analysis (PCA).

We used Principal Component Analysis (PCA) to transform the high-dimensional data (3201 individuals by 4.5 million features) into a lower-dimensional data of 3201 individuals by fewer features using an orthogonal transformation followed by truncating the feature space. This step is essential to reduce the computational burden when using the model to compute predictions for

user-provided inputs. However, conducting PCA on 4.5 million features is challenging because we have to compute a 4.5 million by 4.5 million mutation covariance matrix and then diagonalize that matrix. Instead, we computed the covariance across individuals which is a 3201 by 3201 matrix (referred to as the Gram matrix) and used this matrix to compute the principal components in the 4.5-million-dimensional feature space.

The centered data matrix  $X$  is  $n$  by  $p$ , where  $n$  is the number of individuals (3201 in this case), and  $p$  is the number of SNPs sites (4.5 million in this case). We can use singular value decomposition to express the data matrix as:

$$X = UDV^T$$

where columns of  $U$  are the eigenvectors of the gram matrix ( $XX^T$ ), columns of  $V$  are the eigenvectors of the covariance matrix ( $X^T X$ ), and  $D$  is a  $n$  by  $p$  matrix in which the first  $(n - 1)$  diagonal entries are square-roots of the eigenvalues of  $XX^T$  or  $X^T X$  and all other entries are zero. Our goal is to compute  $V'$ , a truncated version of  $V$  with the first 600 eigenvectors of the mutation covariance matrix, to conduct PCA. Because directly computing  $V$  is computationally expensive, we first compute  $U$ , the eigenvectors of the 3201 by 3201 Gram matrix. We keep the first 600 eigenvectors with the largest eigenvalues to construct a subset of  $U$ ,  $U'$ , as a 3201 by 600-dimensional matrix.  $D'$  is a diagonal 600 by 600-dimensional matrix, where the diagonal entries are the first 600 largest eigenvalues. The choice of truncating to 600 dimensions is based on a sensitivity analysis shown in the Supplementary Table 2.

We can express  $V'$  in terms of  $U'$ ,

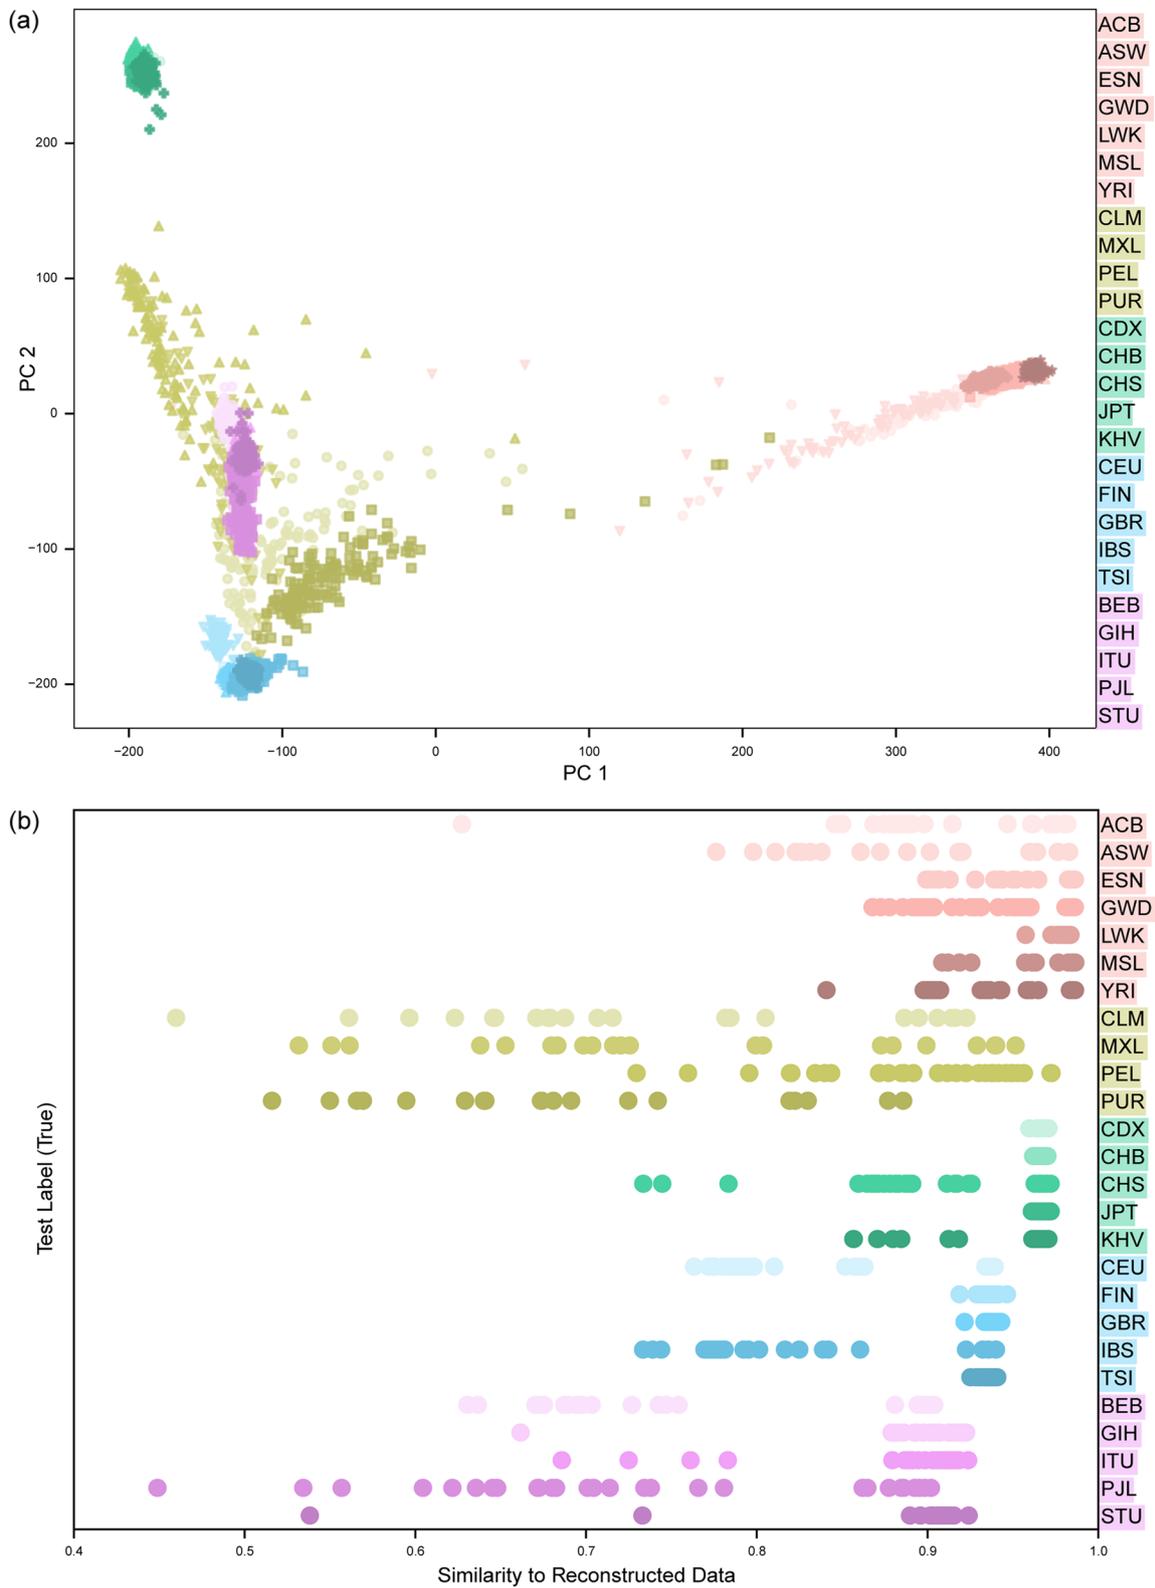
$$(V')^T = (D')^{-1}(U')^T X$$

where we have used the fact that  $(U')^T U' = I$  and  $(D')^{-1}$  is the inverse of matrix  $D'$ . It follows,

$$V' = X^T U'(D')^{-1}$$

(See figure on next page.)

**Fig. 2** PCA of the training data and the heterogeneity within each population group. **a** The first two principal components of the 1KGP data. The marker shape and colors capture the 26 different population groups, with a different color used for each of the 5 different geographic regions. The PCA plot shows a clear separation of the points by the 5 different regions. However, the PCA plot also shows the admixed nature of the populations and that many of the population groups form a continuum of points as opposed to distinct clusters. **b** After splitting the data into training-test (80–20%), we converted each test data point to a vector in the 600-dimensional PCA space obtained from the training data. We used convex optimization then to construct the closest possible vector to the test vector using a linear combination of the mean vectors of the 26 population groups, where the coefficients of the linear combination were constrained to be non-negative and sum to 1. For each test data point, we have plotted the cosine similarity of the vector corresponding to the test data point and its reconstruction. The values range from 0 to 1, where 1 indicates that the data can be perfectly explained by the mean PC vectors of the 26 population groups. Lower values of similarity result from heterogeneity of the data within each population group and lack of population groups in the training data that would explain this heterogeneity through admixture



**Fig. 2** (See legend on previous page.)

$V'$  can be used to project a vector in the original 4.5-million-dimensional feature space to a 600-dimensional space. We will use vectors in this lower-dimensional space to map vectors of alleles to their ancestry [31].

### Ancestry inference

To infer ancestry, our algorithm takes as input the data after PCA (600-dimensional input) and outputs the contributions of the 26 population groups to the ancestry of the input sample. This approach is appropriate because an individual's ancestry can belong to more than one of the population groups.

The predicted ancestry is computed by solving the following constrained optimization problem using convex optimization:

$$\min \|XM - P'\|_2, X \in R^{26}, X \geq 0, \sum_{g=1}^{26} X_g = 1$$

The parameters of interest,  $X$ , is a 26-dimensional non-negative row vector whose elements sum to 1, making it a vector of probabilities.  $X_g$  is the contribution of population group  $g$  to the input.  $M$  is a 26 by 600 matrix in which row  $i$  is computed by projecting all the individuals in population group  $i$  in the training data after pre-processing into the PC space and averaging them.  $M$  then captures the average representation of each population group in the training data.  $P'$  is a row vector of 600-dimensional user-input that has been pre-processed and projected into the PC space of the training data. Note that this optimization can be generalized for  $X$  to be a matrix of  $n$  row vectors, with  $n$  rows of user-input vectors ( $P'$ ) in a form of row matrix in the scAI-SNP method. The optimization problem that we solve is finding the linear combination of the vectors representing each population group in the training data that most closely resembles the input vector with the coefficients constrained to be non-negative and summing to 1. We implemented convex optimization using package `cvxpy` (version 1.3.2) in Python 3.7.4 with a L2 norm to solve for vector  $X$  [32]. In cases where we aimed to classify the input to only one population group, we selected the population group  $g$  for the input with maximum  $X_g$ .

### Imputation

Because scRNA-seq data is expected to be sparse, we need to account for missing genotypes by imputation. To test our imputation approach, we simulated various degrees of missing data. We first split the data randomly into 80–20% proportion of training and test data, respectively, and implemented varying amounts of missing data

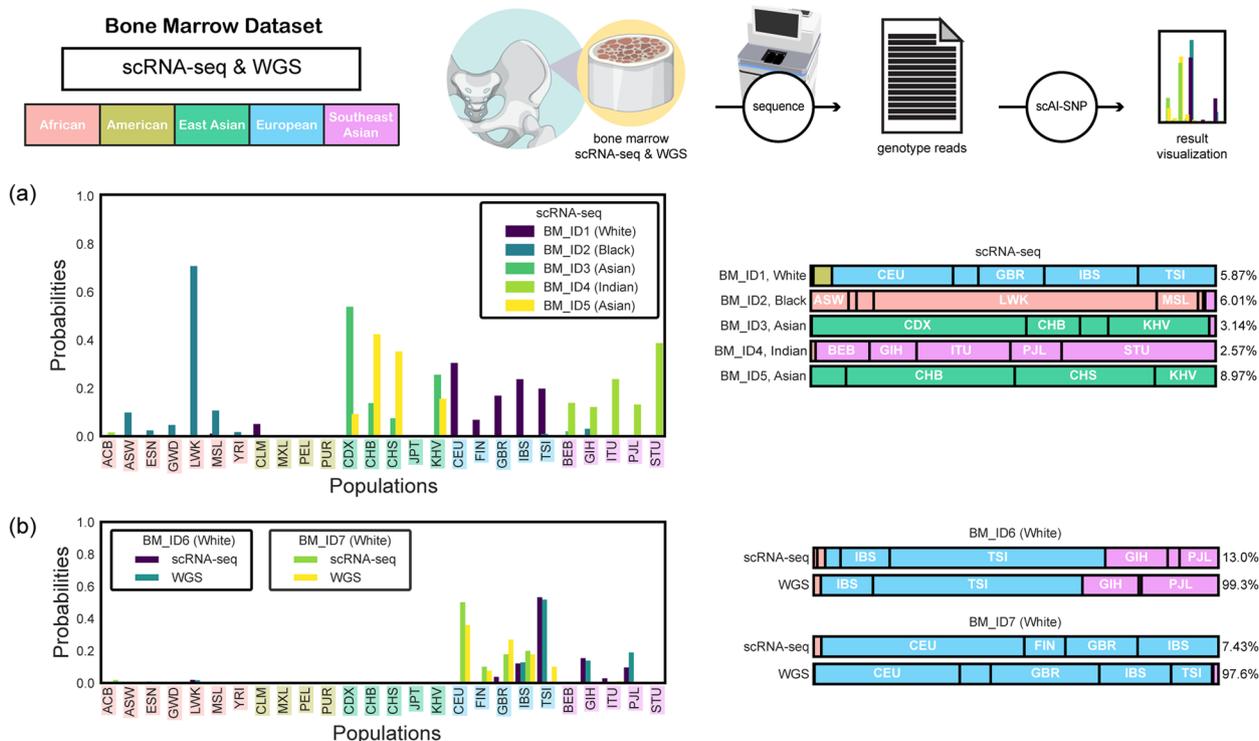
in the test data (90%, 99%, and 99.9% of the allele calls for ancestry-informative SNPs missing). To implement missing data, for each test datapoint, we randomly selected a subset of SNP sites and marked them as missing. This subset can vary from one datapoint to another. We first imputed all the missing values of the SNPs with the corresponding genotype mean for each SNP site in the training data and then centered the data. This ensures that the missing SNP sites are non-informative for dimensionality reduction and classification. After PCA projection, we corrected for the missing data by multiplying the inverse of the proportion of the observed data (for example a factor of 100, for 99% missing) to scale the principal components. Supplementary Fig. 3 shows the first four principal components for the original test data and the imputed and corrected test data. Even with 99% of data missing at random, the principal components in the first four dimensions are quite similar to those of test data. We computed the downstream effects of missing data on the inference accuracy and demonstrate them in Fig. 1c and Supplementary Fig. 1.

### Number of principal components

To determine the ideal number of principal components to use, we conducted a sensitivity analysis in which we split the 3201 individuals into training and test sets at random (80–20% split) and conducted dimensionality reduction and classification with varying numbers of principal components. As shown in Supplementary Table 2, using the first 600 principal components is sufficient as the inference accuracy does not increase significantly with additional principal components.

### Genotyping

Users may use their own genotyping algorithm to prepare inputs for scAI-SNP. The list of 4.5 million SNP sites and their information is provided as a text and Browser Extensible Data (BED) file in the scAI-SNP GitHub page [28]. To genotype SNP sites in single-cell datasets, we extended SComatic. Specifically, users can directly use SComatic to readily prepare their own single-cell data for scAI-SNP. SComatic conducts variant calling of the given 4.5 million sites with various requirements to confidently assign a genotype for each site. It requires 5 reads at minimum for proper genotype extraction, and if the variant base calls to read depth ratio is less than 0.1, the site is considered to be homozygous reference (0/0). Similarly, if the ratio is greater than or equal to 0.1 but less than 0.9, the SNP is considered to be heterozygous mutation (0/1) and homozygous mutation (1/1) if the ratio is greater than or equal to 0.9.



**Fig. 3** Validation of scAI-SNP for ancestry inference. **a** Inferred ancestry from single-cell RNAseq data of bone marrow mononuclear cells from 5 donors with known self-reported race. The vertical bar plots depict the predicted contribution of each of the 26 population groups to each donor. These probabilities are plotted again in horizontally stacked bar plots on the right in a more succinct form. The percentages on the right side are the fraction of the 4.5 million ancestry-informative SNP sites detected in each sample. **b** Same plots as in (a) for 2 samples from donors who self-reported as White. Ancestry inferred from whole genome sequencing of DNA extracted from the peripheral blood of these donors yielded results consistent with those from scRNA-seq samples

**Data sources for testing**

**Bone marrow dataset**

We obtained bone marrow aspirates from 5 healthy individuals (BM\_ID1 to BM\_ID5) of diverse backgrounds. These healthy donor samples were purchased as de-identified samples from the Boston Children’s Hospital. Mono-nuclear cells were isolated from each sample and processed as previously described in Egeren et al. [30]. We then used “Chromium Single Cell 3’ Library & Gel Bead Kit v3” (10xGenomics, #1000075) and followed the Chromium Single Cell 3’ Reagent Kit v3 user guide (10x Genomics, CG000183 Rev A) to generate scRNA-seq libraries. In addition to these 5 samples, we used scRNA-seq data and WGS profiling of bone-marrow cells from a previous study on JAK2-mutant myeloproliferative neoplasms (MPN) [30] as samples BM\_ID6 and BM\_ID7. The WGS had an average depth of 32x. For BM\_ID4 to ID7, CD34 positive cells were enriched from the mono-nuclear cells and used to generate single-cell libraries. To call the alleles for the ancestry-informative SNPs in the WGS data, we used Strelka2 and set a minimum genotype quality (Phred score) of 20 or higher with

a depth of 3 reads or higher. The ancestry-informative SNPs for the samples BM\_ID1 to BM\_ID7 are available for download on the Hormoz Lab GitHub page [28].

**Heart single-cell dataset**

A study of cardiovascular disease that sought to characterize genetic heterogeneity of heart cells conducted single-cell RNA sequencing of 14 donors with unremarkable cardiovascular disease history [31]. CD45 positive cells were enriched from various regions of the heart, processed using the 10x Genomics platform, and sequenced using HiSeq 4000 (Illumina) and NextSeq 500 (Illumina) with a minimum depth of 20,000 to 30,000 read pairs per cell. Each sample had about 31,600 cells on average. Our study used 76 samples that had at least 20,000 ancestry-informative SNPs detected, which were sampled from six different regions of the heart.

**GTEx dataset**

As part of the Genotype-Tissue Expression (GTEx) project, a study produced snRNA-seq profiles of eight tissue types from 16 donors [32]. Each of these snRNA-seq

samples has about 13,700 nuclei with 1,500 reads per nucleus on average. On average about 1.35% of the ancestry-informative SNP sites were detected in the data. We used 10 samples from 4 donors that had at least 20,000 detected sites.

#### **Ovarian cancer MSK SPECTRUM dataset**

We analyzed scRNA-seq data from 160 multi-site primary and metastasized tumor tissues originated from 42 patients with high-grade serous ovarian cancer (HGSOC) [33]. This dataset also had matched WGS data for normal peripheral blood mononucleated cells and tumor cells. Chromium Single-Cell 3' Reagent kit v3 (10x Genomics) was used to sequence the data. We used 268 samples from 26 sites and 41 patients in total. 6 samples that had fewer than 20,000 ancestry-informative SNP sites detected were discarded.

#### **Pan-tissue scATAC-seq dataset**

The pan tissue scATAC-seq data was generated from a study that created a cell atlas of *cis*-regulatory elements [34]. The study team conducted sci-ATAC-seq, a modified scATAC-seq that uses combinatorial indexing, with 92 samples from 30 tissue types [35]. We used 66 samples from 4 donors that were processed and used in the SCoatic study [24]. As with the other datasets, low-quality scATAC-seq samples, which had fewer than 20,000 of the 4.5-million ancestry-informative SNP sites detected, were discarded (47 samples out of 66 representing 22 tissue types passed this quality control).

## **Results**

We trained sCAI-SNP on the 1KGP data as described in Methods. Briefly, the data contains 70 million SNPs across 3201 individuals who belong to one of 26 population groups (Fig. 1a). We identified a subset of 4.5 million SNP sites that were informative for ancestry inference (the most common allele within each population group was not the same across all the groups) and captured each individual's genetic information into a vector of alleles of these 4.5 million SNP sites. We then used principal component analysis (PCA) to reduce the dimensionality of the allele space from 4.5 million to 600. To make an inference, the user provides an input vector corresponding to the alleles of the 4.5 million SNP sites for a given sample. This input can contain missing values if some of the SNP sites were not genotyped. Our algorithm replaces the values of the missing SNP sites with the mean values from the training set. The prediction is computed by multiplying this input vector by a 4.5 million by 600-dimensional matrix to obtain the projection of the input vector to the 600-dimensional PCA space. We appropriately scale the resulting principal components to account for the

missing data by multiplying by the inverse of the proportion of the observed data (for example a factor of 100, for 99% missing). We then use convex optimization to find the linear combination of the mean vectors of each population group that best approximates the input vector. We constrain the coefficients of this linear combination to be non-negative and sum to 1. The resulting 26-dimensional output corresponds to contribution of each of the 26 population groups to the donor from whom the sample was obtained (Fig. 1b).

To validate sCAI-SNP, we first generated synthetic test data from the 1KGP dataset itself. We split the dataset into a training-test set (80–20% split) and used the training set to train the model. We implemented missing alleles in the test data by dropping the value of 90%, 99%, and 99.9% of randomly chosen SNP sites for each test data point. With no missing data, the accuracy was 89%, which dropped to 86% with 99% of the data missing (Supplementary Fig. 1 and Fig. 1c). Therefore, the predictions are remarkably robust to missing data. Importantly, the incorrect predictions of the model are still consistent with expected human demographics and migration patterns. For example, as shown in Fig. 1c, individuals from the ASW (African American in the Southwest of the United States) group are sometimes mistakenly inferred to be from the ESN (Esan, Nigeria) and the YRI (Yoruba in Ibadana, Nigeria) groups. Similarly, individuals from the GBR (British, England and Scotland) group are mostly attributed to CEU (Northwest European Ancestry, the United States). Similarly, CHS (Han Chinese South) population and CHB (Han Chinese in Beijing) populations also demonstrate similarities and mixtures. From known migration histories of these populations, such predictions errors are not necessarily biological and can potentially be attributed to admixtures and limitations in the self-reported ancestry in the 1KGP data. We reasoned that a more complex model with higher accuracy will most likely overfit the 1KGP dataset and not generalize as well. Additionally, our linear model correctly infers the ancestry of mixed individuals (Supplementary Fig. 5). Therefore, we proceeded with our simple linear fit for ancestry inference. An alternative approach would have been to remove the recently admixed populations (such as ASW) from the inference or to predict ancestry up to the 5 geographic regions. We decided against this approach because, as we discuss in more detail in the Discussion, the predicted ancestry can also be used as a proxy for other contexts for an individual, such as environment and socioeconomic status. Therefore, we decided to predict the contributions of all 26 populations groups in the 1KGP dataset including the recently admixed population groups.

To determine the limitation of our method, namely approximating the ancestry of an individual as a linear combination of the mean of the 26 population groups in the 1KGP dataset, we first set out to characterize the heterogeneity in ancestry of individuals within the same population group. To do so, we plotted the first two principal components of the 3201 individuals in the 1KGP dataset. As shown in Fig. 2a, the PCA plot shows distinct groups that correspond primarily to the 5 geographic regions sampled. However, as previously reported [11], we observe heterogeneity within each population group and that many populations form a continuum of points as opposed to distinct clusters. To quantify the degree of heterogeneity in each population group, we split the 1KGP dataset to a training and test datasets (80–20% split) and computed the principal components only using the training data. We then computed the similarity of the test points to the mean of their population group by computing the cosine similarity between the vector corresponding to test point in the PC space and the mean of all vectors in the training data in the same population group as the test point (Supplementary Fig. 4). A value close to 1 indicates a high degree of similarity whereas a value close to 0 indicates low similarity. Certain population groups, particularly those with American and Southeast Asian ancestry, exhibited high heterogeneity in that many individuals in those population groups had low similarity to the population mean.

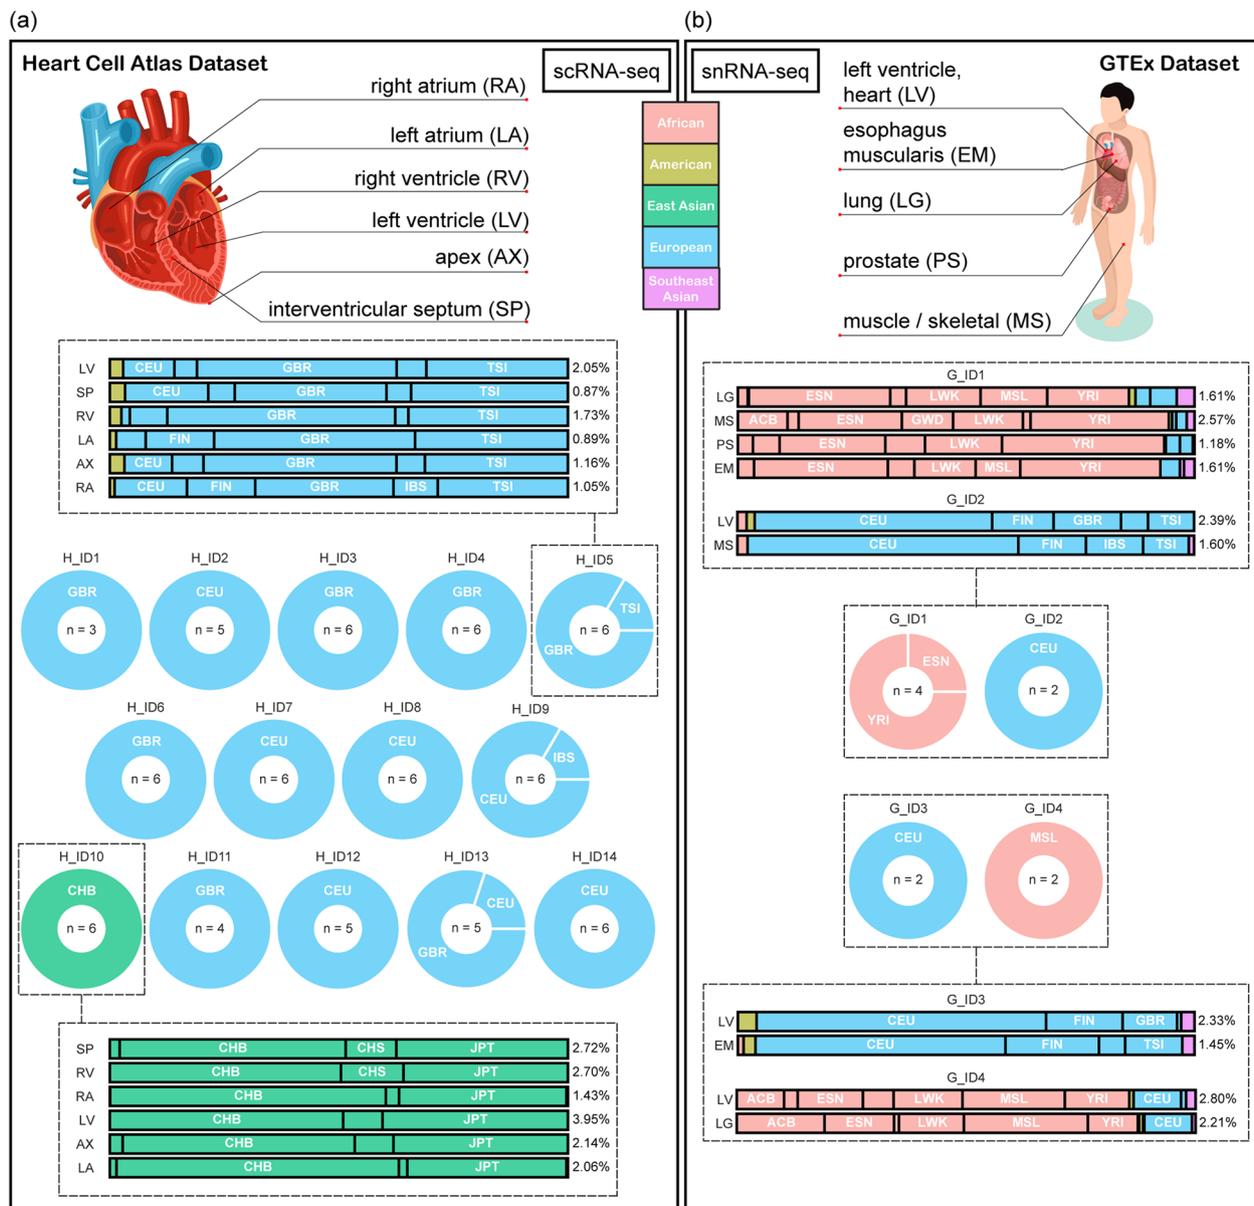
We reasoned that some of this heterogeneity could be explained by admixtures between the 26 population groups present in the training data. To remove this contribution from the heterogeneity measure, we used scAI-SNP to obtain the closest approximation of each test point as a linear combination of the mean vectors of the 26 population groups. Figure 2b shows the cosine similarity between each test vector and its best approximation. Although the level of heterogeneity decreased, individuals from American and Southeast Asian ancestry are still poorly approximated as linear combinations of the mean vectors of the 26 population groups, most likely because these groups are admixtures with populations groups other than the 26 used in the training data. Therefore, the inferred ancestry labels for these population groups (such as CLM, Colombian in Medellin, Colombia, and PJJ, Punjabi in Lahore, Pakistan) should be interpreted with caution because of the high degree of heterogeneity in the training data. We propose ways to supplement the 1KGP dataset in the Discussion section. Nevertheless, for many applications our algorithm is sufficiently accurate especially when inferring up to the 5 primary geographic regions.

To test the accuracy scAI-SNP on single-cell data and on individuals outside of the training data, we collected

bone-marrow biopsies on 7 individuals with self-reported race (3 White, 1 Black, 2 Asians, and 1 Indian; Methods). We performed single-cell RNA sequencing of bone marrow mono-nuclear cells (MNCs) from all these samples and applied scAI-SNP to the data.

As shown in Fig. 3a, the inferred ancestry of each sample is consistent with the self-reported race. To determine whether scAI-SNP predictions from sparse single-cell data is comparable with predictions from whole genome sequencing, we obtained WGS data from peripheral blood of 2 of the 7 individuals. The 4.5 million SNP sites in the WGS data was genotyped using Strelka2 (version 2.9.10) [36]. We obtained allele calls for 99.3% and 97.6% of the 4.5 million ancestry-informative SNP sites of the 2 donors from WGS data compared with 13.0% and 7.43% obtained from their single-cell data respectively. As shown in Fig. 3b, the inferred ancestry of these two individuals from single-cell RNAseq data agree with that inferred from WGS data despite the sparsity of the single-cell data. Taken together, scAI-SNP provides accurate inference of ancestry from single-cell RNAseq data that is consistent to that inferred from WGS data.

Next, we set out to determine whether scAI-SNP can infer ancestry from scRNA-seq data from a wide variety of cell types and tissues. To do so, we applied scAI-SNP to scRNA-seq data from the Heart Cell Atlas dataset [31]. We used data from 14 individuals who had between 3 and 6 scRNA-seq samples (78 in total) from cells sampled from different regions of the heart (Fig. 4 and Methods). These samples had on average ~31,600 cells and were sequenced with a minimum depth of 20,000 to 30,000 read pairs per cells. We detected on average 2.1% of the 4.5 million SNP sites in these samples. 2 samples that had fewer than 20,000 sites (or 0.44% site detected) were not used for further analysis. Figure 4a shows the predicted ancestry for each donor across all their samples as a pie-chart. The inferred ancestry was consistent across different samples from the same individual. The fact that 13 out of the 14 individuals in the Heart Cell Atlas Dataset is of European ancestry highlights the need for inferring ancestry of single-cell data to ensure diversity of atlases. To further validate this finding across more diverse tissues, we inferred ancestry of individuals in the GTEx dataset who had multiple single-nucleus RNA sequencing (snRNA-seq) data obtained from different tissues [32]. Each snRNA-seq sample had about 13,700 nuclei with 1,517 reads per nucleus on average, and we detected about 1.35% of the SNP sites and discarded the samples that didn't have at least 20,000 ancestry-informative SNP sites detected. We applied scAI-SNP to data from individuals that had samples from multiples tissues (4 individuals with 10 samples) as shown in Fig. 4b. Again, we observed consistent inferred ancestries across different



**Fig. 4** scAI-SNP results from single-cell RNA sequencing of a diverse set of cell types and tissues. **a** The inferred ancestry of 14 individuals (H\_ID1 to H\_ID14) from their scRNA-seq data in the Heart Cell Atlas dataset. Each individual had multiple samples obtained from different regions of the heart as shown approximately in the schematic. The pie charts show the distribution of classifications (population group with the highest predicted probability) across all the samples for each individual. The number of samples for each individual is shown in the center of the pie chart. For individuals H\_ID5 and H\_ID10, the detailed predictions are shown as the horizontally stacked probability plots, where the left labels indicate which region of the heart was sampled and the right shows the fraction of ancestry-informative SNP sites that were detected. **b** The inferred ancestry of snRNA-seq data from the GTEx dataset. The labels on the left side of the horizontal bar plots show the region from which the sample was obtained with the legend included with the simplified human schematic. The schematic is an approximation and does not show the exact regions from which the samples were obtained

tissues from the same individual. Taken together, scAI-SNP can consistently infer ancestry from scRNA-seq and snRNA-seq data from cells from a wide variety of tissues.

To determine whether we could accurately infer ancestry from scRNA-seq data of cancer cells, we analyzed the

MSK SPECTRUM Ovarian Cancer dataset [33]. From this dataset, we used data from 41 individuals who had had 3 to 12 scRNA-seq samples and corresponding WGS data (from both tumor and normal cells). The scRNA-seq samples contained about 17,100 cells on average and

had about 2.0% of the ancestry-informative SNP sites detected. As before, 6 samples that had fewer than 20,000 ancestry-informative SNP sites detected were discarded. As shown in Fig. 5, all 41 individuals had consistent ancestry inferred up to the 5 geographic regions. For 39 out of the 41 individuals, WGS data of both tumor and normal cells were available. Among these 39 individuals, 36 individuals' tumor and normal cell ancestry were consistent (the population group with the highest inferred probability was the same). The predicted ancestry of the three individuals that were inconsistent between tumor and normal cells were still within the same geographic region. Similarly, in the cases where the most probable inferred ancestry group from single-cell sample did not match that of the WGS normal data, the most probable groups were still within the same geographic region. Therefore, scAI-SNP can accurately and consistently infer ancestry from both healthy and tumor samples.

Finally, we set out to determine if scAI-SNP could be applied to scATAC-seq data. To do so, we applied our algorithm to scATAC-seq data derived from cells collected from various types of tissues from 4 donors in the Pan-Tissue scATAC-seq dataset [34]. The samples contained on average roughly 114,800 cells and were sequenced to a depth of 6,500 raw reads per nucleus on average. As before, low-quality scATAC-seq samples, which had fewer than 20,000 of the 4.5 million SNP sites detected, were discarded. As shown in Fig. 6, the inferred ancestry using scAI-SNP was consistent across all tissues for each of the donors. The inferred ancestries, up to the 5 major geographic regions, were consistent across all tissue types across all samples, as was the contribution of each of the 26 population groups to each sample qualitatively. Taken together with the results seen with other datasets, scAI-SNP can be applied to different modalities of the single-cell data, including scATAC-seq data.

## Discussion

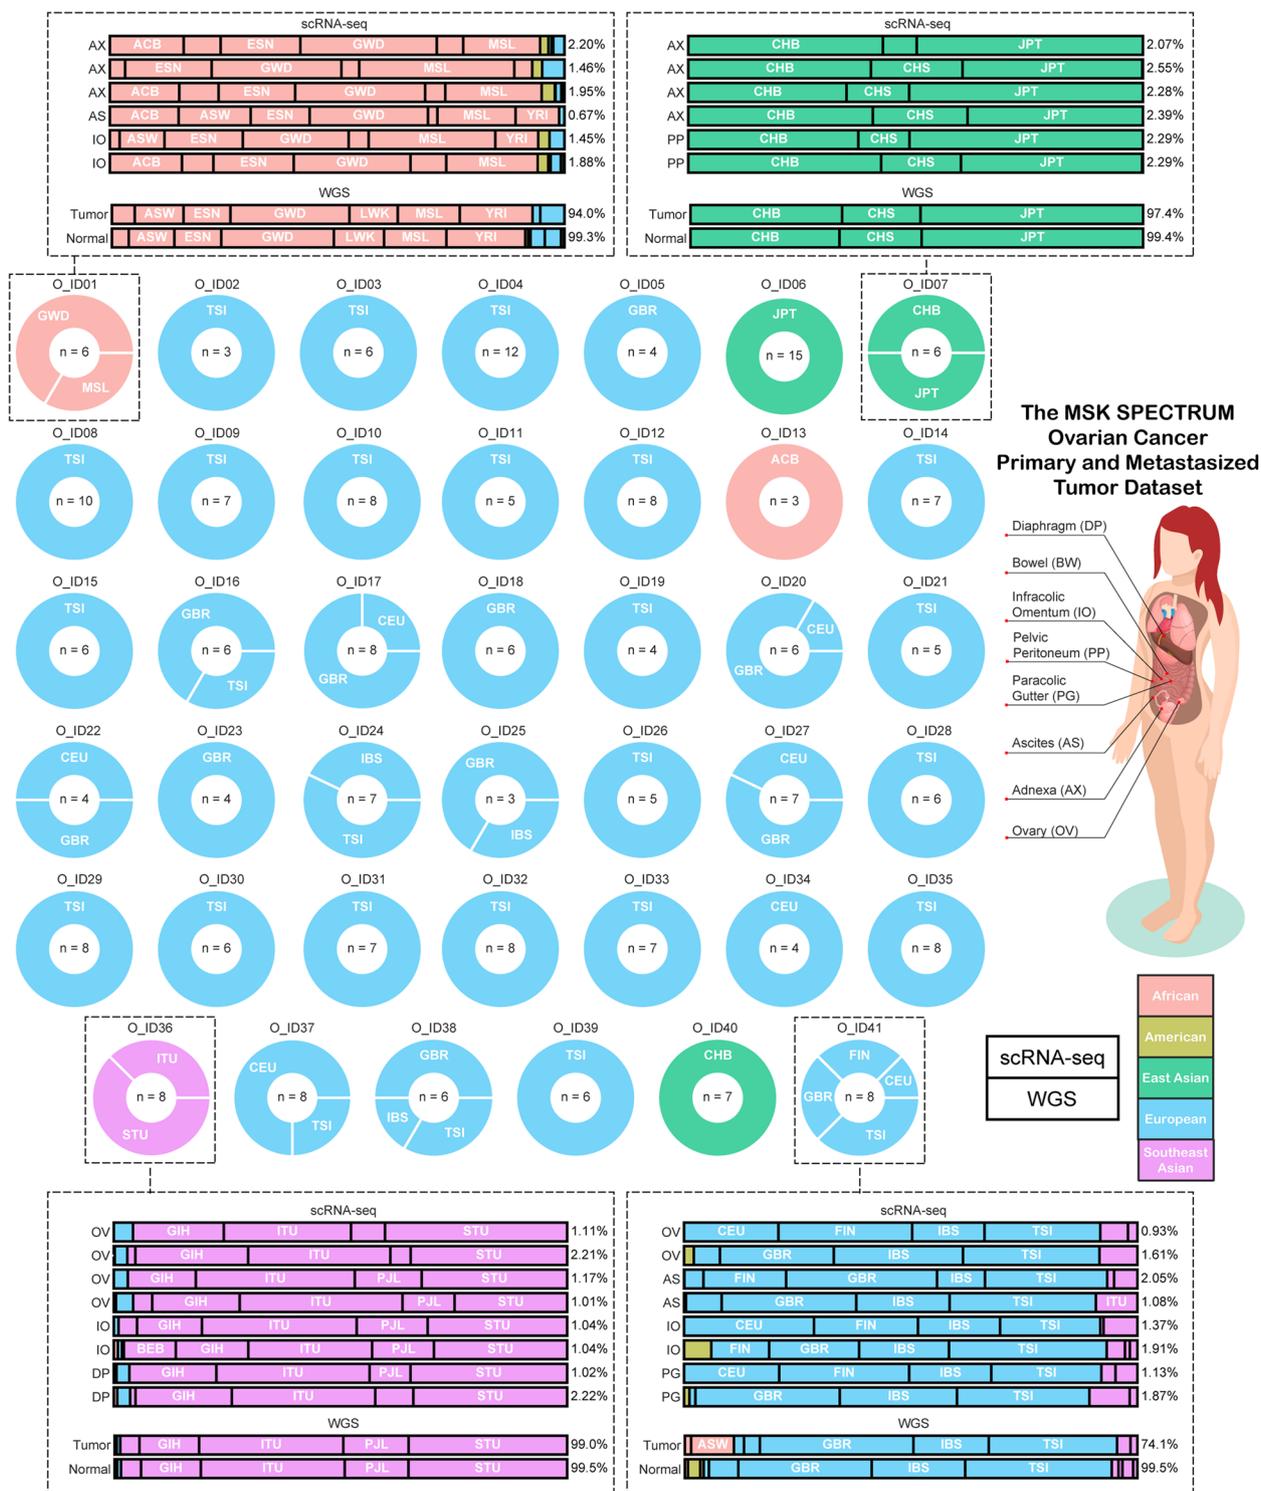
Inferring ancestry from single-cell data is essential for constructing single-cell atlases that reflect human diversity. Here, we introduced a tool, called scAI-SNP, that takes in single-cell genomic data as input, extracts alleles from 4.5 million ancestry-informative SNP sites, and outputs a distribution over 26 ancestry groups, which represents the inferred contribution of each ancestry group to the donor of the sample. We showed that scAI-SNP is accurate and robust to the sparsity present in single-cell data, and can be applied to samples from various cell and tissues and even cancer cells sequenced using WGS, scRNA-seq, snRNA-seq, and scATAC-seq.

The key feature of scAI-SNP is the simplicity in its ancestry inference. We used the data from 3201 individuals in the 1000 Genomes Project (1KGP) to identify 4.5

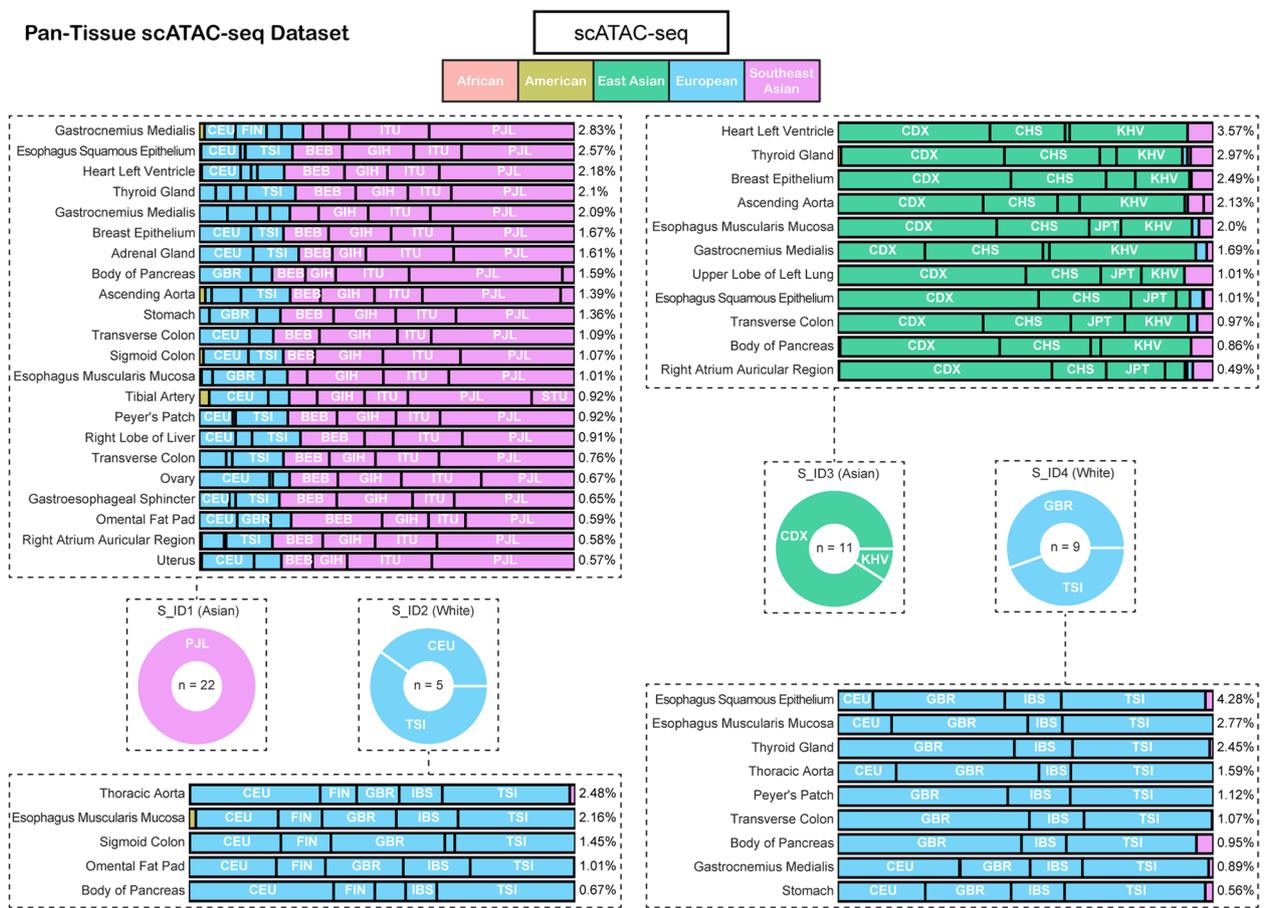
million ancestry-informative SNP sites to learn linear transformations followed by convex optimization that converts an input allele vector of 4.5 million dimensions to a vector of probabilities over the 26 ancestry groups. This simple approach is justified because population admixtures (e.g. individuals with ancestors from different population groups) can be thought of as linear combinations of their parental groups. In addition, this simple approach minimized the risk of over-fitting to the training data. This is especially important because the ground truth in 1KGP comes from self-reported ancestries and is therefore limited in its accuracy.

Our approach has certain limitations. First, our method ignores any linkage disequilibrium between the ancestry-informative SNP sites. Accounting for linkage disequilibrium can significantly improve the ability to distinguish a true mutation from a sequencing error in single-cell data as shown recently by Dou et al. [25], where the authors use their calls for germline mutations to validate global ancestry inference using PCA and carry out local ancestry inference using RFMix [18]. Although linkage disequilibrium is very informative if the goal is to call specific mutations with high confidence, for example for genetic association studies, we reasoned that for global ancestry inference the redundant information provided by the millions of ancestry SNP sites allow us to safely ignore linkage disequilibrium and use a much simpler approach.

Another important limitation of our work is that we cannot distinguish uncertainty in model prediction from admixtures of population groups. This is because the output of the model is a probability distribution over the 26 population groups in the training data (1KGP). A completely uninformative prediction would correspond to a uniform distribution over all 26 ancestry groups. An intermediate prediction of a distribution with significant weight on a subset of the ancestry groups can be interpreted both as model uncertainty across those groups or an individual who is an admixture of those groups. This issue is especially pronounced when the individual is from an ancestry group that is not one of the 26 population groups in the training data. Therefore, the biggest limitation of our method is the limited nature of our training data. 1KGP data has been analyzed for several studies of ancestry, but its 26 populations do not represent an ideal sample of human diversity around the globe because they do not include the full range of genetic diversity. Our analysis of test data points cosine similarity to the closest linear combination of the mean vector of the 26 population groups highlights this. The heterogeneity in many population groups is consistent with missing population groups and that labels for many population groups (such as CLM, Colombian in Medellin, Colombia, and PJI, Punjabi in Lahore, Pakistan) can



**Fig. 5** scAI-SNP results from scRNA-seq and WGS of cancer cells. Pie charts show the inferred ancestry (ancestry group with the largest predicted probability) across all the scRNA-seq tumor samples for a given individual across 41 individuals in the MSK SPECTRUM Ovarian Cancer dataset. The number of samples for each individual is shown in the center of the pie chart. The schematic shows approximately the different regions from which the samples were obtained. For individuals O\_ID01, O\_ID07, O\_ID36, and O\_ID41, the stacked predicted probabilities are shown along with the labels on the left that indicate which tissue was used and on the right that show the fraction of the 4.5 million ancestry-informative SNP sites detected in each sample. Note that multiple scRNA-seq samples may originate from the same site. For those 4 individuals, we also show the inferred ancestry using their respective WGS of normal and tumor tissues



**Fig. 6** scAI-SNP consistently infers ancestry from scATAC-seq data from various types of tissues. Inferred ancestries of 4 donors' scATAC-seq samples from various tissues are shown in the horizontal stacked plots. On the left side of the stacked bars are the description of tissues from which samples were collected and on the right side are the fraction of the 4.5 million ancestry-informative SNP sites detected in each sample. Each donor's inferred ancestry across all of their samples is also summarized in the pie charts in the middle

be misleading by neglecting the significant heterogeneity within those populations. Hence, the users should also consider the predicted probabilities at the regional level (Africa, Europe, East Asia, Southeast Asia, and America) when interpreting the results. In future iterations of scAI-SNP, we hope include more population from other studies of human genetic variation [37], and potentially even data from All of US and UK Biobank [27]. Nevertheless, our simple tool can provide invaluable insight into ancestry information of single-cell datasets and guide the generation of single-cell atlases to represent human genetic diversity.

Finally, we outline broadly the benefits of ensuring diverse ancestry when constructing single-cell atlases. As described in this paper, it is rather easy to identify genetic variants that are predictive of an individual's ancestry. However, this does not imply that traits, in particular neutral traits, can also be predicted from ancestry [38]. Intuitively, when we train models that predict ancestry

from millions of genetic variants (i.e., the ancestry-informative SNPs defined in this work) we are accumulating small contributions to ancestry of each variant across all the variants. The contribution from each variant is positive and can be summed together without cancellation. Although a single variant is not predictive, many of them together can make accurate predictions. A trait is not like ancestry. The contribution of each variant to a trait can be both positive and negative. Therefore, accumulating the contributions of many variants to a trait is no more accurate than predicting the trait from a single variant. Naively, this argument implies that although ancestry can be inferred from genetic variants, ancestry itself is not useful for capturing diversity of single-cell phenotypes.

There are two reasons why this naive argument is not correct. First, rare variants that contribute significantly to a particular trait and are present predominantly in specific population groups due to hard sweeps such as

selection or drift do occur and are clinically important. Examples are polymorphisms associated with Lactase persistence [39], Duffy allele [40] and G6PD deficiency that reduce the risk of malaria infection [41], and EPAS1 gene in Tibetans for adaptation to high altitude [42]. In the clinic, some variants are predictive of propensity for disease, such as the 8q24 variant, which is associated with increased risk of prostate cancer and more prevalent in African Americans compared with individuals of European ancestry [43]. Such variation across population groups may also result in incompatibility of reference ranges computed using one population group when applied to another. For example, blood counts of individuals of European descent are significantly different compared with those of African descent [7]. Taken together, single-cell atlases need to incorporate and account for rare genetic variants across population groups.

Second, ensuring diversity of ancestry can indirectly ensure diversity of environment and socioeconomic status of the donors of atlases. For example, consider polygenic scores that aggregate the contributions of many genetic variants to a trait. The accuracy of polygenic risk scores varies significantly with context, such as age and income, in addition to ancestry [44]. In the cases where context is not recorded, ancestry can act as a proxy for environmental and socioeconomic factors. This would allow us, for example, to study impact of environment on biological traits, elucidating the interactions between context and genetic variation. Diversity also allows us to remove noise from linkage disequilibrium when associating variants with traits. This is highlighted by the fact that accuracy of polygenic scores decreases when applied to individuals with increasing genetic distance to the training cohort [45]. Ideally ancestry should be supplemented with race and ethnicity. Race and ethnicity are even more closely correlated with environment and socioeconomic status than ancestry, and with other social determinants of health such as racism and discrimination [46]. Constructing single-cell atlases from donors with diverse ancestry, alongside race and ethnicity, should ensure that the scientific discoveries made using such as atlases and their eventual use in the clinic result in improved and equitable health outcomes.

#### Abbreviations

SNP	Single-nucleotide polymorphism
scRNA-seq	Single-cell RNA sequencing
scATAC-seq	Single-cell Assay for Transposase-Accessible Chromatin sequencing
GWAS	Genome-Wide Association Study
WGS	Whole Genome Sequencing
FRAPPE	Fast Randomized Approximate Principal component Extraction
PCA	Principal Component Analysis
1KGP	1000 Genomes Project
NYGC	New York Genome Center
BED	Browser Extensible Data file format

GTEx	Genotype-Tissue Expression project
MSK SPECTRUM	Memorial Sloan Kettering Cancer Center's spatiotemporal determinants of High Grade Serous Ovarian Cancer evolution, treatment and response
ASW	African American in the Southwest of the United States
ESN	Esan, Nigeria
YRI	Yoruba in Ibadan, Nigeria
GBR	British, England and Scotland
CEU	Northwest European Ancestry, the United States
CHS	Han Chinese South
CHB	Han Chinese in Beijing
CLM	Colombian in Medellin, Colombia
PJL	Punjabi in Lahore, Pakistan
MNC	Mono-nuclear cell
snRNA-seq	Single-nucleus RNA sequencing
ENA	European Nucleotide Archive
GEO	Gene Expression Omnibus
AnVIL	Analysis Visualization and Informatics Lab-space

#### Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s44330-025-00029-4>.

Supplementary Material 1.  
Supplementary Material 2.

#### Acknowledgements

We thank Sasha Gusev for helpful and insightful feedback on the manuscript.

#### Authors' contributions

S.C.H. and S.H. were responsible in designing and developing scAI-SNP and conducting relevant literature review. S.C.H. and S.H. were the primary contributors in writing the manuscript. S.C.H. trained scAI-SNP, generated inference model results, and was responsible in compiling the results in figures and writing. S.C.H. created the GitLab package. F.M. was responsible for applying the trained scAI-SNP for the test datasets, writing sections related to the test datasets, and revising the SComatic GitHub package to facilitate scAI-SNP users. I.C.-C. was responsible for providing analytical feedback, funding acquisition, and supervision. S.H. initiated the research question, was responsible for funding acquisition and supervision, and is the corresponding author. All authors read, edited, and contributed to writing of the final manuscript.

#### Funding

This work was supported by funding from the National Institutes of Health (NIH) National Heart, Lung, and Blood Institute grant numbers R01HL158269, the Chan Zuckerberg Foundation (Seed Grant) CZF2019-002433, and the Chan Zuckerberg Initiative (CZI) Seed Networks Supplemental Funding Opportunity (Tissue Diversity Supplement) 2020-224559. F.M. and I.C.-C. thank EMBL for funding.

#### Data availability

The code and scripts used for ancestry inference and plots can be found on our lab's GitHub page (<https://github.com/hormoz-lab/scAI-SNP>) [28]. The variants of 30x WGS data from the 1000 Genomes Project are publicly available at this link (<https://www.internationalgenome.org/data-portal/data-collection/30x-grch38>) [47]. The scRNA-seq and whole-genome sequencing data used for the validation shown in Figure 3 (BM\_ID6 and BM\_ID7, also referred to as ET1 and ET2) is available via dbGAP under study accession number phs002308. For the other bone marrow scRNA-seq data (BM\_ID1 to BM\_ID7), the genotype data is available on our lab's GitHub page (<https://github.com/hormoz-lab/scAI-SNP>) [28]. The scRNA-seq data of the Heart Cell Atlas is available on ENA (European Nucleotide Archive) with study accession number ERP123138 (<https://www.ebi.ac.uk/ena/browser/view/ERP123138>) [48]. The sequencing data of GTEx samples is available on AnVIL (Analysis Visualization and Informatics Lab-space) and requires a data access application via dbGap under study accession number phs000424 ([https://anvil.terra.bio/#works/paces/anvil-datastorage/AnVIL\\_GTEx\\_V9\\_hg38](https://anvil.terra.bio/#works/paces/anvil-datastorage/AnVIL_GTEx_V9_hg38)) [49]. The WGS and scRNA-seq of ovarian cancer dataset (MSK SPECTRUM) are available (<https://www.synapse.org/#!Synapse:syn11111111>).

se.org/#/Synapse:syn25569736/wiki/612269) [50]. The scATAC-seq data is available in the GEO database under accession number GSE184462. The reference genome used (GRCh38) is available (<https://hgdownload.soe.ucsc.edu/downloads.html>) [51].

## Declarations

### Ethics approval and consent to participate

All human data used in the study are second use from published data except for the bone marrow dataset detailed in methods section. For the bone marrow dataset, the study was approved by and conducted in accordance with the Declaration of Helsinki protocol (Dana-Farber Cancer Institute IRB protocol no. 01–206 and Massachusetts General Hospital protocol 13–583). All patients provided informed consent. The healthy donor samples were purchased as de-identified samples from the Boston Children's Hospital.

### Consent for publication

Not applicable.

### Competing interests

The authors declare no competing interests.

Received: 26 September 2024 Accepted: 1 May 2025

Published online: 19 May 2025

## References

- Regev A, Teichmann SA, Lander ES, Amit I, Benoist C, Birney E, et al. The human cell atlas. *eLife*. 2017;6:e27041.
- Rozenblatt-Rosen O, Stubbington MJT, Regev A, Teichmann SA. The human cell atlas: from vision to reality. *Nature*. 2017;550(7677):451–3.
- Rozenblatt-Rosen O, Regev A, Oberdoerffer P, Nawy T, Hupalowska A, Rood JE, et al. The human tumor atlas network: charting tumor transitions across space and time at single-cell resolution. *Cell*. 2020;181(2):236–49.
- Stranger BE, Forrest MS, Dunning M, Ingle CE, Beazley C, Thorne N, et al. Relative impact of nucleotide and copy number variation on gene expression phenotypes. *Science*. 2007;315(5813):848–53.
- Heyn H, Moran S, Hernando-Herraez I, Sayols S, Gomez A, Sandoval J, et al. DNA methylation contributes to natural human variation. *Genome Res*. 2013;23(9):1363–72.
- Hancock AM, Witonsky DB, Alkorta-Aranburu G, Beall CM, Gebremedhin A, Sukernik R, et al. Adaptations to climate-mediated selective pressures in humans. *PLoS Genet*. 2011;7(4): e1001375.
- Beutler E, Waalen J. The definition of anemia: what is the lower limit of normal of the blood hemoglobin concentration? *Blood*. 2006;107(5):1747–50.
- Hershman D, Weinberg M, Rosner Z, Alexis K, Tiersten A, Grann VR, et al. Ethnic neutropenia and treatment delay in African American women undergoing chemotherapy for early-stage breast cancer. *J Natl Cancer Inst*. 2003;95(20):1545–8.
- Peng L, Yang J, Lu X, Okada T, Kondo T, Ruan C, et al. Effects of biological variations on platelet count in healthy subjects in China. *Thromb Haemost*. 2004;91(02):367–72.
- Eller LA, Eller MA, Ouma B, Kataaha P, Kyabaggu D, Tumusiime R, et al. Reference intervals in healthy adult Ugandan blood donors and their impact on conducting international vaccine trials. *PLoS One*. 2008;3(12): e3919.
- Wojcik GL, Graff M, Nishimura KK, Tao R, Haessler J, Gignoux CR, et al. Genetic analyses of diverse populations improves discovery for complex traits. *Nature*. 2019;570(7762):514–8.
- Pakstis AJ, Speed WC, Soundararajan U, Rajeevan H, Kidd JR, Li H, et al. Population relationships based on 170 ancestry SNPs from the combined Kidd and Seldin panels. *Sci Rep-UK*. 2019;9(1):18874.
- Kidd KK, Speed WC, Pakstis AJ, Furtado MR, Fang R, Madbouly A, et al. Progress toward an efficient panel of SNPs for ancestry inference. *Forensic Sci Int Genetics*. 2014;10:23–32.
- Kosoy R, Nassir R, Tian C, White PA, Butler LM, Silva G, et al. Ancestry informative marker sets for determining continental origin and admixture proportions in common populations in America. *Hum Mutat*. 2009;30(1):69–78.
- Alexander DH, Novembre J, Lange K. Fast model-based estimation of ancestry in unrelated individuals. *Genome Res*. 2009;19(9):1655–64.
- Pritchard JK, Stephens M, Donnelly P. Inference of population structure using multilocus genotype data. *Genetics*. 2000;155(2):945–59.
- Price AL, Patterson NJ, Plenge RM, Weinblatt ME, Shadick NA, Reich D. Principal components analysis corrects for stratification in genome-wide association studies. *Nat Genet*. 2006;38(8):904–9.
- Maples BK, Gravel S, Kenny EE, Bustamante CD. RFMix: a discriminative modeling approach for rapid and robust local-ancestry inference. *Am J Hum Genet*. 2013;93(2):278–88.
- Tang H, Peng J, Wang P, Risch NJ. Estimation of individual admixture: analytical and study design considerations. *Genet Epidemiology*. 2005;28(4):289–301.
- Raj A, Stephens M, Pritchard JK. fastSTRUCTURE: variational inference of population structure in large SNP data sets. *Genetics*. 2014;197(2):573–89.
- Patterson N, Price AL, Reich D. Population structure and eigenanalysis. *Plos Genet*. 2006;2(12):e190.
- McVean G. A genealogical interpretation of principal components analysis. *PLoS Genet*. 2009;5(10): e1000686.
- Durbin RM, Altshuler D, Durbin RM, Abecasis GR, Bentley DR, Chakravarti A, et al. A map of human genome variation from population-scale sequencing. *Nature*. 2010;467(7319):1061–73.
- Muyas F, Sauer CM, Valle-Inclán JE, Li R, Rahbari R, Mitchell TJ, et al. De novo detection of somatic mutations in high-throughput single-cell profiling data sets. *Nat Biotechnol*. 2024;42(5):758–767.
- Dou J, Tan Y, Kock KH, Wang J, Cheng X, Tan LM, et al. Single-nucleotide variant calling in single-cell sequencing data with Monopogen. *Nat Biotechnol*. 2024;42(5):803–812.
- Arriaga-Mackenzie IS, Matesi G, Chen S, Ronco A, Marker KM, Hall JR, et al. Summix: a method for detecting and adjusting for population structure in genetic summary data. *Am J Hum Genet*. 2021;108(7):1270–82.
- Privé F. Using the UK Biobank as a global reference of worldwide populations: application to measuring ancestry diversity from GWAS summary statistics. *Bioinformatics*. 2022;38(13):3477–80.
- Hong SC. scAI-SNP: a method for inferring ancestry from single-cell data. 2023. Available from: <https://github.com/hormoz-lab/scAI-SNP>.
- Byrská-Bishop M, Evani US, Zhao X, Basile AO, Abel HJ, Regier AA, et al. High-coverage whole-genome sequencing of the expanded 1000 Genomes Project cohort including 602 trios. *Cell*. 2022;185(18):3426–3440.e19. Available from: <https://www.sciencedirect.com/science/article/pii/S0092867422009916>.
- Egeren DV, Escabi J, Nguyen M, Liu S, Reilly CR, Patel S, et al. Reconstructing the lineage histories and differentiation trajectories of individual cancer cells in myeloproliferative neoplasms. *Cell Stem Cell*. 2021;28(3):514–23.
- Litvínková M, Talavera-López C, Maatz H, Reichart D, Worth CL, Lindberg EL, et al. Cells of the adult human heart. *Nature*. 2020;588(7838):466–72.
- Eraslan G, Drokhlyansky E, Anand S, Fiskin E, Subramanian A, Slyper M, et al. Single-nucleus cross-tissue molecular reference maps toward understanding disease gene function. *Science*. 2022;376(6594):eab14290.
- Vázquez-García I, Uhlitz F, Ceglita N, Lim JL, Wu M, Mohibullah N, et al. Ovarian cancer mutational processes drive site-specific immune evasion. *Nature*. 2022;612(7941):778–86.
- Zhang K, Hocker JD, Miller M, Hou X, Chiou J, Poirion OB, et al. A single-cell atlas of chromatin accessibility in the human genome. *Cell*. 2021;184(24):5985–6001.
- Cusanovich DA, Daza R, Adey A, Pliner HA, Christiansen L, Gunderson KL, et al. Multiplex single-cell profiling of chromatin accessibility by combinatorial cellular indexing. *Science*. 2015;348(6237):910–4.
- Kim S, Scheffler K, Halpern AL, Bekritsky MA, Noh E, Källberg M, et al. Strelka2: fast and accurate calling of germline and somatic variants. *Nat Methods*. 2018;15(8):591–4.
- Bergström A, McCarthy SA, Hui R, Almarri MA, Ayub Q, Danecek P, et al. Insights into human genetic variation and population history from 929 diverse genomes. *Science*. 2020;367(6484):eaay5012.
- Edge MD, Rosenberg NA. A general model of the relationship between the apportionment of human genetic diversity and the apportionment of human phenotypic diversity. *Hum Biol*. 2015;87(4):313–37.

39. Anguita-Ruiz A, Aguilera CM, Gil Á. Genetics of lactose intolerance: an updated review and online interactive world maps of phenotype and genotype frequencies. *Nutrients*. 2020;12(9): 2689.
40. Hamblin MT, Thompson EE, Rienzo AD. Complex signatures of natural selection at the duffy blood group locus. *Am J Hum Genet*. 2002;70(2):369–83.
41. Tishkoff SA, Varkonyi R, Cahinhinan N, Abbes S, Argyropoulos G, Destro-Bisol G, et al. Haplotype diversity and linkage disequilibrium at human G6PD: recent origin of alleles that confer malarial resistance. *Science*. 2001;293(5529):455–62.
42. Simonson TS, Yang Y, Huff CD, Yun H, Qin G, Witherspoon DJ, et al. Genetic evidence for high-altitude adaptation in Tibet. *Science*. 2010;329(5987):72–5.
43. Amundadottir LT, Sulem P, Gudmundsson J, Helgason A, Baker A, Agnarsson BA, et al. A common variant associated with prostate cancer in European and African populations. *Nat Genet*. 2006;38(6):652–8.
44. Hou K, Xu Z, Ding Y, Harpak A, Pasaniuc B. Calibrated prediction intervals for polygenic scores across diverse contexts. *medRxiv*. 2023;2023.07.24.23293056.
45. Ding Y, Hou K, Xu Z, Pimplaskar A, Petter E, Boulier K, et al. Polygenic scoring accuracy varies across the genetic ancestry continuum. *Nature*. 2023;618(7966):774–81.
46. Borrell LN, Elhawary JR, Fuentes-Afflick E, Witonsky J, Bhakta N, Wu AHB, et al. Race and genetic ancestry in medicine — a time for reckoning with racism. *N Engl J Med*. 2021;384(5):474–80.
47. 30x WGS data from the 1000 Genomes Project. Available from: <https://www.internationalgenome.org/data-portal/data-collection/30x-grch38>. Cited 2024 Feb 29.
48. scRNA-seq data of the heart cell atlas on ENA. Available from: <https://www.ebi.ac.uk/ena/browser/view/ERP123138>. Cited 2024 Feb 29.
49. The sequencing data of GTEX samples. Available from: [https://anvil.terra.bio/#workspaces/anvil-datastorage/AnVIL\\_GTEX\\_V9\\_hg38](https://anvil.terra.bio/#workspaces/anvil-datastorage/AnVIL_GTEX_V9_hg38). Cited 2024 Feb 29.
50. The WGS and scRNA-seq of ovarian cancer data set (MSK SPECTRUM). Available from: <https://www.synapse.org/#!Synapse:syn25569736/wiki/612269>. Cited 2024 Feb 29.
51. Reference genome data (GRCh38). Available from: <https://hgdownload.soe.ucsc.edu/downloads.html>. Cited 2024 Feb 29.

## Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.