



OPEN

Enabling population assignment from cancer genomes with SNP2pop

Qingyao Huang ^{1,2} & Michael Baudis ^{1,2*}

In many cancers, incidence, treatment efficacy and overall prognosis vary between geographic populations. Studies disentangling the contributing factors may help in both understanding cancer biology and tailoring therapeutic interventions. Ancestry estimation in such studies should preferably be driven by genomic data, due to frequently missing or erroneous self-reported or inferred metadata. While respective algorithms have been demonstrated for baseline genomes, such a strategy has not been shown for cancer genomes carrying a substantial somatic mutation load. We have developed a bioinformatics tool for the assignment of population groups from genome profiling data for both unaltered and cancer genomes. Despite extensive somatic mutations in the cancer genomes, consistency between germline and cancer data reached of 97% and 92% for assignment into 5 and 26 ancestral groups, respectively. Comparison with self-reported meta-data estimated a matching rate between 88–92%, mostly limited by interpretation of self-reported ethnicity labels compared to the standardized mapping output. Our SNP2pop application allows to assess population information from SNP arrays as well as sequencing platforms and to estimate the population structure in cancer genomics projects, to facilitate research into the interplay between ethnicity-related genetic background, environmental factors and somatic mutation patterns in cancer biology.

Cancer arises from the accumulation of genomic aberrations in dividing cells of virtually all types of proliferating tissues (somatic variations). The irregular cellular expansion and other hallmarks of cancer can result from a plethora of mechanisms affecting multiple cellular processes¹. Some oncogenetic pathways can be initiated by exogenous factors, e.g. tobacco smoke or ultraviolet radiation². However, exposure to carcinogenics varies in its effects for people from different genetic background, which suggests that somatic mutations can be influenced by inherited (“germline”) variations^{3,4}.

Cancer studies have reported significant world-wide variation in incidence and prognosis between ethnicity groups^{5–8}. While such differences have been attributed to unequal social and economical circumstances which influence risk factors and therapeutic interventions, several studies have shown impact of population specific genomic variants with predisposing effects on malignant transformation and phenotypic behaviour^{9–12}. Due to the late onset of most cancers, even high-penetrance Mendelian-type variants may not be purged by natural selection and can accumulate in particular populations. Such variants may play key roles in cancer development¹³. Notably, mutations on BRCA1/2 genes confer a high risk to develop breast and ovarian carcinomas. Three founder mutations in Ashkenazi Jewish population cause the BRCA1/2 mutation prevalence to be 10-fold higher than all sporadic mutations in the general population^{14,15}. Mitochondrial aldehyde dehydrogenase (ALDH2) encodes an enzyme in alcohol metabolism. One variant, ALDH2*504Lys, which increases risk for alcohol-related liver, colorectal and esophageal cancer by alcohol consumption, has 36% prevalence in East Asian populations^{16,17}.

Many other studies have reported prevalent genetic variants in specific population groups which may contribute to the observed disparities in occurrence and prognosis among populations^{18–20}. As patients of European ancestry are overrepresented in medical research and clinical studies, we have a biased view on genetic aetiology in diseases and available therapeutic approaches may not benefit patients with other genetic backgrounds. Studying less explored ancestral backgrounds may lead to discovery of new biological mechanisms and therapeutic prospects. Although high-impact monogenic germline variants are rarely expected in cancer, polygenic variation models for breast cancer which estimate the combined risk profile of multiple loci have been developed, suggesting the potential of exploring genome-wide risk profiles²¹. Another study on germline background of

¹Institute of Molecular Life Science, University of Zurich, Winterthurerstrasse 190, 8057, Zurich, Switzerland. ²SIB, Swiss Institute of Bioinformatics, Winterthurerstrasse 190, 8057, Zurich, Switzerland. *email: mbaudis@imls.uzh.ch

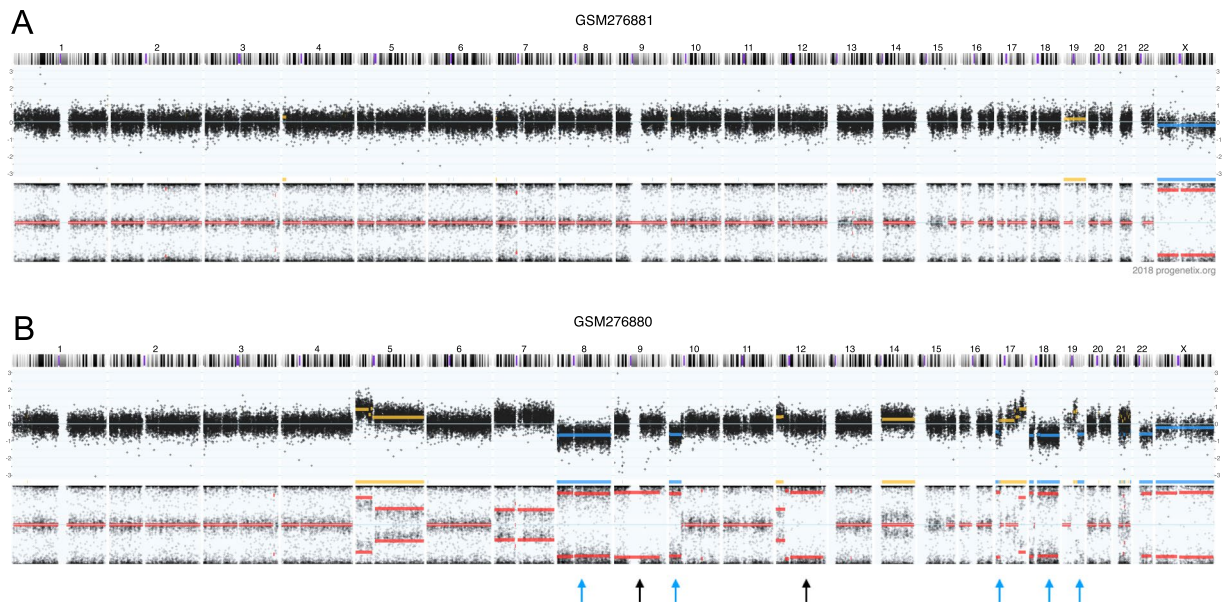


Figure 1. CNV tracks for a paired normal/cancer samples deposited in GEO with sample IDs GSM276881 and GSM276880. GSM276880 is a glioblastoma sample. GSM276881 is its peripheral blood control. The CNV tracks are consisted of two panels. Upper panel indicates the total copy number, represented by logR ratio ($\log_2 \frac{\text{copy number}}{\text{reference copy number}}$ at any probe) Namely, $\log R = 0$ indicates a position with normal two copies, $\log R > 0$ indicates a copy number gain and $\log R < 0$ indicates a copy number loss. Lower panel indicates the allele specific copy number, represented by B-allele frequency (BAF, $\frac{\text{No. } B \text{ allele copy}}{\text{Total copy number}}$). Any given SNP position can have a value between 0 and 1. A line at 0.5 indicates a heterozygous region. Compared to the unaltered genome (A), the cancerous counterpart (lower) has copy number loss in chr8, 10p, 18, 19qter, 22q and copy-neutral loss of heterozygosity in chr9,12q, accounting for 18.2% of genome.

cancer genomes has also identified disease-associated chromosomal regions from only seven individual samples with linkage mapping²². This type of studies can be conducted population-wise, with sufficient number of samples from the same population or ethnicity group.

With the increasing number of available genome profiles and the decreasing cost to genotype clinical samples, the stratification between patients' genetic backgrounds has become feasible with the promise to guide therapeutic strategies and improve the clinical prognoses. Since several studies have demonstrated the relevance of considering an individual's genomic origin for preventive screening (reviewed in Foulkes *et al.*¹⁵), information about the population background of cancer patients may be an additional factor for individual therapeutic decisions as well as for the stratification of clinical study cohorts. A meta-analysis addressing the interplay between genetic background, cancer development and therapeutic responses is desirable, not only for robust statistical associations in molecular target identification, but also for the rational design of studies incorporating informative biosamples.

For many data repositories, "population group" of a sample can be assumed based on a geographical location associated with the sample. Alternatively, a self-reported "race" category is commonly used in the U.S. census data^{23,24}. A biosample's geographic origin is often approximated using the location of the study's research facility or the contact address of its main authors²⁵. However, while these data can be easily retrieved, they may not provide an accurate representation of patients' origins for the purpose of population-specific ancestry mapping. Self-reported data are often inconsistent across studies, vague in category description (e.g. "white", "black" v.s. "Caucasian", "African") and misleading when patients have incomplete awareness of the migration and admixture histories of their ancestors. Overall, when associating oncogenic molecular signatures with germline variations, information from the above sources lacks in relevant detail and consistency.

A better approach to population assessment would be the computational estimate of ancestry with population-specific genomic variants. This has been shown previously for germline profiles, achieving 90% accuracy to distinguish three populations, African American, Asian and Caucasian, by using as few as 100 population-diverging single nucleotide polymorphisms (SNPs)²⁶, and nowadays is a standard methodology with claimed better granularity behind a number of commercial "ancestry" services, e.g. 23andme.com, myheritage.com and ancestry.com. We hypothesise that a similar strategy can be applied to cancer genome data, despite the additional cancer-related somatic mutations which leads to both information loss (e.g. large scale homozygous or allelic deletions) and added noise (e.g. somatic mutations masking germline variants). An example of a cancer genome containing copy number loss and copy-neutral loss of heterozygosity (CN-LOH) events and its paired normal sample is shown in Fig. 1. It is a typical case of altered genome with copy number loss (blue arrow) with CN loss regions partly recovered by doubling of the second allele (CN-LOH, black arrow) at multiple genomic locations. In addition to a general test of feasibility of inferring population background from the noisy cancer genome data, we continue with benchmarking population mapping procedures for heterogeneous datasets from different genotyping platforms, with the aim of integrating cross-platform cancer genotyping data for meta-analysis.

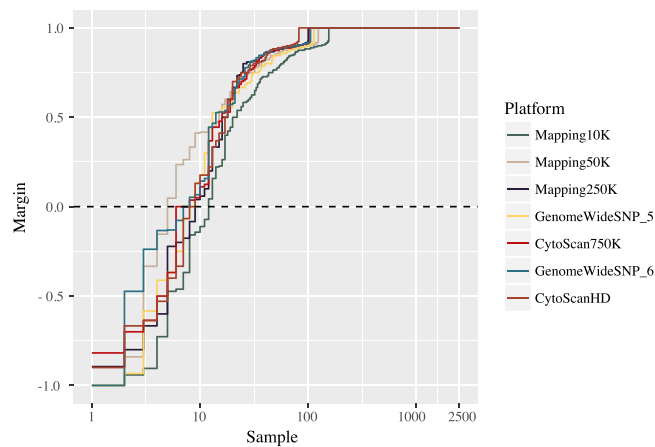


Figure 2. Estimating the model accuracy with reference samples with SNPs extracted from genotyping platforms of seven different resolution levels. Here we use margin as an estimator for the accuracy of the prediction model with random forest classifier output. It is defined as difference between the proportion of votes for the correct class minus maximum proportion of votes for the other classes. A positive value indicates correct prediction. For 2504 samples in the 1000 Genome Project, the number of incorrectly predicted samples ranged from 4 to 11, corresponding to an error rate of 0.16% to 0.4%. Note that this is cross validation error on reference dataset, rather than real data.

Results

We developed SNP2pop pipeline to assign genomic data to ancestral groups defined by 1000 Genomes Project²⁷, i.e. five continental groups and 26 population groups, respectively: 1. Admixed American (AMR) includes: Colombians from Medellin, Colombia (CLM), Mexican Ancestry from Los Angeles USA (MXL), Peruvians from Lima, Peru (PEL) and Puerto Ricans from Puerto Rico (PUR); 2. African (AFR) includes: African Caribbeans in Barbados (ACB), Americans of African Ancestry in SW USA (ASW), Esan in Nigeria (ESN), Gambian in Western Divisions in the Gambia (GWD), Luhya in Webuye, Kenya (LWK), Mende in Sierra Leone (MSL), and Yoruba in Ibadan, Nigeria (YRI). 3. East Asian (EAS) includes: Chinese Dai in Xishuangbanna, China (CDX), Han Chinese in Beijing, China (CHB), Southern Han Chinese (CHS), Japanese in Tokyo, Japan (JPT) and Kinh in Ho Chi Minh City, Vietnam (KHV); 4. European (EUR) includes Utah Residents with Northern and Western European Ancestry (CEU), Finnish in Finland (FIN), British in England and Scotland (GBR), Iberian Population in Spain (IBS) and Toscani in Italia (TSI); 5. South Asian (SAS) includes: Bengali from Bangladesh (BEB), Gujarati Indian from Houston, Texas (GIH), Indian Telugu from the UK (ITU), Punjabi from Lahore, Pakistan (PJT) and Sri Lankan Tamil from the UK (STU).

We benchmarked this tool with various normal and cancer samples from independent datasets to demonstrate the feasibility and reliability of this approach.

Cross-platform benchmarking. We first used the original data from 1000 Genomes to validate the level of resolution needed for accurate population assignment from the pipeline. Taking the sequencing data of 2504 samples and extracting the SNPs from the nine genotyping platforms of interest gave rise to the dataset in this benchmarking experiment. The number of SNPs per platform ranged from 10 204 (Affymetrix Mapping 10 K) to 934 946 SNPs (Affymetrix Genome Wide SNP 6). For all nine genotyping platforms (of seven levels of resolution), the model performed equally well in capturing the informative SNPs and predicting the population category. The assignment of 2504 samples from 1000 Genomes Project into 5 continental groups had low margins of error for all genotyping platforms (Fig. 2). To benchmark the separation of 26 groups, we applied a repeated cross validation (CV) on both random forest and multinomial linear regression model for the assignment task. Here, the CV accuracy rate ranged between 75–85% depending on platform resolution. When combining pairs of frequently mixed groups into 10 categories, a CV accuracy rate rose to >99% (Fig. S3 and Table S3).

Benchmarking normal genome profile assignment with HapMap data. To validate the accuracy of our tool to map population origins from non-cancer SNP datasets, we used 112 samples from the HapMap project²⁸ which is not included in the reference set, but acquired through GEO. After a further examination for family relatedness, we excluded 67 first-degree relatives to samples from 1000 Genomes project. By HapMap metadata, the 45 samples came from three distinct population groups (at the level of 26 population groups, from three different super population groups): CEU, CHB and YRI. Except for one sample with “CEU” label predicted as “AMR”, the assignment of the rest of samples all matched their ethnicity information indicated in the meta-data (Table 1).

Paired cancer-normal comparison. One of the highlights in this tool lies in the determination of population origin from cancer genome profiles carrying various acquired genomic aberrations. Since the non-cancer samples could be correctly assigned according to HapMap categories, we then validated the cancer genome based assignments in samples where normal genome profiles of the same patients (e.g. from peripheral blood or

	CEU	CHB	YRI
AFR	0	0	14
AMR	1	0	0
EAS	0	6	0
EUR	24	0	0

Table 1. Comparison of HapMap metadata and predicted population group. Each column indicates HapMap population labels. Each row indicates the label predicted by the tool. CEU stands for Utah residents with Northern and Western European ancestry from the CEPH collection, CHB stands for Han Chinese in Beijing, China. YRI stands for Yoruba in Ibadan, Nigeria; AFR stands for African; AMR stands for Admixed American; EAS stands for East Asian; EUR stands for European.

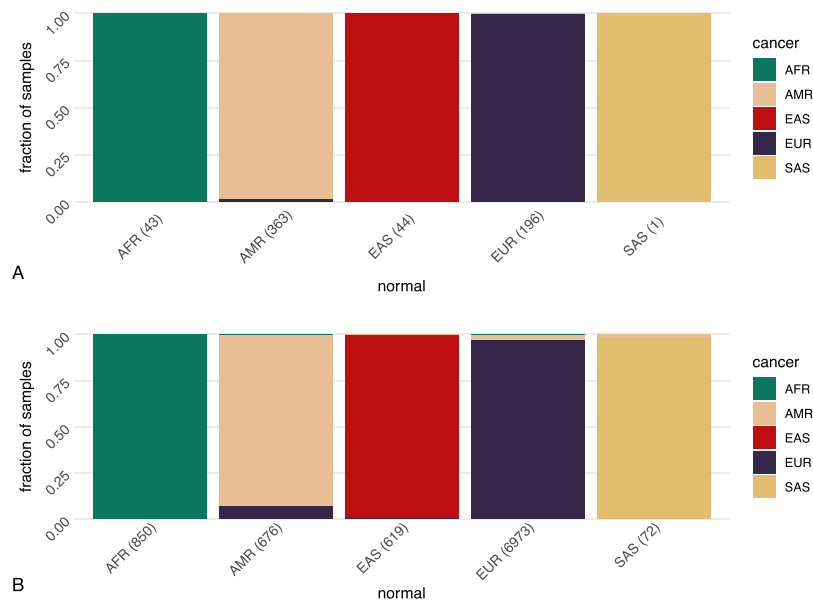


Figure 3. Accuracy of assignment with paired tumor and normal samples from GEO and TCGA project. 647 (score > 0.2) individuals from GEO (**A**) and 8216 (score > 0.2) individuals from TCGA project (**B**) with paired tumor and normal samples were examined. The columns are the prediction groups for the normal samples and the color codes indicate the proportion of each group predicted from their respective cancer samples.

non-cancer tissue samples) were available as reference. We performed the validation with two independent data sources — GEO and TCGA project.

GEO data. From the GEO repository, we retrieved all paired normal and cancer samples from 1145 individuals and compared the outcome of the population assignment. When including all 1145 individuals, 92.1% of the normal samples matched with paired tumor samples. With an increasing confidence score, the matching accuracy increased (Table S4). The error rate dropped from 15.9% to 1.5% comparing samples with score range “0.2–0.4” to those with “0–0.2”. After setting a threshold of normal samples with score > 0.2, 95.8% accuracy could be achieved for the remaining 688 individuals. When also setting the score threshold for cancer samples to > 0.2, 98.9% of the 647 remaining samples could be matched correctly (Fig. 3A). This comparison suggested a high accuracy rate in population assignment for cancer samples and an increase in the level of accuracy with a lower admixture background of the individual.

TCGA data. We used the genomic array profiling data from TCGA project, all of which originate from the array platform Genome Wide SNP 6. We extracted all 18380 samples (of 9190 individuals), where a normal tissue control for the respective tumor sample was available. 8924 out of 9190 (97.1%) individuals had matched tumor/normal categories. After setting a threshold of normal samples with score > 0.2, 98.4% accuracy (8413 out of 8549 individuals) can be achieved. When also setting the score threshold for cancer samples to > 0.2, 8195 of the 8216 (99.7%) remaining samples could be matched correctly (Fig. 3B). The score and error relation was depicted in the Table S5.

For population assignment into 26 groups, 8522 out of 9190 (92.7%) pairs had matching categories (Fig. S4). With the same score thresholding as described above, the accuracy increased to 97% (7374 out of 7602) and 99.1% (6948 out of 7010) respectively.

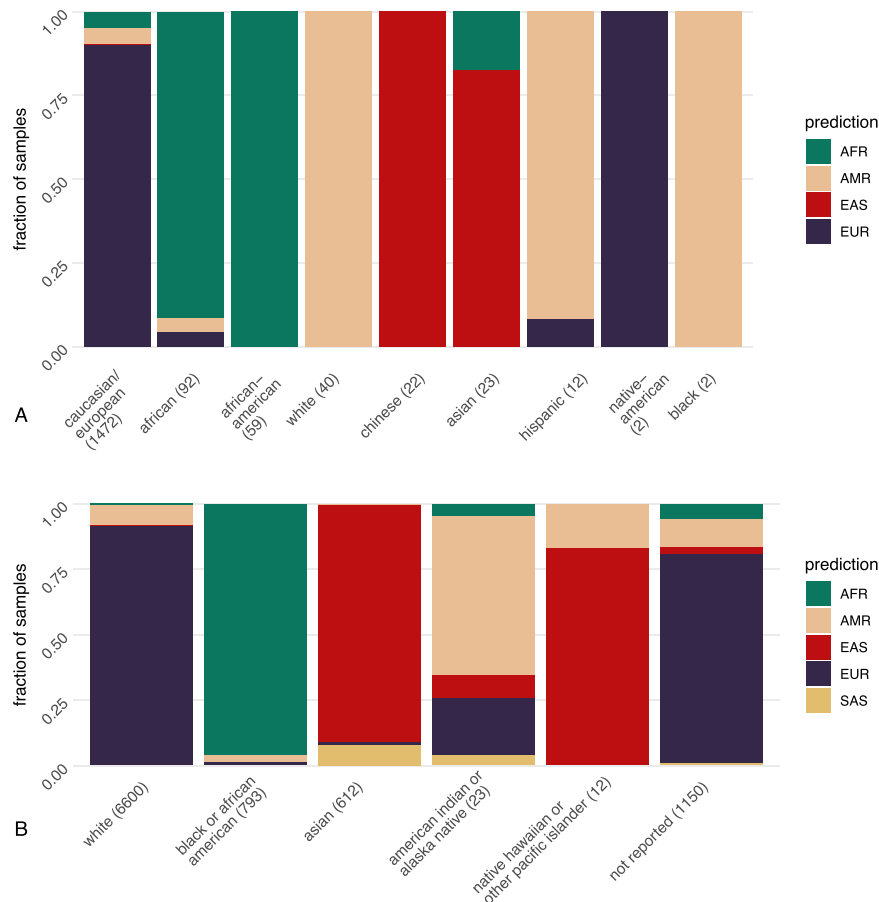


Figure 4. Accuracy of assignment with self-report metadata from GEO and TCGA project. **(A)** The composition of assigned super population groups of 1724 samples from GEO with extracted population/ethnicity metadata in nine groups: “african”(92), “african-american”(59), “black”(6), “caucasian/european”(1472), “white”(40), “asian”(23), “chinese”(12), “hispanic”(12) and “native american”(2). **(B)** The composition of assigned super population groups of 9190 samples in six “race” categories from TCGA metadata: “American Indian or Alaska native”(23), “Asian”(612), “Black or African American”(793), “Native Hawaiian or other Pacific Islander”(12), “White”(6600) or “Not Reported”(1150).

Modification of SNP status during carcinogenesis. We further scrutinized the source of noise and information loss during carcinogenesis. Using the same dataset, we performed a SNP-by-SNP comparison between the paired normal and cancer samples and summarized the SNPs which changed from heterozygous to homozygous (Het2Homo), from homozygous to heterozygous (Homo2Het) or stayed identical (Fig. S5). The Het2Homo (6.3% on average) transition occurred in case of allele loss or copy-neutral loss of heterozygosity and constituted the larger part to noise in cancer samples. The Homo2Het (5.6% on average) transition could come from loss of both alleles causing BAF to appear at 0.5 or less frequently when somatic point mutations coincide with germline polymorphism sites. In this section, We benchmarked the cancer/normal assignment consistency with presence of these two types of SNP status modifications. We also indirectly addressed the issue of tumor sample purity here, as mixed cancer/normal samples would be assigned to same categories as either of the pure cell populations.

Comparison with self-reported ethnicity metadata. To assess the overall accuracy of self-reported population information, we compared the meta-data from GEO and TCGA with our benchmarked results.

GEO data. We retrieved a total of 1724 samples with intelligible self-reported metadata from GEO. We extracted the population-implicating keywords, which formed nine groups: “african”(92), “african-american”(59), “black”(6), “caucasian/european”(1472), “white”(40), “asian”(23), “chinese”(12), “hispanic”(12) and “native american”(2) (Fig. 4). Specifically, “african” and “african-american” samples were mostly assigned to “AFR” group (91.3% and 93.6% respectively). The 1472 “caucasian/european”-labeled samples were assigned to “EUR” (90.2%) with small fraction assigned to “AFR”(4.6%) or “AMR”(5.0%). 11 out of 12 “hispanic” samples were assigned to “AMR” with one as “EUR”. All 40 “white” samples were assigned to “AMR”. These 40 samples derived from the same study, so the patients were likely from a similar ethnicity background and reported as “white” by the study. The “asian” samples were mostly assigned to “EAS” group (19/23). All 12 “chinese” samples were correctly assigned to “EAS”. Two ‘native-american’ samples were both assigned to “EUR”. Counting all samples, a meta-data matching rate of

88.3% was achieved. This indicated the existing heterogeneity in ethnicity description across studies and a certain degree of inaccuracy in self-reported ethnicity information.

TCGA data. From the test of paired cancer/normal samples in TCGA, we also compared our results of 9190 normal samples with the “race” attribute provided by the metadata. There, six categories were distinguished: “American Indian or Alaska native”(23), “Asian”(612), “Black or African American”(793), “Native Hawaiian or other Pacific Islander”(12), “White”(6600) or “Not Reported”(1150). The fraction of assignment results in each of these categories were shown in the bar plots in Fig. 4B). Most samples reported as “white” were assigned to “EUR”(91.4%) with 7.7% “AMR”; Most samples reported as “Black or African American” were assigned to “AFR”(95.6%) with “AMR”(3.0%) and EUR(1.0%); The “Asian”-labeled samples were composed of 90.5% “EAS” and 8.2% “SAS”. 14 out of the 23 “American Indian or Alaska Natives”-labeled samples were assigned to “AMR” and 5 to “EUR”. In the 12 “Native Hawaiian or other Pacific Islanders” samples, 10 were assigned to “EAS” and 2 to “AMR”. The samples were also assigned to one of the 26 population groups (Fig. S6), but there was no more meta-data for a comparison on this level. If we expect the following three “race” groups to match the assignment: “White” to “EUR”, “Black or African American” to “AFR” and “Asian” to “EAS” and “SAS”, then a total matching rate of 92.4% was achieved with 95.6% in African group, 98.6% in Asian group and 91.4% in European group. Despite the high matching percentage on the level of super population group in the self-report groups with large sample numbers (“White”, “Black or African American” and “Asian”), one may still argue that with the prediction outcome from our tool, the “race” information defined in the project could be well complemented and adapted for a quantitative measure for genetic background assessment.

Discussion

The presented method resolves two issues that have been missing in the earlier approaches to estimate population structure from genotyping data: 1. it allows the integration of platforms with heterogeneous, non-overlapping SNP positions for cross-platform meta-analysis; 2. it confirms the possibility of deriving the population origin using data from aberrant cancer genomes in addition to those of unaltered genomes.

In our study, we establish the robustness of deriving population despite platform heterogeneity by cross-platform benchmarking. We show that for all the genotyping platforms used in the study, our pipeline achieves similarly high prediction accuracy, even for the lowest resolution platform with around 10 000 SNPs. This helps to integrate genotyping data generated with different platforms for meta-analysis, where the heterogeneity in resolution does not pose great impact on the parameters of interest (e.g. cancer copy number variation studies).

We demonstrate the feasibility of incorporating aberrant cancer genomes in a step-wise manner. First, we benchmark the method using samples with unaltered genomes (non-cancer) which are different from the reference dataset, and show a high matching accuracy. In addition, we simulate different levels of LOH from unaltered genomes and observed mis-assignment or decrease of confidence score only at high coverage of LOH (LOH in >60% genome) (Fig. S7). Then, we employ paired normal/cancer samples and benchmark the assignment consistency between them. With two independent data resources (GEO and TCGA), we show that the population group estimated from unaltered genome of the same individual matches that from the aberrant cancer genome with a high accuracy rate.

From the paired sample comparison, we do observe a certain degree of mismatches. These often occur in samples with a highly admixed background (indicated by lower score), rather than related to whether they are normal or cancer samples as discussed above. Furthermore, the mismatch occurs mostly between the AMR and EUR labels, which is partly expected as the 1000 Genomes Project describes “AMR” group as “Admixed American”. Clearly defining indigenous American population as reference is impractical due to the admixture events in the last hundreds of years in South America. Such admixture events inevitably pose a challenge when establishing stable reference groups for assessing the impact of population background on secondary phenotypes like mutational patterns in cancer. As many human populations are still limitedly represented or missing and major contributions to human genetic variations are emerging²⁹, genomic population models will be amenable for re-assortment. Accordingly, our SNP2pop tool allows user-specified labeled reference data upload and assign population labels to the unknown samples according to the provided reference groups.

To conclude, we have developed a bioinformatics tool to assign population group based on SNP genotyping array data. We demonstrate its feasibility and accuracy on cancer samples, where somatic mutations may obfuscate part of the ancestry related SNP signal. This work can facilitate the re-analysis and meta-analysis of available cancer data by grouping samples with similar genetic background to study the potential genetic predisposition to cancer. In addition, our method provides the basis for subsequent haplotype phasing and refinement of genomic landscape for emerging somatic variation. With this tool, researchers will be able to integrate cancer genome profiling data from multiple resources to better assess the contribution of population background in population-specific mutation patterns occurring in cancer.

Methods

We used whole genome sequencing data from the 1000 Genomes Project²⁷ as reference data, which contains 2504 individuals from 26 population groups out of 5 continental groups. Allele information for SNP positions measured by the 9 selected array platforms was retrieved. In order to achieve between-study consistency for selection of informative SNPs, we used admixture model³⁰ to pre-compute parameters optimized for each genotyping platform based on the reference data. The allele frequency and ancestry fraction parameters were projected to the incoming cancer dataset of the same platform. Then, by applying a random forest classification, we assigned the population label to the highest voted group and produced a score for the difference between highest and second highest percentage votes. The overview of the pipeline is shown in Fig. 5. The tool is accessible for direct use as

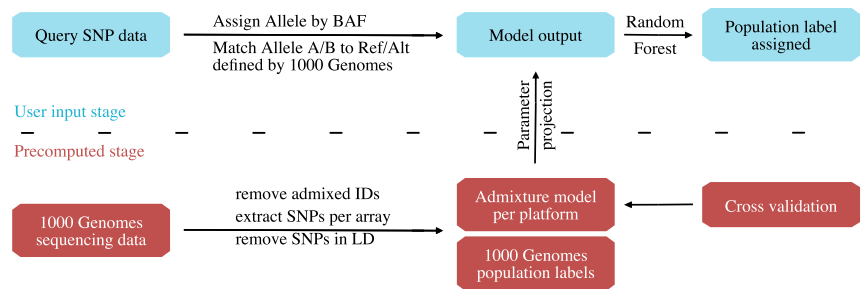


Figure 5. Pipeline of our tool to derive population assignment for individual cancer samples. The development workflow of the tool starts with retrieving the sequencing data from 1000 Genomes Project, removing IDs with admixture origin, extracting SNPs according to the genotyping array platforms and removing SNPs in linkage disequilibrium (LD). The generated PLINK files are stored as an intermediate reference file set. This file set is used further to train an admixture model with the model parameter matrix (allele frequency and ancestry fraction) stored for reference. Finally, a random forest classification is trained for the reference data based on their annotated population labels. These pre-computation steps are the basis of the tool, which saves time and preserves inter-study consistency. Now, when the new genotyping data comes in, the tool will assign the allele by B-allele frequency (BAF), check for the reference/alternative base for each SNP to match the reference data, assign the missing alleles (homozygous SNP from incoming data, thereby alternative base not available) based on reference SNP annotation. When the data is cleaned and compatible, the tool will use the pre-computed model parameters to analyze the admixture in the incoming dataset. Finally, with the model output, the population group of the incoming data is assigned. Users can choose to output the model parameters as well as the population groups.

a Docker image “<https://hub.docker.com/r/baudisgroup/snp2pop>” on DockerHub and its source codes can be found on Github in project “<https://github.com/baudisgroup/snp2pop>”.

Data preparation. *Reference data preparation.* The SNP positions for each platform were acquired from Affymetrix annotation files. The allele information was extracted for all positions with vcfTools³¹. The 12 mislabeled or admixed individuals were removed from the reference dataset, leaving 2492 individuals. The SNP positions with duplicated rsIDs in annotation files were removed. The reference or alternative strands were swapped to match the SNP array annotation. Sites with minor allele frequency (MAF) of less than 5 percent were removed. SNPs were subsequently pruned for independence using PLINK 1.9³². Specifically, a sliding window of 50 bp and a 5 bp shift of window at each step of pruning, with the variance inflation factor ($VIF = 1/(1 - R^2)$, with R^2 equal to the multiple correlation coefficient for a SNP regressed on all other SNPs) at 1.5. The result files were stored as PLINK output for each platform in .bed, .bim and .fam formats, of which the .bim files were used to extract SNP positions from target data.

Target data preparation. The SNP array data were processed with ACNE R package³³ to extract allele-specific copy numbers as B-allele frequencies (BAF). SNPs were labeled as homozygous A, heterozygous AB or homozygous B by the BAF value in ranges 0–0.15, 0.15–0.85 or 0.85–1, respectively, to allow both for noise and expected aneuploidy in the biosamples.

Data used for benchmarking are accessed through arrayMap database³⁴, using a collection of re-processed genotyping series from the Gene Expression Omnibus (GEO) repository³⁵, and the TCGA data repository³⁶ from <https://cancergenome.nih.gov>. The samples deposited in the former were submitted by individual research studies, whereas the latter coordinated cancer genome data generation and processing with centralized protocols.

Admixture model. While many approaches use principle component analysis (PCA) to select informative SNPs for population assignment, deriving these SNPs prior to clustering methods results in varying sets of SNPs between datasets. Here, we used the reference panel to train an admixture statistical model³⁰, to derive the allele frequency in theoretical ancestor groups for each SNP. The number of theoretical ancestor groups (K) was chosen to be 9 by cross-validation. The ancestry fraction plot for reference individuals demonstrates a proper information extraction to distinguish the five continental categories. By projecting the model parameters to a new sample with the corresponding platform, a robust and consistent output with 9 ancestry fractions was generated. The ratio of the 9 fractions in each individual in the reference data of 26 population groups for each platform is shown in Fig. S1.

Random forest label assignment. The nine ancestry fractions from the reference population were used to build a random forest model in each platform to predict the 26 population categories. The score was calculated as the difference in percentage votes between the best and the second best predicted labels. We performed repeated cross validation experiments (5-fold cross validation with 10 repetitions), and discovered twelve individuals assigned into a category different from the labels defined in their meta-data, which might occur due to a platform-specific detection of rare alleles that segregate in these samples. They were removed from reference list of the model. Finally, 2492 samples were used as training set in the classification model (removed IDs and information are found in Table S1). We compared random forest (RF) learning method with a classical multinomial

linear regression (MLR) model in terms of performance and computation time. The computation time is 3–5 fold slower than MLR (Table S2) but the performance in terms of accuracy is moderately higher (Fig. S2). The classification strategy used in the tool is first getting votes for all 26 population groups, and generating a highest vote group as the prediction result. Then, users have options to further generate a prediction out of 5 super-population groups or 10 broader population groups as described in Section “Cross-platform benchmarking”, combining votes of groups belonging to the same super-population or broader population group, and output the group with highest vote as the prediction outcome, and a score which allows for closer scrutiny for potential admixtures.

Data availability

The docker version of the tool is provided in DockerHub (hub.docker.com/r/baudisgroup/snp2pop). The code is available through the repository at github.com/baudisgroup/snp2pop.

Received: 9 July 2019; Accepted: 4 March 2020;

Published online: 16 March 2020

References

- Hanahan, D. & Weinberg, R. Hallmarks of cancer: the next generation. *Cell* **144**, 646–674 (2011).
- Alexandrov, L. B. *et al.* Signatures of mutational processes in human cancer. *Nature* **500**, 415 (2013).
- Lynch, H. & de la Chapelle, A. Hereditary colorectal cancer. *N Engl J Med* **348**, 919–932 (2003).
- Zhang, J., Nichols, K. & Downing, J. Germline mutations in predisposition genes in pediatric cancer. *N Engl J Med* **374**, 1391 (2016).
- Parkin, D. M., Pisani, P. & Ferlay, J. Global cancer statistics. *CA: a cancer journal for clinicians* **49**, 33–64 (1999).
- Danaei, G. *et al.* Causes of cancer in the world: comparative risk assessment of nine behavioural and environmental risk factors. *Lancet* **366**, 1784–1793 (2005).
- Siegel, R., Miller, K. & Jemal, A. Cancer statistics, 2017. *CA Cancer J Clin* **67**, 7–30 (2017).
- Siegel, R. L., Miller, K. D. & Jemal, A. Cancer statistics, 2018. *CA: a cancer journal for clinicians* **68**, 7–30 (2018).
- Amundadottir, L. T. *et al.* A common variant associated with prostate cancer in european and african populations. *Nat. genetics* **38**, 652 (2006).
- Stacey, S. N. *et al.* Common variants on chromosomes 2q35 and 16q12 confer susceptibility to estrogen receptor-positive breast cancer. *Nat. genetics* **39**, 865 (2007).
- Tenesa, A. *et al.* Genome-wide association scan identifies a colorectal cancer susceptibility locus on 11q23 and replicates risk loci at 8q24 and 18q21. *Nat. genetics* **40**, 631 (2008).
- Wu, C. *et al.* Genetic variants on chromosome 15q25 associated with lung cancer risk in chinese populations. *Cancer research* **69**, 5065–5072 (2009).
- Frank, S. A. Genetic predisposition to cancer - insights from population genetics. *Nat. Rev. Clin. Oncol.* **5**, 764 EP - Review Article (2004).
- Miki, Y. *et al.* A strong candidate for the breast and ovarian cancer susceptibility gene brca1. *Science* **266**, 66–71 (1994).
- Foulkes, W. D., Knoppers, B. M. & Turnbull, C. Population genetic testing for cancer susceptibility: founder mutations to genomes. *Nat. Rev. Clin. Oncol.* **13**, 41 EP - Review Article (2015).
- Li, H. *et al.* Refined geographic distribution of the oriental ald2* 504lys (nee 487lys) variant. *Annals human genetics* **73**, 335–345 (2009).
- Brooks, P. J., Enoch, M.-A., Goldman, D., Li, T.-K. & Yokoyama, A. The alcohol flushing response: an unrecognized risk factor for esophageal cancer from alcohol consumption. *PLoS medicine* **6**, e1000050 (2009).
- Keenan, T. *et al.* Comparison of the genomic landscape between primary breast cancer in african american versus white women and the association of racial differences with tumor recurrence. *J. Clin. Oncol.* **33**, 3621–3627, <https://doi.org/10.1200/JCO.2015.62.2126> PMID: 26371147 (2015).
- Deng, J. *et al.* Comparative genomic analysis of esophageal squamous cell carcinoma between asian and caucasian patient populations. *Nat. Commun.* **8**, 1533, <https://doi.org/10.1038/s41467-017-01730-x> (2017).
- Zhang, W., Edwards, A., Flemington, E. K. & Zhang, K. Racial disparities in patient survival and tumor mutation burden, and the association between tumor mutation burden and cancer incidence rate. *Sci. Reports* **7**, 13639, <https://doi.org/10.1038/s41598-017-13091-y> (2017).
- Pharoah, P. D. P. *et al.* Polygenic susceptibility to breast cancer and implications for prevention. *Nat. Genet.* **31**, 33 EP - Article (2002).
- Albrechtsen, A. *et al.* Relatedness mapping and tracts of relatedness for genome-wide data in the presence of linkage disequilibrium. *Genetic epidemiology* **33**, 266–274 (2009).
- U.S. Census Bureau. Disparities in stem employment by sex, race, and hispanic origin, <https://www.census.gov/library/publications/2013/acs/acs-24.html> (2013).
- U.S. Census Bureau. Income, poverty, and health insurance coverage in the united states: 2012, <https://www.census.gov/library/publications/2013/demo/p60-245.html> (2013).
- Cai, H. *et al.* Progenetix: 12 years of oncogenomic data curation. *Nucleic acids research* **42**, D1055–D1062 (2014).
- Turakulov, R. & Easteal, S. Number of snps loci needed to detect population structure. *Hum. heredity* **55**, 37–45 (2003).
- Auton, A. *et al.* A global reference for human genetic variation. *Nature* **526**, 68–74 (2015).
- International HapMap Consortium. The international hapmap project. *Nature* **426**, 789–796 (2003).
- Sherman, R. M. *et al.* Assembly of a pan-genome from deep sequencing of 910 humans of african descent. *Nat. Genet.* <https://doi.org/10.1038/s41588-018-0273-y> (2018).
- Alexander, D., Novembre, J. & Lange, K. Fast model-based estimation of ancestry in unrelated individual. *Genome Res* **19**, 1655–1664 (2009).
- Danecek, P. *et al.* The variant call format and vcf tools. *Bioinformatics* **27**, 2156–2158 (2011).
- Purcell, S. *et al.* Plink: a tool set for whole-genome association and population-based linkage analyses. *The Am. J. Hum. Genet.* **81**, 559–575 (2007).
- Ortiz-Estevéz, M., Bengtsson, H. & Rubio, A. Acne: a summarization method to estimate allele-specific copy numbers for affymetrix snp arrays. *Bioinformatics* **26**, 1827–1833 (2010).
- Cai, H., Kumar, N. & Baudis, M. arraymap: a reference resource for genomic copy number imbalances in human malignancies. *PLoS One* **7**, e36944 (2012).
- Edgar, R., Domrachev, M. & Lash, A. Gene expression omnibus: Ncbi gene expression and hybridization array data repository. *Nucleic Acids Res* **30**, 207–210 (2002).
- Cancer Genome Atlas Research Network. Comprehensive genomic characterization defines human glioblastoma genes and core pathways. *Nature* **455**, 1061–1068 (2008).

Acknowledgements

We thank Paula Carrio Cordo for the metadata curation, Bo Gao and Damian Szklarczyk for helpful discussions. We thank the Zurich Bioinformatics group and ELIXIR Population Genomics for valuable comments. Q.H. has been supported by University of Zurich with “Forschungskredit CanDoc” fellowship.

Author contributions

Q.H. developed the tool and performed benchmarking. M.B. conceived the project. Both authors contributed to data assembly and writing and editing of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-020-61854-x>.

Correspondence and requests for materials should be addressed to M.B.

Reprints and permissions information is available at www.nature.com/reprints.

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



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, 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 license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license 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 license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020