Faster model-based estimation of ancestry proportions

Cindy G. Santander[1], Alba Refoyo Martinez[2] and Jonas Meisner[3,4,*]

[1]Department of Biology, University of Copenhagen, Denmark
[2]Center for Health Data Science, University of Copenhagen, Denmark
[3]Mental Health Centre Copenhagen, Copenhagen University Hospital, Denmark
[4]Novo Nordisk Foundation Center for Basic Metabolic Research, University of Copenhagen, Denmark

*Corresponding author

July 8, 2024

Abstract

Ancestry estimation from genotype data in unrelated individuals has become an essential tool in population and medical genetics to understand demographic population histories and to model or correct for population structure. The `ADMIXTURE` software is a widely used model-based approach to account for population stratification, however, it struggles with convergence issues and does not scale to modern human datasets or to the large number of variants in whole-genome sequencing data. Likelihood-free approaches optimize a least square objective and have gained popularity in recent years due to their scalability. However, this comes at the cost of accuracy in the ancestry estimates. We present a new model-based approach, `fastmixture`, which adopts aspects from likelihood-free approaches for parameter initialization, followed by a mini-batch expectation-maximization procedure to model the standard likelihood. We demonstrate in a simulation study that the model-based approaches of `fastmixture` and `ADMIXTURE` are significantly more accurate than recent and likelihood-free approaches. We further show that `fastmixture` runs approximately 20 times faster than `ADMIXTURE` on both simulated and empirical data from the 1000 Genomes Project such that our model-based approach scales to much larger sample sizes than previously possible. Our software is freely available at `https://github.com/Rosemeis/fastmixture`.

Introduction

For the past two decades, unsupervised ancestry estimation has been a crucial element in studies on human evolutionary genetics and genome-wide association studies [1, 2]. Estimating global ancestry proportions in unrelated individuals and corresponding ancestral allele frequencies has been a classic way of correcting for population structure and understanding demographic processes that have shaped the evolutionary history of modern populations. Ancestry estimation from genotype data emerged with the model-based clustering approach proposed with `STRUCTURE` [3], where individuals are proportionally assigned to an assumed number of ancestral populations. They modeled the probability of the observed genotype data given ancestry proportions and ancestral allele frequencies using a Bayesian approach. Due to scalability issues, the Bayesian model was later replaced by maximum likelihood models [4, 5].

In the era of big data, where ever-growing cohorts contain thousands of individuals with genotype data for millions of genetic variants, classic state-of-the-art model-based approaches, such as `ADMIXTURE` and `STRUCTURE`, fail to scale due to computational intractability. Over the past decade, there have been multiple attempts to scale unsupervised ancestry estimation. These efforts have primarily been rooted in likelihood-free approaches, with a few exceptions that have attempted to scale the standard model-based approach. Model-based and likelihood-free approaches have been shown to be connected within the framework of matrix factorization, where both approximate an observed genotype matrix by lower rank matrices under different assumptions and constraints [6]. In likelihood-free approaches, it is common to optimize a different least square objective using either non-negative matrix factorization (NMF) [7, 8] or alternating least square (ALS) [9, 10]. Meanwhile, for the model-based approaches, a stochastic variational inference algorithm [11] and a neural network autoencoder [12] have been proposed in recent times.

As large-scale cohorts begin to expand and include more cosmopolitan representations of individuals around the globe, obtaining accurate ancestry estimates has become a crucial task of modern genomics. This is essential to properly correct for population structure in genome-wide association studies and mitigate the bias in the over-representation of European-descent samples in current cohorts due to the ever-increasing focus on genomics for population health applications [13, 14]. While convergence within a reasonable time frame is highly desirable for these extensive datasets, recent attempts to improve scalability most often compromise accuracy in the ancestry estimation, as demonstrated here. The field still lacks a method that scales to larger datasets while being as accurate as the classic state-of-the-art approaches.

We introduce `fastmixture`, a novel model-based method for estimating ancestry proportions and ancestral allele frequencies in unrelated individuals. We leverage randomized singular value decomposition (SVD) for initializing the ancestry proportions and ancestral allele frequencies, followed by a mini-batch accelerated scheme to speed up the convergence of the expectation-maximization (EM) algorithm. We demonstrate in an extensive simulation study and on real data that `fastmixture` significantly outperforms the original `ADMIXTURE` software in terms of speed while maintaining higher accuracy compared to recently developed approaches for ancestry estimation.

Material and methods

We define a diallelic genotype matrix of N individuals and M variants or single-nucleotide polymorphisms (SNPs) as $\mathbf{G} \in \{0, 1, 2\}^{N \times M}$, which corresponds to the minor allele counts. We describe the estimation of ancestry proportions and ancestral allele frequencies with K ancestral sources

as a low-rank matrix factorization problem such that $\mathbf{G} \approx 2\mathbf{Q}\mathbf{P}^T$, with $\mathbf{Q} \in [0,1]^{N\times K}$ and the constraint $\sum_{k=1}^{K} q_{ik} = 1$, for $i = 1, \ldots, N$, and $\mathbf{P} \in [0,1]^{M\times K}$. The matrix factorization problem can also be interpreted as the estimation of individual allele frequencies from the genotype matrix, $\mathbf{H} = \mathbf{Q}\mathbf{P}^T$. Here, $h_{ij} = \sum_{k}^{K} q_{ik}p_{jk}$ is the individual allele frequency of individual i at variant j assuming K ancestral sources. The individual allele frequency will correspond to the underlying parameter in a binomial sampling process of a genotype conditioned on population structure.

Likelihood model

We estimate ancestry proportions and ancestral allele frequencies by maximizing the likelihood model introduced in `ADMIXTURE` [5]. The log-likelihood model is defined as follows, assuming independence for individuals and variants:

$$L(\mathbf{Q}, \mathbf{P}) = \sum_{i}^{N}\sum_{j}^{M} g_{ij}\log(h_{ij}) + (2 - g_{ij})\log(1 - h_{ij})\,. \tag{1}$$

We utilize the expectation-maximization (EM) algorithm of `frappe` [4] and `ADMIXTURE` [5] to maximize the log-likelihood, where the EM updates at iteration t for entries of $\mathbf{Q}$ and $\mathbf{P}$ are defined as follows, respectively:

$$q_{ik}^{(t+1)} = \frac{q_{ik}^{(t)}}{2M}\sum_{j=1}^{M}\left(\frac{g_{ij}p_{jk}^{(t)}}{h_{ij}^{(t)}} + \frac{(2 - g_{ij})(1 - p_{jk}^{(t)})}{1 - h_{ij}^{(t)}}\right)\,, \tag{2}$$

$$p_{jk}^{(t+1)} = \left(p_{jk}^{(t)}\sum_{i=1}^{N}\frac{g_{ij}q_{ik}^{(t)}}{h_{ij}^{(t)}}\right)\Bigg/\left(\sum_{i=1}^{N}\frac{g_{ij}q_{ik}^{(t)}p_{jk}^{(t)}}{h_{ij}^{(t)}} + \frac{(2 - g_{ij})q_{ik}^{(t)}(1 - p_{jk}^{(t)})}{1 - h_{ij}^{(t)}}\right)\,. \tag{3}$$

The EM algorithm is notorious for its slow convergence rate and to expedite this process, we employ the `SQUAREM` SqS3 acceleration scheme [15, 16]. The SqS3 acceleration scheme combines multiple EM updates into a larger jump in parameter space at the expense of an increased computational cost per iteration. The convergence criteria of the EM algorithm is defined by the difference in log-likelihood between two successive iterations:

$$L(\mathbf{Q}^{(t+1)}, \mathbf{P}^{(t+1)}) - L(\mathbf{Q}^{(t)}, \mathbf{P}^{(t)}) < \epsilon\,, \tag{4}$$

with ϵ being a user-defined threshold.

SVD initialization

Multiple approaches for inferring population structure are connected under the umbrella of matrix factorization, each incorporating different conditions and constraints [6]. Therefore, we leverage the efficiency of likelihood-free approaches to provide a better initialization of $\mathbf{Q}$ and $\mathbf{P}$ to aid the convergence rate of the EM algorithm in comparison to a standard random initialization.

We initialize $\mathbf{Q}$ and $\mathbf{P}$ using individual allele frequencies estimated from randomized SVD performed on the genotype matrix, combined with an alternating least squares (ALS) approach. The randomized SVD is defined as $\mathbf{G}_C \approx \mathbf{U}_{[1:D]}\mathbf{S}_{[1:D]}\mathbf{V}_{[1:D]}^T$, which extracts the top D singular values and singular vectors, with $\mathbf{G}_C$ being the centered genotype matrix. The individual allele frequency

of individual i at variant j is then approximated using the SVD as $\hat{h}_{ij} = f_j + \frac{1}{2}\sum_{d=1}^{D} u_{id}s_d v_{jd}$, where f_j is the minor allele frequency of variant j. We use $D = K - 1$, as the top $K - 1$ singular vectors will capture the population structure of K distinct populations [17]. By initializing $\mathbf{P}$ randomly, we can then factorize the individual allele frequency matrix, $\hat{\mathbf{H}}$, using an ALS approach to iteratively estimate both $\mathbf{Q}$ and $\mathbf{P}$, which minimizes the following least square objective:

$$\min_{\mathbf{Q},\mathbf{P}} \|\hat{\mathbf{H}} - \mathbf{Q}\mathbf{P}^T\|_F^2 \,, \tag{5}$$

where $\|.\|_F$ is the Frobenius norm. The concept is similar to `ALStructure` [9] and `SCOPE` [10], which instead rely on a latent subspace estimation step for approximating the individual allele frequencies. The convergence criteria of the ALS approach is defined by the root mean square error (RMSE) of $\mathbf{Q}$ matrices between two successive iterations:

$$\sqrt{\frac{1}{NK}\sum_{i=1}^{N}\sum_{k=1}^{K}\left(q_{ik}^{(t+1)} - q_{ik}^{(t)}\right)^2} < \delta \tag{6}$$

with δ being a user-defined threshold. The initialization of $\mathbf{Q}$ and $\mathbf{P}$ is fully described in the algorithm 1.

Algorithm 1 `fastmixture` initialization

Given: $\mathbf{G}$, $\mathbf{f}$, D, K
$\mathbf{G}_C \leftarrow \text{center}(\mathbf{G}, \mathbf{f})$ ▷ Centered genotype matrix
$\mathbf{U}_{[1:D]}$, $\mathbf{S}_{[1:D]}$, $\mathbf{V}_{[1:D]} \leftarrow \text{randomizedSVD}(\mathbf{G}_C, D)$ ▷ Randomized SVD
$\hat{\mathbf{H}} = \frac{1}{2}\mathbf{U}_{[1:D]}\mathbf{S}_{[1:D]}\mathbf{V}_{[1:D]}^T + \mathbf{f}$ ▷ Individual allele frequencies based on SVD
$\mathbf{P} \in [0,1]^{M\times K}$ initialized randomly
while not converged **do** ▷ Alternating least squares for solving equation 5
    $\mathbf{Q} \leftarrow \hat{\mathbf{H}}\mathbf{P}(\mathbf{P}^T\mathbf{P})^{-1}$
    $\mathbf{P} \leftarrow \hat{\mathbf{H}}^T\mathbf{Q}(\mathbf{Q}^T\mathbf{Q})^{-1}$
    if $\text{rmse}(\mathbf{Q}, \mathbf{Q}_{prev}) < \delta$ **then** ▷ Convergence criteria for ALS
        break
    end if
    $\mathbf{Q}_{prev} \leftarrow \mathbf{Q}$
end while
Return: $\mathbf{Q}$, $\mathbf{P}$

Mini-batch optimization

To further accelerate the convergence rate of the EM algorithm, we introduce simple mini-batch updates inspired by stochastic gradient descent and stochastic EM algorithms [18]. In each iteration, we randomly split the M variants into B batches, and perform accelerated SqS3 EM updates sequentially in each of the B batches. The strategy of sub-sampling variants was also explored in `TeraStructure` [11]. This results in the entries of $\mathbf{Q}$ being updated B times across the batches, while the entries of $\mathbf{P}$ are still only updated once per cycle. Following the mini-batch updates, full SqS3 EM updates are applied to stabilize the parameters in every iteration. Given the stochastic nature of our mini-batch updates, we halve the number of batches B every time the log-likelihood estimate

(Equation 1) fluctuates between iterations. The algorithm will therefore gradually converge towards a standard SqS3 algorithm for $B \to 1$. Our proposed mini-batch setting resembles a mini-batch gradient descent approach, where the batch size is increased over time. For a detailed description of the `fastmixture` method, please refer to Algorithm 2.

Algorithm 2 `fastmixture` estimation

Given: $\mathbf{G}$, $\mathbf{f}$, K, B

$\mathbf{Q}$, $\mathbf{P} \leftarrow$ initialization($\mathbf{G}$, $\mathbf{f}$, K) ▷ Algorithm 1: SVD initialization
$L_{prev} \leftarrow L(\mathbf{Q}, \mathbf{P})$ ▷ Log-likelihood (equation 1)
while not converged **do**
  Randomly split M variants into B batches
  for $b = 1, \ldots, B$ **do**
    $\mathbf{Q}, \mathbf{P}_{[b]} \leftarrow \text{SqS3}(\mathbf{G}_{[b]}, \mathbf{Q}, \mathbf{P}_{[b]})$ ▷ Mini-batch SqS3 EM updates
  end for
  $\mathbf{Q}, \mathbf{P} \leftarrow \text{em}(\mathbf{G}, \mathbf{Q}, \mathbf{P})$ ▷ Standard EM updates (equation 2 and 3)
  $\mathbf{Q}, \mathbf{P} \leftarrow \text{SqS3}(\mathbf{G}, \mathbf{Q}, \mathbf{P})$ ▷ Full SqS3 EM updates
  if $L(\mathbf{Q}, \mathbf{P}) < L_{prev}$ **then**
    $B = B/2$ ▷ Halve the number of batches
  end if
  if $B = 1$ and $L(\mathbf{Q}, \mathbf{P}) - L_{prev} < \epsilon$ **then** ▷ Convergence criteria for EM
    break
  end if
  $L_{prev} \leftarrow L(\mathbf{Q}, \mathbf{P})$
end while

Return: $\mathbf{Q}$, $\mathbf{P}$

Implementation details

The `fastmixture` software is implemented as a multithreaded command-line tool written in Python and Cython, which takes binary `PLINK` [19] files as input. It utilizes the NumPy library [20] for efficient array manipulation. We assume that the user has performed standard preprocessing, (e.g., variant filtering based on a minor allele frequency threshold). The genotype matrix is stored in an 8-bit integer format. For algorithm 1, we read the centered genotypes in chunks to reduce the memory consumption and perform randomized SVD as introduced in `PCAone` [21]. This approach minimizes the memory footprint of `fastmixture`, which is primarily dominated by the genotype matrices of $MN + 8CN$ bytes, with C being the chunk size of variants used in the randomized SVD. This allows our method to handle large-scale datasets effectively. The individual allele frequency matrix obtained in the randomized SVD, $\hat{\mathbf{H}}$, is computed implicitly through the singular matrices in the ALS approach to further reduce the memory footprint. Throughout the study, we use a starting point of $B = 32$ mini-batches to speed up the convergence of the EM algorithm, which works well across all the tested scenarios. We set the convergence criteria of the ALS approach and the EM algorithm to $\delta = 1.0 \times 10^{-4}$ and $\epsilon = 0.5$, respectively. Given the computational expense of the log-likelihood estimation step (equation 1), which requires a full pass over the data, we only evaluate it every fifth iteration along with convergence checks to reduce the number of computations. If the entries of $\mathbf{Q}$ and $\mathbf{P}$ are out of domain during the ALS or accelerated EM updates, we simply map them back to their domain through truncation and projection procedures.

Simulations

We simulate genotypes using the `msprime` [22] backwards-in-time coalescent model and infer true ancestral tracts using `tspop` [23]. To evaluate the capabilities of `fastmixture`, we assume different demographic models in three different scenarios, all featuring an admixture event where the true ancestry proportions are known. In each scenario, we simulate a genetic segment of 250Mb using a constant recombination rate of 1.28×10^{-8} and a mutation rate of 2.36×10^{-8} for 1,000 individuals. We visualize the different demographic models in Figure 1A, S1A, and S2A. A census event precedes the admixture event in all simulations to track the true ancestry of the inherited segments in the sampled individuals. The length of the segments each individual has inherited from each source population are then aggregated to estimate the ground truth ancestry proportions using `tspop`. Scenarios A and B are constructed from simple demographic models, where we assume a constant population size of 10,000 for all simulated populations, while Scenario C extends the out-of-Africa model [24] with an additional admixture event (American-Admixture) [25]. We perform standard filtering on minor allele frequencies at a threshold of 0.05, resulting in datasets consisting of 689,563 SNPs, 687,107 SNPs, and 500,114 SNPs for Scenario A, B, and C, respectively. A quick overview of the simulated datasets is provided in the supplementary material.

1000 Genomes Project

We also evaluate our `fastmixture` software in the phase 3 release of the 1000 Genomes Project (1KGP) [26]. The dataset consists of genotype data of 2,504 individuals from 26 populations across the world, assigned to five super-populations: AFR (African ancestry), AMR (American ancestry), EAS (East Asian ancestry), EUR (European ancestry), SAS (South Asian ancestry). We keep diallelic SNPs with a minor allele frequency greater than the standard threshold of 0.05, resulting in a total of 6,864,700 SNPs. Due to computational complexity and runtime considerations for comparative analyses, we also construct a downsampled dataset, which we refer to as "1KGP Down". This dataset is obtained by random downsampling or thinning the SNP set of the full dataset by a factor of ten, such that the downsampled dataset consists of 686,470 SNPs.

Results

We evaluate and compare the performance of `fastmixture` against widely used software for estimating ancestry proportions: `ADMIXTURE` (v1.3.0) [5], `Neural ADMIXTURE` (v1.4.1) [12] and `SCOPE` [10]. All software are executed with their default parameters. `ADMIXTURE` and `Neural ADMIXTURE` are model-based approaches, like ours. `ADMIXTURE` maximizes the likelihood model using a block relaxation method combined with a quasi-Newton acceleration scheme, while `Neural ADMIXTURE` employs a neural network autoencoder approach. On the other hand, `SCOPE` is a likelihood-free approach that estimates the ancestry proportions and ancestral allele frequencies using an alternating least square (ALS) method, differing from the maximum likelihood approaches of the other tools.

In our simulation study, we assess these tools based on the root mean square error (RMSE) (Table 1) and the Jenson-Shannon divergence (JSD) (Table S2) between the estimated ancestry proportions and the ground truth ancestry proportions. We also report the log-likelihood estimates (equation 1) (Table S3) for both the simulated and the empirical datasets of the 1000 Genomes Project. The different measures of assessment in the simulation study are further detailed in the supplementary material.

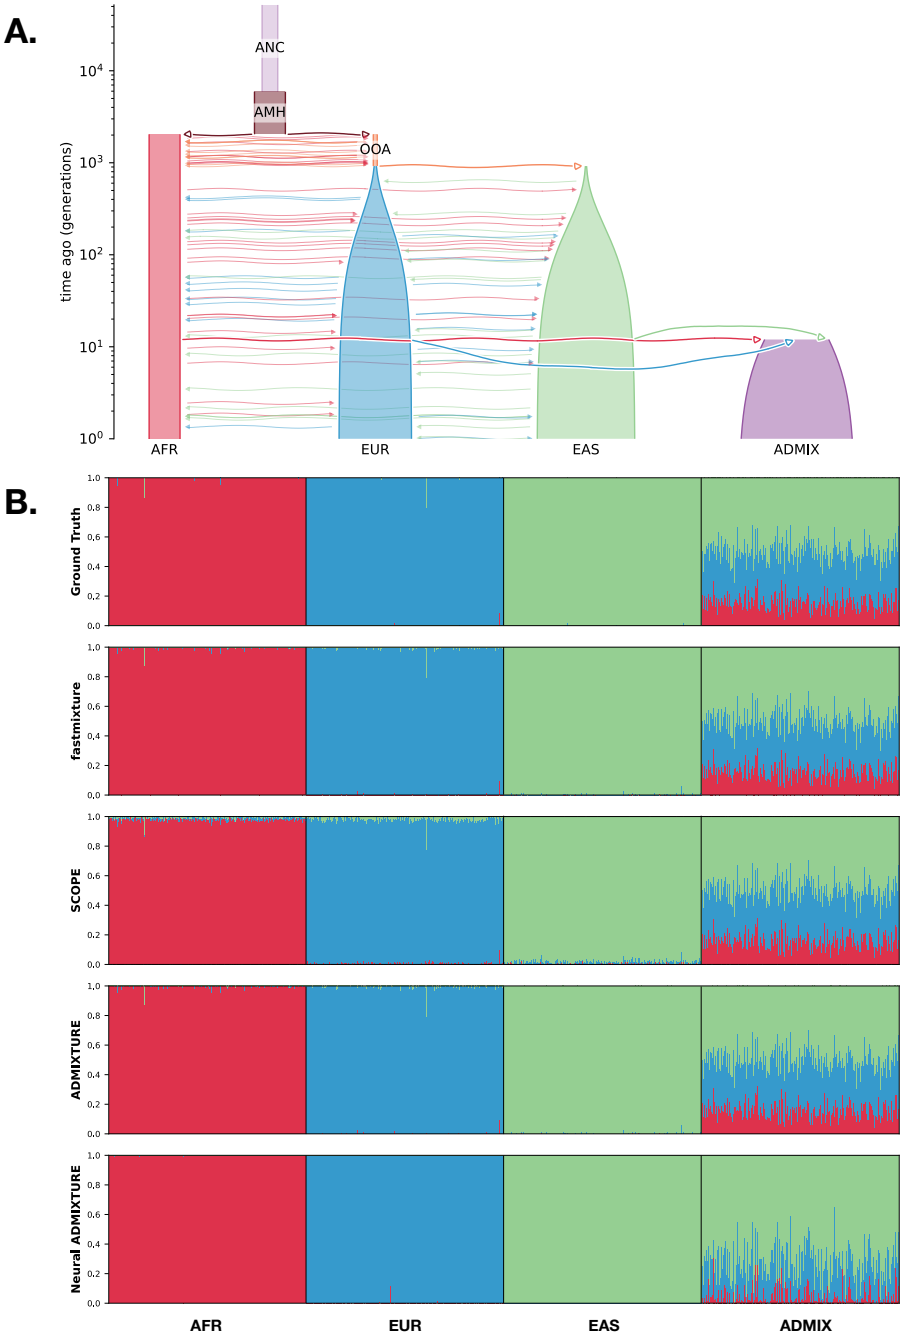

Figure 1: A. Demographic model of *Scenario C* (American admixture [25]) with 250 individuals sampled from each of the four populations. **B.** Admixture plots of the ancestry proportions in the simulated individuals for $K = 3$ with the ground truth at the top followed by the different software using their run with the highest log-likelihood.

Scenario	fastmixture	Neural ADMIXTURE	ADMIXTURE	SCOPE
A	0.0369 (4.4e-5)	0.1289 (0.0200)	0.0367 (3.1e-7)	0.0677 (7.8e-5)
B	0.0279 (6.2e-5)	0.3690 (0.0022)	0.0279 (2.2e-7)	0.0661 (8.2e-5)
C	0.0085 (6.6e-7)	0.1010 (0.0372)	0.0085 (4.5e-8)	0.0186 (3.3e-6)

Table 1: Root mean square error (RMSE) measures for estimated ancestry proportions in the three different simulation scenarios for the evaluated methods against the ground truth. The mean across five different runs is reported with the corresponding standard deviation in parenthesis.

Additionally, we compare the computational runtimes of `fastmixture` to the three other software (Figure 3 and Table S1). Notably, `ADMIXTURE` exhibits significant scalability issues, being approximately 20 times slower than the other approaches across all evaluated datasets. For instance, `ADMIXTURE` fails to complete a run within 24 hours for $K = 5$ on the full 1KGP dataset. In contrast, `fastmixture` has comparable runtimes to the two other faster approaches in the simulation study. `SCOPE` is the fastest of all evaluated approaches, showcasing the appealing choice of optimizing the least square objective, a step also utilized for parameter initialization in `fastmixture`.

Ancestry estimation in simulation studies

Under the simple demographic model of Scenario A (Figure S1), `ADMIXTURE` and `fastmixture` were the two most accurate software followed by `SCOPE`. When only considering the scalable methods, `fastmixture` outperformed `SCOPE` in terms of RMSE (Table 1) and JSD (Table S2), where `SCOPE` appeared to produce highly noisy ancestry proportions in individuals of the unadmixed populations. On all accounts, `Neural ADMIXTURE` severely underperformed. Notably when admixture was introduced from more than 3 sources, as in Scenario B (Figure S2), `Neural ADMIXTURE` failed to detect POP4 as a separate population from POP3 and modeled the ADMIX population incorrectly as a separate unadmixed population. For Scenario B, both `ADMIXTURE` and `fastmixture` were again more accurate than `SCOPE`, while `fastmixture` was ∼26x faster than `ADMIXTURE` (Figure 3 and Table S1).

Using the American-Admixture demographic model in Scenario C (Figure 1), we consistently observed that `ADMIXTURE` and `fastmixture` perform similarly in accuracy outputting results closest to the ground truth (Table 1 and Table S2) with `fastmixture` being ∼17x faster than `ADMIXTURE`. While across all scenarios `SCOPE` exhibited the fastest runtimes, it also consistently underperformed in accuracy as measured with RMSE and JSD, however, markedly better than `Neural ADMIXTURE`.

Ancestry estimation in the 1000 Genomes Project

In addition to our simulation study, we also applied `fastmixture` to empirical data of the 1000 Genomes Project (1KGP): both a downsampled version and the full set of variants. Specifically, the downsampled version was the only way to assess the performance of `ADMIXTURE` within a reasonable time frame in regards to convergence (Figure 2). Here we observed, similarly in the simulation study, that `fastmixture` performed comparably to `ADMIXTURE` and the two methods achieved the highest log-likelihoods, followed by `SCOPE` and then `Neural ADMIXTURE` (Table S3), even though `Neural ADMIXTURE` maximizes the log-likelihood in its approach. For the full 1KGP dataset, we only compared `fastmixture` to the other two scalable approaches, `SCOPE` and `Neural ADMIXTURE` (Figure S3). We demonstrated again that `fastmixture` was the best performing methods in terms of

achieving the highest log-likelihood. Strikingly, `Neural ADMIXTURE` failed to accurately distinguish ancestral contributions in the AMR and SAS super-populations in both the full dataset and in 1KGP Down. Moreover, we saw a correlation of $r^2 = 0.72$ between the estimated ancestry proportions in the downsampled and the full dataset for `Neural ADMIXTURE`, indicating poor robustness. This is in strong contrast to both `fastmixture` and `SCOPE` where the results between the full dataset and the downsampled one were highly correlated ($r^2 \approx 1$). We observed the same similar pattern of increased noise in the estimated ancestry proportions of the unadmixed individuals from `SCOPE`, which was consistent across the two 1KGP datasets.

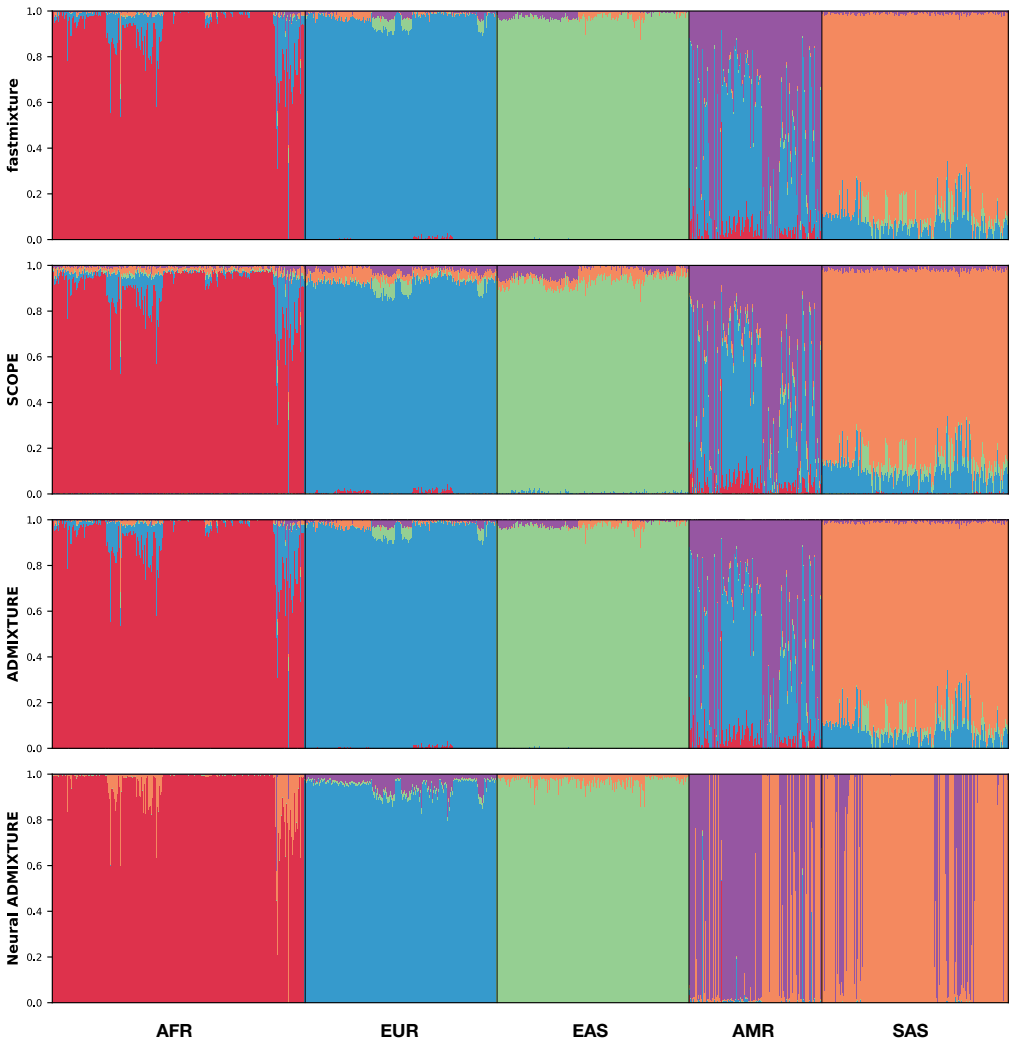

Figure 2: Admixture plots of the estimated ancestry proportions for $K = 5$ in the downsampled version of the 1000 Genomes Project dataset (1KGP Down) by the different software using their run with the highest log-likelihood.

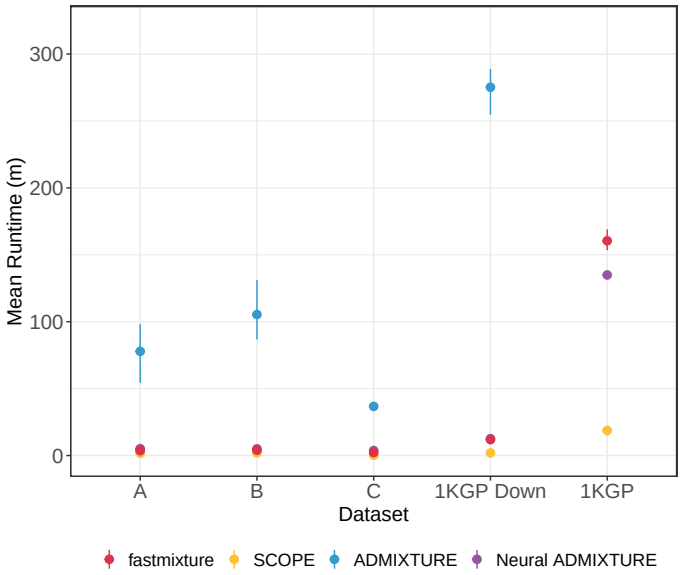

Figure 3: Runtime (in minutes) for the different software measured across five runs in the three simulated scenarios (A, B, and C), and in the two versions of the 1000 Genomes Project dataset: one including all individuals (1KGP), and the downsampled version (1KGP Down). The intervals display the full range of runtimes across the five runs. All analyses have been performed on a cluster with an Intel(R) Xeon(R) Gold 6152 CPU using 64 threads.

Discussion

We have presented, `fastmixture`, our novel method and software for ancestry estimation in genotype data using our extended model-based approach. We demonstrate that our approach performs comparably to `ADMIXTURE` while being ∼20x faster on average across all evaluated datasets. Among the four methods assessed, `ADMIXTURE` and `fastmixture` stand out as the most accurate approaches as shown in our simulation study. It is well known that `ADMIXTURE` has scalability issues for large sample sizes and whole-genome sequencing data, which we also demonstrate using the 1KGP datasets. We have not been able to run `ADMIXTURE` on the full 1KGP SNP set within a reasonable time frame ($<$ 24 hours) and have therefore only run it on the downsampled version with 10x less SNPs, a common procedure in population genetic studies. `fastmixture` has runtimes comparable to the recently introduced `Neural ADMIXTURE` software, which employs an autoencoder framework for speeding up the ancestry estimation. Note that the testing of the software has exclusively been conducted in a CPU-based setup. However, `Neural ADMIXTURE` struggles across all evaluated datasets and consistently performs the poorest among the evaluated methods, as it can only manage to model unadmixed populations. This is with an exception in Scenario B, where it even models the unadmixed POP3 and POP4 populations incorrectly as one. `SCOPE` emerges as the fastest approach across all evaluated datasets, demonstrating excellent scalability. However, its optimization of a simpler least squares objective compromises its ability to accurately estimate ancestry proportions, where it can be difficult to distinguish real admixture signals from noise.

When optimizing a different objective, it is expected that the log-likelihood estimates for `SCOPE`

would be lower compared to model-based approaches. However, based on the RMSE and JSD measures against the ground truth in the simulation study, our results showcase that the least square objective used in likelihood-free approaches, such as `SCOPE`, is not an optimal choice in comparison to the likelihood model. Our findings suggest that the added noise in the ancestry proportions estimated in `SCOPE` will only increase for scenarios with a larger K and for more complex demographic models. Furthermore, we observed a general trend of major optimization issues for the `Neural ADMIXTURE` software across all evaluated scenarios. A critical issue appears to be in their convergence evaluation, where log-likelihood estimates are normalized across individuals and variants in their mini-batch training setup, causing premature convergence. This premature convergence negatively impacts their performance, despite `Neural ADMIXTURE` erroneously reporting faster runtimes than it would achieve at optimal solutions.

While our `fastmixture` software does not entirely solve the scalability issues of model-based approaches, it represents a significant step by enabling researchers to estimate accurate ancestry proportions for much larger sample sizes and whole-genome sequencing data. It will also facilitate a more feasible exploration of increased numbers of ancestral sources. We anticipate that `fastmixture` will be the preferred alternative to `ADMIXTURE` in future population genetic studies, and it will also allow researchers to correct for population structure in genome-wide association studies of moderate sample sizes, leveraging our accurate estimates of ancestry proportions.

Data availability

The study uses genotype data from the phase 3 release of the 1000 Genomes Project [26], publicly available at `http://ftp.1000genomes.ebi.ac.uk/vol1/ftp/release/20130502/`. The results (trees and ancestry estimations) and processed genotype data files used in the study are available on Zenodo (`https://doi.org/10.5281/zenodo.12683372`) for reproducibility.

Code availability

The `fastmixture` software is open-source and freely available on GitHub (`https://github.com/Rosemeis/fastmixture`). Scripts for reproducing the entire simulation study are available on Zenodo (`https://doi.org/10.5281/zenodo.12683372`).

Author contributions

J.M. has conceived the study and implemented the method. C.G.S., A.R.M., and J.M. have performed the analyses, discussed the results, and contributed to the final manuscript.

Acknowledgements

J.M. is supported by a grant from the Lundbeck Foundation (R380-2021-1225).

Competing interest statement

The authors declare no competing interests.

References

[1] J. Novembre, T. Johnson, K. Bryc, Z. Kutalik, A. R. Boyko, A. Auton, A. Indap, K. S. King, S. Bergmann, M. R. Nelson, *et al.*, "Genes mirror geography within europe," *Nature*, vol. 456, no. 7218, pp. 98–101, 2008.

[2] J. Marchini, L. R. Cardon, M. S. Phillips, and P. Donnelly, "The effects of human population structure on large genetic association studies," *Nature genetics*, vol. 36, no. 5, pp. 512–517, 2004.

[3] J. K. Pritchard, M. Stephens, and P. Donnelly, "Inference of population structure using multilocus genotype data," *Genetics*, vol. 155, no. 2, pp. 945–959, 2000.

[4] H. Tang, J. Peng, P. Wang, and N. J. Risch, "Estimation of individual admixture: analytical and study design considerations," *Genetic Epidemiology: The Official Publication of the International Genetic Epidemiology Society*, vol. 28, no. 4, pp. 289–301, 2005.

[5] D. H. Alexander, J. Novembre, and K. Lange, "Fast model-based estimation of ancestry in unrelated individuals," *Genome research*, vol. 19, no. 9, pp. 1655–1664, 2009.

[6] B. E. Engelhardt and M. Stephens, "Analysis of population structure: a unifying framework and novel methods based on sparse factor analysis," *PLoS genetics*, vol. 6, no. 9, p. e1001117, 2010.

[7] E. Frichot, F. Mathieu, T. Trouillon, G. Bouchard, and O. François, "Fast and efficient estimation of individual ancestry coefficients," *Genetics*, vol. 196, no. 4, pp. 973–983, 2014.

[8] J. Meisner and A. Albrechtsen, "Inferring population structure and admixture proportions in low-depth ngs data," *Genetics*, vol. 210, no. 2, pp. 719–731, 2018.

[9] I. Cabreros and J. D. Storey, "A likelihood-free estimator of population structure bridging admixture models and principal components analysis," *Genetics*, vol. 212, no. 4, pp. 1009–1029, 2019.

[10] A. M. Chiu, E. K. Molloy, Z. Tan, A. Talwalkar, and S. Sankararaman, "Inferring population structure in biobank-scale genomic data," *The American Journal of Human Genetics*, vol. 109, no. 4, pp. 727–737, 2022.

[11] P. Gopalan, W. Hao, D. M. Blei, and J. D. Storey, "Scaling probabilistic models of genetic variation to millions of humans," *Nature genetics*, vol. 48, no. 12, pp. 1587–1590, 2016.

[12] A. Dominguez Mantes, D. Mas Montserrat, C. D. Bustamante, X. Giró-i Nieto, and A. G. Ioannidis, "Neural admixture for rapid genomic clustering," *Nature Computational Science*, pp. 1–9, 2023.

[13] A. R. Martin, M. Kanai, Y. Kamatani, Y. Okada, B. M. Neale, and M. J. Daly, "Clinical use of current polygenic risk scores may exacerbate health disparities," *Nature genetics*, vol. 51, no. 4, pp. 584–591, 2019.

[14] Y. Wang, K. Tsuo, M. Kanai, B. M. Neale, and A. R. Martin, "Challenges and opportunities for developing more generalizable polygenic risk scores," *Annual review of biomedical data science*, vol. 5, no. 1, pp. 293–320, 2022.

[15] R. Varadhan and C. Roland, "Simple and globally convergent methods for accelerating the convergence of any em algorithm," *Scandinavian Journal of Statistics*, vol. 35, no. 2, pp. 335–353, 2008.

[16] L. Skotte, T. S. Korneliussen, and A. Albrechtsen, "Estimating individual admixture proportions from next generation sequencing data," *Genetics*, vol. 195, no. 3, pp. 693–702, 2013.

[17] N. Patterson, A. L. Price, and D. Reich, "Population structure and eigenanalysis," *PLoS genetics*, vol. 2, no. 12, p. e190, 2006.

[18] S. Ruder, "An overview of gradient descent optimization algorithms," *arXiv preprint arXiv:1609.04747*, 2016.

[19] C. C. Chang, C. C. Chow, L. C. Tellier, S. Vattikuti, S. M. Purcell, and J. J. Lee, "Second-generation plink: rising to the challenge of larger and richer datasets," *Gigascience*, vol. 4, no. 1, pp. s13742–015, 2015.

[20] C. R. Harris, K. J. Millman, S. J. Van Der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, *et al.*, "Array programming with numpy," *Nature*, vol. 585, no. 7825, pp. 357–362, 2020.

[21] Z. Li, J. Meisner, and A. Albrechtsen, "Pcaone: fast and accurate out-of-core pca framework for large scale biobank data," *bioRxiv*, pp. 2022–05, 2022.

[22] F. Baumdicker, G. Bisschop, D. Goldstein, G. Gower, A. P. Ragsdale, G. Tsambos, S. Zhu, B. Eldon, E. C. Ellerman, J. G. Galloway, *et al.*, "Efficient ancestry and mutation simulation with msprime 1.0," *Genetics*, vol. 220, no. 3, p. iyab229, 2022.

[23] G. Tsambos, J. Kelleher, P. Ralph, S. Leslie, and D. Vukcevic, "link-ancestors: fast simulation of local ancestry with tree sequence software," *Bioinformatics Advances*, vol. 3, no. 1, p. vbad163, 2023.

[24] S. Gravel, B. M. Henn, R. N. Gutenkunst, A. R. Indap, G. T. Marth, A. G. Clark, F. Yu, R. A. Gibbs, . G. Project, C. D. Bustamante, *et al.*, "Demographic history and rare allele sharing among human populations," *Proceedings of the National Academy of Sciences*, vol. 108, no. 29, pp. 11983–11988, 2011.

[25] S. R. Browning, B. L. Browning, M. L. Daviglus, R. A. Durazo-Arvizu, N. Schneiderman, R. C. Kaplan, and C. C. Laurie, "Ancestry-specific recent effective population size in the americas," *PLoS genetics*, vol. 14, no. 5, p. e1007385, 2018.

[26] . G. P. Consortium *et al.*, "A global reference for human genetic variation," *Nature*, vol. 526, no. 7571, p. 68, 2015.

Supplementary material

Assessment measures

We here describe the different assessment measures used in the study. The log-likelihood estimate is already defined in equation 1. $\mathbf{Q}$ refers to the estimated ancestry proportions and $\mathbf{Q}^*$ refers to the simulated ground truth ancestry proportions.

Root mean square error (RMSE)

$$\mathrm{RMSE}(\mathbf{Q}^*, \mathbf{Q}) = \frac{1}{NK} \sum_{i=1}^{N} \sum_{k=1}^{K} (q_{ik}^* - q_{ik})^2 \,. \tag{S1}$$

Jensen-Shannon divergence (JSD)

$$\mathrm{JSD}(\mathbf{Q}^*, \mathbf{Q}) = \frac{1}{2} \big(\mathrm{KLD}(\mathbf{Q}^*, \mathbf{X}) + \mathrm{KLD}(\mathbf{Q}, \mathbf{X}) \big) \,, \tag{S2}$$

where $\mathbf{X} = \frac{1}{2}(\mathbf{Q}^* + \mathbf{Q})$ and $\mathrm{KLD}(\cdot, \cdot)$ is the Kullback–Leibler divergence:

$$\mathrm{KLD}(\mathbf{Q}, \mathbf{X}) = \frac{1}{N} \sum_{i=1}^{N} \sum_{k=1}^{K} \left(q_{ik} \log \frac{q_{ik}}{x_{ik}} \right) \,. \tag{S3}$$

Simulated data overview

Scenario A

- $N = 1{,}000$, with 250 individuals from each population (POP1, POP2, POP3, ADMIX)
- $M = 689{,}563$

Scenario B

- $N = 1{,}000$, with 200 individuals from each population (POP1, POP2, POP3, POP4, ADMIX)
- $M = 687{,}107$

Scenario C

- $N = 1{,}000$, with 250 individuals from each population (AFR, EUR, EAS, ADMIX)
- $M = 500{,}114$

Supplementary figures

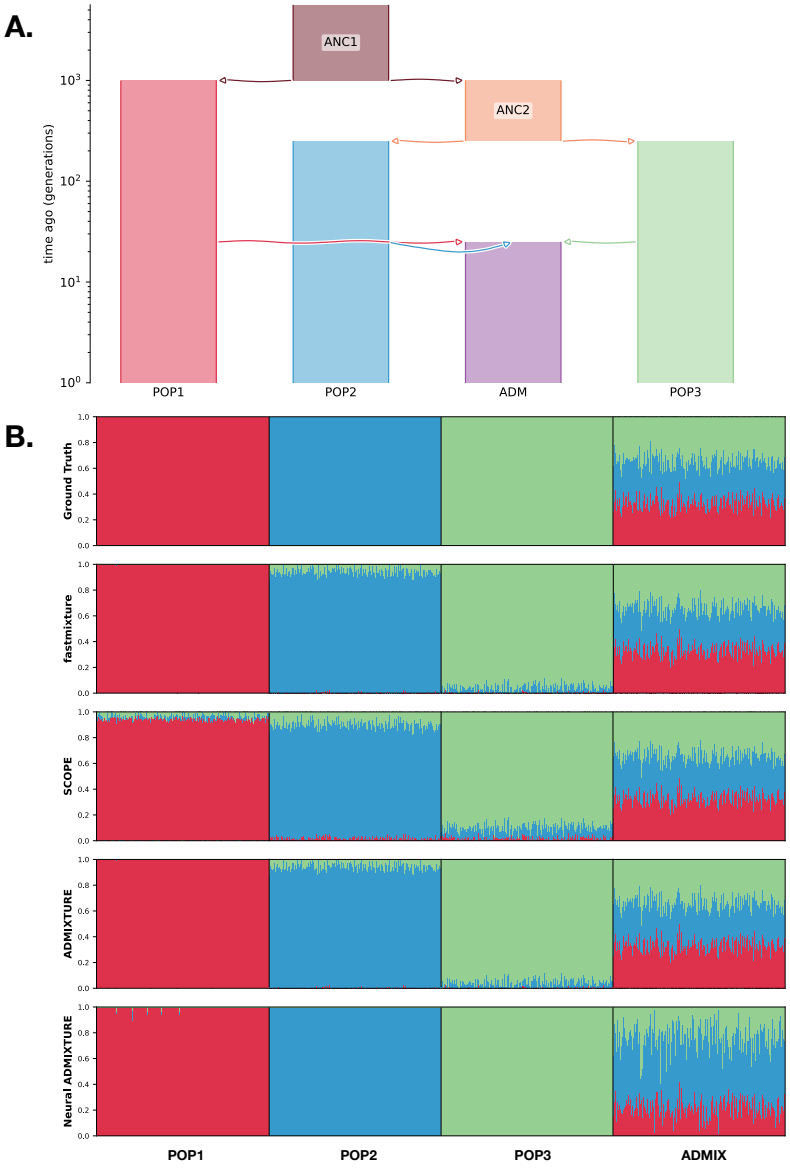

Figure S1: **A.** Demographic model of *Scenario A* with 250 individuals sampled from each of the four populations. **B.** Admixture plots of the ancestry proportions in the simulated individuals for $K = 3$ with the ground truth at the top followed by the different software using their run with the highest log-likelihood.

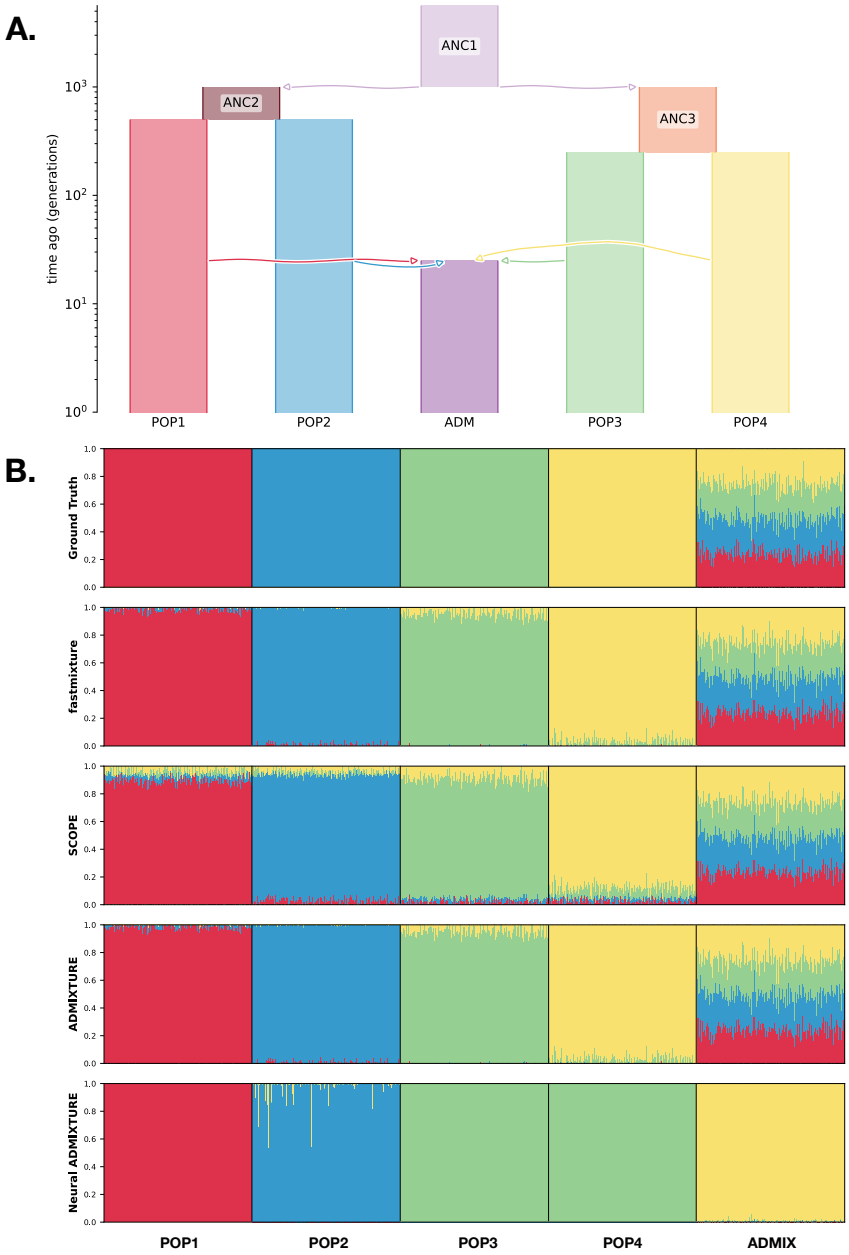

Figure S2: **A.** Demographic model of *Scenario B* with 200 individuals sampled from each of the five populations. **B.** Admixture plots of the ancestry proportions in the simulated individuals for $K = 4$ with the ground truth at the top followed by the different software using their run with the highest log-likelihood.

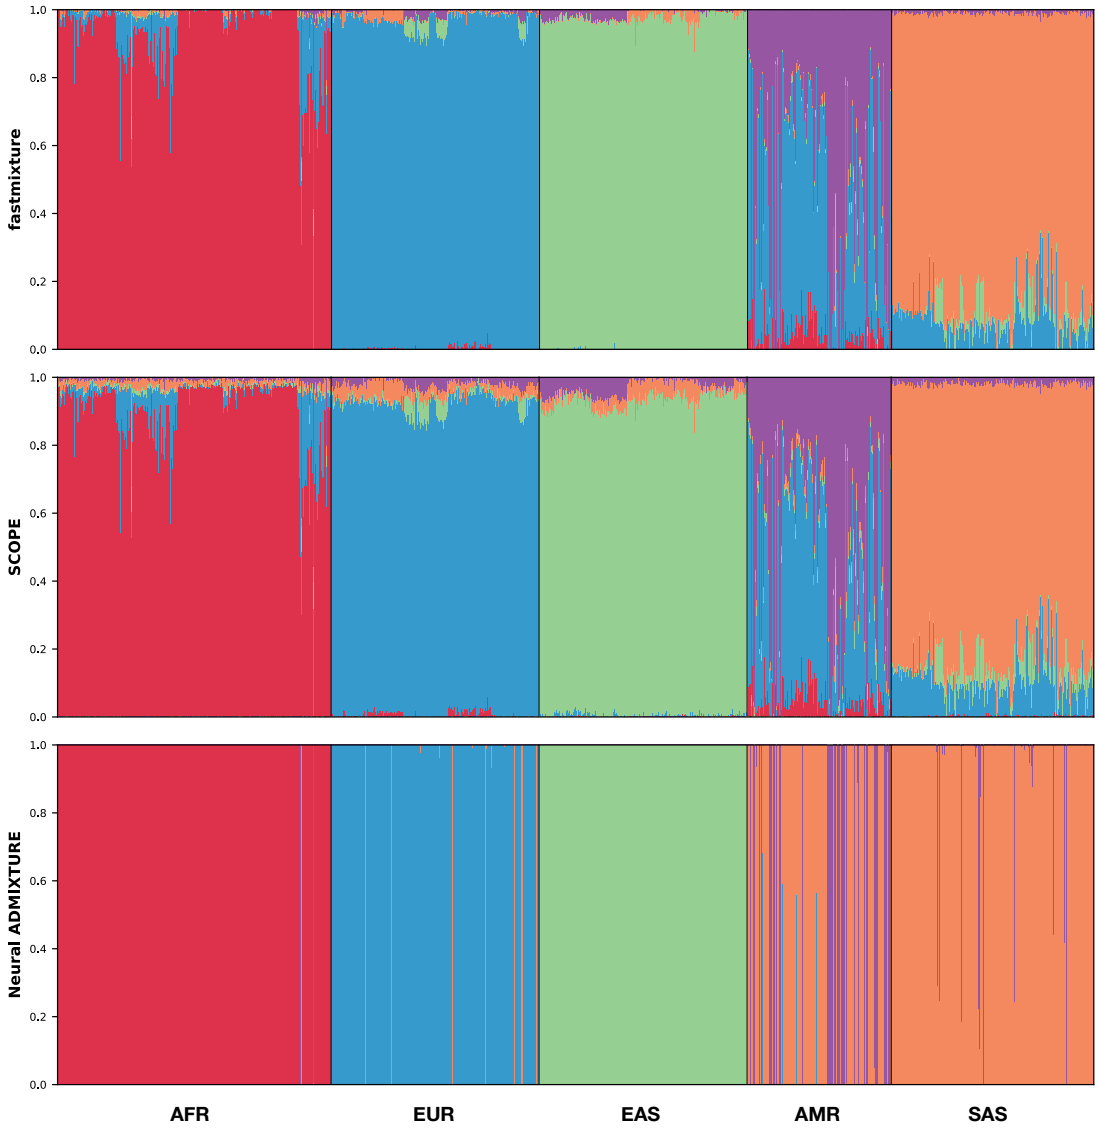

Figure S3: Admixture plots of the estimated ancestry proportions for $K = 5$ in the full version of the 1000 Genomes Project dataset (1KGP) by the different software using their run with the highest log-likelihood.

Supplementary tables

Dataset	`fastmixture`	`Neural ADMIXTURE`	`ADMIXTURE`	`SCOPE`
A	3.62 (0.29)	5.13 (0.13)	77.850 (17.97)	1.69 (0.08)
B	4.00 (0.20)	4.98 (0.12)	105.34 (18.33)	1.96 (0.10)
C	2.23 (0.28)	3.75 (0.12)	36.72 (2.37)	0.24 (0.01)
1KGP Down	11.87 (0.93)	12.55 (0.20)	275.24 (12.77)	1.97 (0.25)
1KGP	160.48 (6.75)	134.92 (2.46)	-	18.69 (2.27)

Table S1: Computational runtimes (in minutes) for the different simulated and empirical datasets for the evaluated methods. The mean across five different runs is reported with the corresponding standard deviation in parenthesis. All analyses have been performed on a cluster with an Intel(R) Xeon(R) Gold 6152 CPU using 64 threads.

Scenario	`fastmixture`	`Neural ADMIXTURE`	`ADMIXTURE`	`SCOPE`
A	0.0095 (1.8e-5)	0.0211 (0.0070)	0.0094 (1.3e-7)	0.0252 (3.0e-5)
B	0.0080 (2.7e-5)	0.2075 (0.0044)	0.0079 (1.5e-7)	0.0331 (4.2e-5)
C	0.0008 (1.2e-6)	0.0187 (0.0129)	0.0008 (4.5e-8)	0.0058 (1.1e-6)

Table S2: Jensen-Shannon divergence (JSD) measures for the estimated ancestry proportions in the three different simulated datasets for the evaluated methods against the ground truth. The mean across five different runs is reported with the corresponding standard deviation in parenthesis.

Dataset	`fastmixture`	`Neural ADMIXTURE`	`ADMIXTURE`	`SCOPE`
A	-637729492.8 (0.1)	-640197951.4 (39763.0)	-637729492.4 (0.5)	-637845454.1 (10.0)
B	-629360598.6 (0.3)	-633930570.4 (37694.0)	-629360598.5 (1.6)	-629538330.5 (38.1)
C	-429286156.1 (4.5)	-433792567.4 (627402.9)	-429286155.6 (4.1)	-429536651.1 (25.0)
1KGP Down	-1509950459.3 (1.4)	-1531090301.0 (228164.6)	-1509950464.3 (7.5)	-1510431876.5 (254.2)
1KGP	-15097015806.1 (1.5)	-15342582278.2 (3449977.8)	-	-15101701104.4 (2807.9)

Table S3: Log-likelihood measures in the different simulated and empirical datasets for the evaluated methods. "1KGP" refers to data from the 1000 Genomes Project, where "1KGP (down)" is the downsampled variant. The mean across five different runs is reported with the corresponding standard deviation in parenthesis. `ADMIXTURE` is not able to finish a single run within 24 hours when using the full SNP set of the 1KGP.