Interpreting f-statistics and admixture graphs: theory and examples

Mark Lipson[1,2]

[1]Department of Genetics, Harvard Medical School, Boston, MA 02115, USA

[2]Department of Human Evolutionary Biology, Harvard University, Cambridge, MA 02138, USA

Email: mlipson@genetics.med.harvard.edu

Abstract

A popular approach to learning about admixture from population genetic data is by computing the allele-sharing summary statistics known as f-statistics. Compared to some methods in population genetics, f-statistics are relatively simple, but interpreting them can still be complicated at times. In addition, f-statistics can be used to build admixture graphs (multi-population trees allowing for admixture events), which provide more explicit and thorough modeling capabilities but are correspondingly more complex to work with. Here, I discuss some of these issues to provide users of these tools with a basic guide for protocols and procedures. My focus is on the kinds of conclusions that can or cannot be drawn from the results of f_4-statistics and admixture graphs, illustrated with real-world examples involving human populations.

Keywords: f-statistics, admixture graphs, admixture, parameter estimation

Introduction

f-statistics (Reich et al., 2009; Patterson et al., 2012) are a widely used toolkit for making inferences about phylogeny and admixture from population genetic data, particularly in humans. The statistics measure correlations in allele frequencies among sets of two, three, or four populations. Observed values reflect degrees of shared ancestry and can serve as a means for testing hypotheses regarding population split orders and past gene flow events under historical models.

As compared to some other common methods in population genetics, f-statistics are quite simple and flexible, but interpreting them is not always straightforward. Additionally, one of the primary applications of f-statistics is in building admixture graphs (i.e., phylogenetic trees augmented with admixture events) with more than four populations, which introduces a greater level of complexity. In this note, I hope to clarify some of these potential difficulties and provide a range of tips for practitioners. Some of the topics have previously been addressed in other places (as cited) but are covered here as well for the sake of completeness.

f-statistics and admixture

Basic definitions and properties

More complete introductions to f-statistics have been published elsewhere (Reich et al., 2009; Patterson et al., 2012; Lipson et al., 2013), but the following are some basics that are used in other sections of the paper. The most general definition is that of the f_4-statistic $f_4(A, B; C, D)$, which measures the average correlation in allele frequency differences between (i) populations A and B and (ii) populations C and D (i.e., $(p_A - p_B) * (p_C - p_D)$,

for allele frequencies p, typically averaged over many biallelic single-nucleotide polymorphisms [SNPs]). This f_4-statistic is the same as the (perhaps more familiar) D-statistic up to a normalization factor. If the four populations are related by the (unrooted) phylogeny $((A, B),(C, D))$, then the expected value of $f_4(A, B; C, D)$ will be zero, while the expected values of $f_4(A, C; B, D)$ and $f_4(A, D; B, C)$ will be positive. Simple algebra shows that

$$f_4(A, B; C, D) = f_4(C, D; A, B),$$

$$f_4(A, B; C, D) = -f_4(B, A; C, D) = -f_4(A, B; D, C),$$

$$f_4(A, B; C, D) = f_4(A, C; B, D) + f_4(A, D; C, B).$$

The other two basic definitions are of the f_2- and f_3-statistics, which can be formulated as $f_2(A, B) = f_4(A, B; A, B)$ and $f_3(A; B, C) = f_4(A, B; A, C)$.

The most important usage for f-statistics is in the context of admixture. If a population C has a mixture of ancestry derived from sources C' and C'' in proportions α and $(1 - \alpha)$, then in expectation,

$f_4(A, B; C, D) = \alpha f_4(A, B; C', D) + (1 - \alpha) f_4(A, B; C'', D).$

Expected values of f-statistics can be visualized in terms of overlapping paths in an admixture graph (Fig. 1; a more extensive illustration is also given in Figure 2 of Patterson et al. (2012)). In the case of admixture, the equations can be derived by forming linear combinations from the constituent ancestry components; the typical expected value is a branch length times a mixture proportion (Fig. 1).

Figure 1. Expected values of f_4-statistics under specified admixture graph models. (A) The expected value of $f_4(A, B; C, D)$ is given by the intersection between the path from A to B with the path from C to D. Under the model shown, $\mathrm{E}[f_4(A, B; C, D)] = 0$. (B) The expected value of $f_4(A, D; B, C)$ is given by the intersection between the path from A to D with the path from B to C. Under the model shown, $\mathrm{E}[f_4(A, D; B, C)] = y$. (C) With population C admixed, the path from B to C can be decomposed into two components. Under the model shown, with a proportion of α B-related ancestry and $1 - \alpha$ D-related ancestry, the former yields a path (lighter red) that has a weight of α but does not intersect the path from A to D, while the latter yields a path (darker red) that has a weight of $1 - \alpha$ and intersects the path from A to D over the branch with length y. In total, $\mathrm{E}[f_4(A, D; B, C)] = (1 - \alpha)y$.

An important point is that, unlike F_{ST} (and normalized D-statistics, at least approximately), the values of f-statistics depend on the absolute allele frequencies of the SNPs used to calculate them (cf. Lipson et al. (2013)). For example, adding fixed sites to the SNP set will shrink f-statistics toward zero. As a result, when comparing multiple f-statistics, it is important that each one should be computed on the same set of SNPs (or as similar as possible). In applications involving ancient DNA, where missing data is common, I typically make the assumption that the SNPs covered for each individual or population are a random subset with respect to allele frequency. By contrast, comparisons across different genotyping arrays are likely to be biased.

Interpreting non-zero f-statistics

If a set of four populations are unadmixed relative to each other, then some permutation of them will yield an f_4-statistic of zero (in expectation), as in Fig. 1A. Equivalently, if all three permutations of f_4-statistics for a certain set of four populations are (significantly) non-zero, then at least one of the populations must be admixed; this is one of the most common signals of admixture used in the literature. In this paper, I will use the example of a quartet consisting of four present-day human populations: Mixe (from Mexico), Han, French, and Baka (hunter-gatherers from Cameroon). As with all Native Americans, Mixe are known to be descended from an (ancient) admixture event involving eastern and western Eurasian lineages, in proportions of approximately 70% and 30% (Raghavan et al., 2014). I computed the three possible f_4-statistics for this quartet and obtained significantly non-zero values, with the signs as expected based on the known history (Table 1). (These and all results in the paper are computed from previously published whole-genome sequence data (Mallick et al., 2016; Fan et al., 2019), on a set of ~1.1 million autosomal SNPs (Mathieson et al., 2015), using the implementation in ADMIXTOOLS (Patterson et al., 2012), including standard errors estimated by block jackknife.)

Table 1. Observed f_4-statistics (values and Z-scores for difference from zero) for the example populations.

Populations				f_4(A, B; C, D)	
A	**B**	**C**	**D**	**Value**	**Z-score**
Mixe	Baka	Han	French	0.011	27.1
Mixe	French	Han	Baka	0.013	35.8
Mixe	Han	Baka	French	-0.0025	-8.9

In this case, there is prior knowledge available about the admixture in Mixe, but in general, without additional information, the existence of such a quartet does not identify which of the four populations is admixed. Here, for example, it could also be that Han is

admixed with most of its ancestry related to Mixe but a small amount related to Baka, and likewise for the other two (see further discussion in the admixture graph sections below). In real-world applications, it can also be true that more than one population is admixed, making the interpretation more complicated. Sometimes, in fact, two admixture events together can cause an f_4-statistic to be close to zero and thereby mask the signal of admixture (at first glance).

Another observation is that as depicted in Fig. 1, f_4-statistics are not only zero or non-zero but also carry quantitative information about amounts of shared drift between populations. One implication is that populations sharing more drift (i.e., yielding longer intersecting paths in an admixture graph) will have greater-magnitude f_4-statistics associated with them. In some cases, this can allow one to identify which of a set of candidate populations is the best proxy (in this phylogenetic sense) for a component of ancestry represented in a test population. However, in practice, this procedure is complicated by the fact that the maximizing reference is not necessarily the closest proxy if the source for that component was itself admixed (Pickrell et al., 2014). A related point is that if a certain signal is weak compared to the noise in the data—for example, if the shared drift branch is short—then one may not have enough power to identify it. Finally, f-statistics can be subject to certain kinds of biases and batch effects (to varying degrees, as with other methods) arising from SNP ascertainment, sample type and processing (ancient versus present-day, sequencing platform, etc.), and other aspects of the data, so it is important to keep such factors in mind when interpreting results.

Admixture graphs: modeling and inference procedure

Fitting an admixture graph

In addition to their stand-alone usage, f-statistics can serve as a means to fit admixture graphs from allele frequency data. In this context, an admixture graph consists of an ordering of population splits, positions of admixture events, branch length parameters, and mixture proportions. Given the first two, the third and fourth can be inferred by solving a system of equations (linear in terms of the branch lengths) in which observed f-statistic values are matched to their expectations in terms of the model parameters. For example, one such equation for the model in Fig. 1B would be $f_2(B, C) = x + y + z$. With n populations, there are $3 \times \binom{n}{4}$ possible f_4-statistics, $3 \times \binom{n}{3}$ possible f_3-statistics, and $\binom{n}{2}$ possible f_2-statistics, but many of these are linearly dependent. In fact, there are a total of $\binom{n}{2}$ linearly independent f-statistic equations (possible bases include (1) the set of all f_2-statistics, and (2) the set of all f_2- and f_3-statistics with a given population in the first position). More extensive descriptions of the admixture graph inference process can be found in Patterson et al. (2012); Lipson et al. (2013); Leppälä et al. (2017).

The software I typically use to build admixture graphs is *qpGraph* (also referred to as ADMIXTUREGRAPH) (Patterson et al., 2012). In *qpGraph*, the user manually specifies the topology of the model, and the program then solves for the optimal values of the parameters. In theory, one might wish to search the entire space of all topologies and parameter values (for a given number of admixture events) to find the best-fitting model, but the size of the space (exponential in the number of populations) makes this impractical for larger graphs (Leppälä et al., 2017). The set of basis statistics used for fitting is the set (2) alluded to in the previous paragraph, with the first population listed in the input file as the "base" population.

In its standard mode, *qpGraph* attempts to minimize the quantity $S(G) = 1/2(g - f)'Q^{-1}(g - f)$, known as the "score" of the model, where f is the vector of observed basis f-statistics (of length $\binom{n}{2}$), g is the vector of predicted f-statistics under the model, and Q is the (estimated) covariance matrix of the statistics. Assuming multivariate normal errors, the score gives the negative log-likelihood of the model. To help insure that Q^{-1} does not become unstable, one can use the "diag" input parameter to add a small number (0.0001 works well in my experience, but smaller values may be sufficient as well) to the diagonal entries of Q. The program can also be run using simple least-squares optimization without the Q matrix ("lsqmode"), but in this case highly correlated statistics will be treated as independent for the sake of the fitting, and the score will no longer represent a log-likelihood, both of which make the full objective function preferable. Other input parameters I typically set are "outpop: NULL" (meaning no specified outgroup population in which SNPs are required to be polymorphic) and "lambdascale: 1" (leaving the f-statistics in typical units rather than scaling into approximate F_{ST}).

By default, *qpGraph* utilizes the set of SNPs that have genotype calls for at least one individual in each population in the model. With low-coverage data (for example, in some ancient DNA applications), this can result in losing the majority of the sites in the initial data set. The program allows an option to use all SNPs instead ("allsnps: YES" or "use-allsnps: YES," in which case each basis statistic is computed on as many sites as possible for the two or three populations involved), but this mode can give unreliable results, in particular when the base population is highly diverged. To the best of my knowledge, this effect is caused by greater absolute noise when estimating larger-magnitude basis statistics, such that the small relative fluctuations in empirical f-statistics caused by modest changes in the SNP set become substantial in the context of the admixture graph. In my own work, my preference has always been to avoid using the all-SNPs option. If this

causes an undesirable loss of coverage, then the best approach given the current imple-
mentation of *qpGraph* is probably to set as the base a population that (a) is not highly
diverged from the others in the model, and (b) preferably has multiple individuals with
diploid data (again to reduce the magnitudes of the statistics).

Parameters and constraints

An important consideration is whether the system of equations used to infer the param-
eters of an admixture graph is over- or under-determined. As mentioned above, a model
with n populations has $\binom{n}{2}$ linearly independent constraints (i.e., equations). In the ab-
sence of admixture, there are $2n - 3$ parameters, which is the number of branches in an
unrooted binary tree with n leaf nodes (with the settings I have described, *qpGraph* results
should not depend on where the root of a graph is specified). Converting a population
from unadmixed to admixed adds two parameters: one for the mixture proportion and
one for the split position of the new source of ancestry. Thus, with a admixture events,
the total number of free parameters is $2n + 2a - 3$. One point to note is that in the case
of an admixed population with two unsampled sources (which is the typical scenario), the
three branch lengths surrounding the admixture event (in Fig. 2A, from the node "East1"
to "East2," from "West1" to "West2," and from "pAM1" to Mixe) cannot be determined
individually but instead form a single compound parameter $\alpha^2 x + (1-\alpha)^2 y + z$ (where α is
the mixture proportion, x and y are the branch lengths to the two corresponding sources,
and z is the terminal branch length). The only exception (to my knowledge) is the case
in which at least three populations are included that can be modeled as having different
proportions of ancestry from the same two sources, which allows the branch lengths to be
solved for individually.

Even if the inequality $\binom{n}{2} \geq 2n + 2a - 3$ is satisfied for an admixture graph as a

whole, there can be some parameters that are not uniquely determined because of repetition across the different equations caused by multiple populations in phylogenetically equivalent positions. Further discussion of this phenomenon can be found in the example sections below. Additionally, the notion of sufficient constraint is not absolute. A model can have enough populations in distinct positions in order to be able to estimate a mixture proportion, but if two of them are only slightly separated, then the precision of the estimate will generally be lower. Similarly, if one of the populations providing the constraint is itself admixed, then the power will often be reduced.

Fit quality

To my knowledge, no absolute measure of model fit has been developed for admixture graphs, but there are several ways to evaluate how well a given model fits the data (see also Lipson and Reich (2017); Lipson et al. (2017, 2020)). The following discussion is tailored for *qpGraph*, but the ideas also apply more generally. First, the program returns a list of residual poorly-predicted f-statistics and their Z-scores (drawn from the set of all possible f-statistics, not only those in the basis), which can give a good sense for the performance of the model and some idea of which populations are responsible for the greatest inaccuracies. There is no general rule for what threshold constitutes a significantly non-zero residual; the situation is complicated because there are many statistics being tested simultaneously, but many of those are also correlated with each other.

Deviations between model predictions and the observed data can be caused either by an incorrectly specified topology or un-modeled admixture. In the first case, assuming that the program does not get stuck at a local optimum, it will try to move the populations as close as possible to their correct positions but will be constrained by the input topology.

Thus, an incorrectly specified split order usually manifests as an inferred length-zero internal branch; when such branches (i.e., trifurcations) appear in the results, the order of splits should be adjusted and re-tried. (The default *qpGraph* visualization output rounds branch lengths to the nearest integer, so some non-zero-length but very short branches may initially appear as zero.) As noted in the *f*-statistics section above, however, one may not have sufficient power to resolve short branches, so some sets of three lineages may be found to be statistically consistent with forming a trifurcation, with all three possible split orders having similar fit quality.

In the case of un-modeled admixture, the observed deviations could potentially reflect admixture in one of multiple different populations. Often one can gain information by examining the full list of residuals and noting which populations occur repeatedly. Another approach is to remove one population from the model and see if the fit improves, although even if it does, that could imply either that the population in question had un-modeled admixture or that it provided a constraint enabling the detection of un-modeled admixture among the other populations.

The score of the final graph is also returned as an output from the program, so it can be used to compare the fit quality of different models with the same set of populations, preferring the one with the lower score. (If the equations being fit were independent, then one could apply a chi-squared test for the overall fit, but in practice they are heavily correlated. *qpGraph* returns a naive degrees of freedom count and *p*-value alongside the score, but they are not well calibrated.) As above, while this approach provides a useful heuristic, evaluating statistical significance is complicated, and I do not have a rigorous set of recommendations. One recent direction that seems promising is using the score to compare alternative models with the same populations and same number of admixture events. In that case, the score difference can be interpreted in an AIC/BIC framework,

with the likelihood difference as a Bayes factor (Leppälä et al., 2017; Flegontov et al., 2019; Shinde et al., 2019). The same idea could also be applied in cases with unequal numbers of free parameters—for example, adding one admixture event and testing whether the score improvement is significant. However, defining the change in degrees of freedom is not straightforward in this situation: as noted above, a new admixture event creates two additional parameters in the model, but that does not account for whether the admixture comes from a pre-specified source or from a source that is allowed to be located anywhere in the graph. Finally, the score can additionally be used to compute confidence intervals on parameters (by considering the likelihood as a function of different values of a single branch length or mixture proportion), although it is worth keeping in mind that the results are model-dependent.

Admixture graphs: examples

Four populations

The first examples I will present are four-population admixture graphs containing Mixe, Han, French, and Baka. Given the observed non-zero f_4-statistics in Table 1, there must be at least one admixture event present in order to fit the data. However, in light of the discussions above about determining which population is admixed and about parameters and constraints in admixture graphs, it would be expected that these models should be insufficiently constrained to determine which population is admixed. Indeed, they have $\binom{4}{2} = 6$ constraints but $2(4) + 2(1) - 3 = 7$ free parameters. And, as expected, the inferred model in which Mixe is specified as admixed fits the data perfectly (i.e., a set of branch length and mixture proportion parameters can be chosen so that the six basis f-statistics are predicted exactly by the model, yielding $S(G) = 0$; Fig. 2A), but perfect fits can

also be obtained when the other three populations are (incorrectly) specified as admixed instead (Fig. 2B–D).

Figure 2. Four-population admixture graphs modeling (A) Mixe, (B) Baka, (C) Han, or (D) French as admixed. All four versions provide perfect fits to the data (exact agreement between observed and predicted f-statistics). In this and all following figures, branch lengths (in f-statistic units, multiplied by 1000) are rounded to the nearest integer.

Interestingly, for some parameter values, the admixed population can be determined even with only four populations in the model: if a negative f_3-statistic can be formed for some triple, then the "target" population of the statistic must be admixed. To give an example, I replaced Mixe with Kyrgyz in the four-population model. With Kyrgyz modeled as admixed, the fit is perfect as before (Fig. 3A). With Baka modeled as admixed, however, the fit is very poor, with residuals up to $Z = 27$ (Fig. 3B). The most extreme residual is the statistic f_3(Kyrgyz; Han, French), which has an observed value of -0.0064 ($Z = 27$ for difference from zero) but can only be negative if Kyrgyz (in the position of the test population in a "three-population test" for admixture) is admixed (Reich et al., 2009; Patterson et al., 2012).
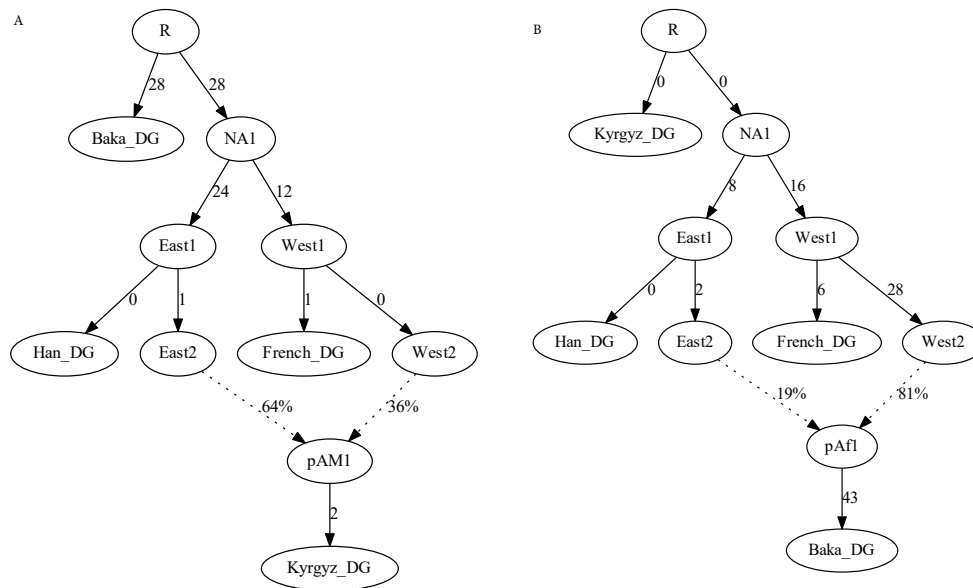
Figure 3. Four-population admixture graphs with Kyrgyz in place of Mixe, modeling either (A) Kyrgyz or (B) Baka as admixed. The first provides a perfect fit to the data, whereas the second has residuals up to $Z = 27$.

Another note is that in these examples, I have been focusing on the primary signal of deep eastern/western Eurasian admixture in Mixe. The other populations are also

admixed in their own ways; for example, all of the non-Africans have small proportions of Neanderthal ancestry, and Baka are admixed with ancestry related to nearby Bantu-speaking farmers (Fan et al., 2019). However, the first signal is not evident in the data without deeper outgroups present, and the second without other African populations. Conversely, if the model contained several sub-Saharan African populations plus Mixe as the only non-Africans, then the primary signal in our examples here would not be visible. In some ways, this inability to detect certain admixture events is beneficial, as it means that models can be constructed so as to focus on events of interest while ignoring some that are outside the desired scope.

Five populations

In general, in order to be able to solve for the parameters of an admixture graph including one admixture event, it is necessary to use at least five populations, providing $\binom{n}{2} = 10$ constraints for the $2n + 2a - 3 = 9$ free parameters. Concurrently, in contrast to the four-population examples above, having five populations present allows one to determine which of the populations is admixed, as long as the topological relationships of the populations are all unique relative to the true mixing sources. More detail on this last point can be found elsewhere (Pease and Hahn, 2015; Lipson and Reich, 2017). A simple version of this statement is that, at least in the case of a single admixture event, one four-population subset will be unadmixed, whereas the other four subsets will include the admixed population. Similarly, in order to solve for a given mixture proportion in a larger graph, there must four populations present (aside from the admixed one in question) in distinct positions, yielding a non-redundant five-population subgraph; three populations in distinct positions allows one to detect the signal of admixture but not to determine the proportion uniquely.

As an example, I added Ulchi (from the Amur River Basin of northeastern Asia) as a fifth population alongside the four from above. Ulchi splits closer to the eastern Eurasian source population for Mixe than does Han, which provides the additional degree of constraint. The five-population model is a good fit to the data ($Z = 1.9$ for the most significant residual), but not a perfect one (Fig. 4A). By contrast, if Baka are modeled as admixed instead of Mixe, the fit is poor ($Z = 4.7$; Fig. 4B). I also show an example where the topology is incorrectly specified, with Han closer than Ulchi to the eastern Eurasian source population for Mixe (Fig. 4C); this version fits poorly ($Z = 5.7$), and the branch connecting the split positions of Ulchi and Han collapses to length zero.

Figure 4. Five-population admixture graphs. (A) Standard four-population example plus Ulchi; all f-statistics are predicted to within 1.9 standard errors of their observed values. (B) Same five populations, but with Baka modeled as admixed; residual statistics are present up to $Z = 4.7$ (C) Same five populations, with Mixe modeled as admixed, but with the positions of Han and Ulchi reversed; residual statistics are present up to $Z = 5.7$. (D) Original four populations plus Hungarian, with Baka modeled as admixed; all f-statistics are predicted to within 1.2 standard errors of their observed values.

Having five populations present (with a single admixture event) also provides the

ability to infer uniquely optimal parameter values. In the four-population example model, the initial estimate of eastern Eurasian ancestry in Mixe was 71%, but with the proportion manually set at 75%, the fit is still perfect (Fig. 5A). Outside of a certain range of mixture proportions (dependent on the values of the branch lengths), the fit will become worse, but within a finite interval, the likelihood is entirely flat. In terms of f_4-statistics, the observed non-zero value is being fit as equal to a branch length in the admixture graph times the mixture proportion (as in Fig. 1C), but without additional constraint, that product can remain the same while the branch length and mixture proportion covary (where the range is determined by bounds on the individual parameter values, e.g., positivity). With five populations, however, there is a unique optimal solution; for example, if I set the mixture proportion at 70% eastern Eurasian ancestry (as compared to the point estimate of 76% in the five-population model), there are residuals up to $Z = 2.6$ (Fig. 5B), and the score is more than 10 units worse. Even in the example above with Kyrgyz (i.e., a four-population model where the admixed population can be determined because of a negative f_3-statistic; Fig. 3), the parameters remain not uniquely determined.
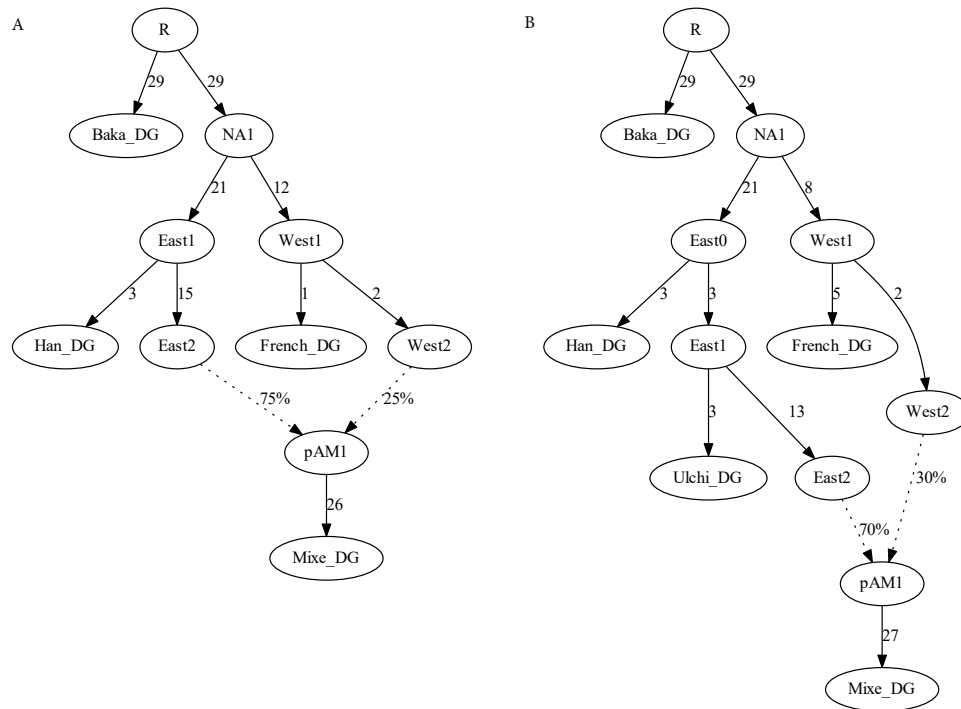
Figure 5. Admixture graphs with pre-specified mixture proportion parameters. (A) Four-population model, with the proportion locked at 75%; the fit is perfect. Note that the branch lengths shift slightly relative to Fig. 2A. (B) Five-population model, with the proportion locked at 70%; residual statistics (indicating a need for more eastern Eurasian ancestry in Mixe) are present up to $Z = 2.6$.

Finally, in Fig. 4D, I show a model with the original four populations plus Hungarian instead of Ulchi. Although there are five populations present, French and Hungarian can be modeled as sister groups, so equations relating parameters in the graph to statistics of the form $f_2(\text{French}, X)$ and $f_2(\text{Hungarian}, X)$ are linearly dependent (up to their terminal branch lengths) and hence do not contribute fully independent constraints. This can be seen in the results, as Baka can successfully be modeled as the admixed population (with residuals up to $Z = 1.2$ reflecting small observed asymmetries between French and Hungarian). This contrasts with Ulchi, which has a distinct phylogenetic position from Han (relative to the other populations in the model) and thus adds new constraints

(although it is worth noting again that a population with only a slightly different position would add constraint but only weakly).

Discussion

Most of the results in this paper pertaining to admixture graphs have been formulated from the perspective of the *qpGraph* software, but other fitting methods are also available. At the level of optimization scheme, the results have assumed that models are fit based on a distance metric (specifically, f-statistics). There are also other possibilities; for example, the *TreeMix* algorithm (Pickrell and Pritchard, 2012) is conceptually similar, and the results are comparable, but it is based on a maximum-likelihood framework.

Different methods also take different approaches to automation and the selection of which populations to model as admixed. *qpGraph* leaves the choice of how many admixture events to include (and which populations are admixed) up to the user; some guidelines pertaining to this choice have been discussed above. For smaller models, it can also be possible to search some or all of the full graph space (Shinde et al., 2019) to determine best-fitting topologies for a given number of admixture events (for example, using the similar *admixturegraph* implementation in R (Leppälä et al., 2017)). *MixMapper* (Lipson et al., 2013) provides an intermediate level of automation by attempting to infer an unadmixed sub-model and then fitting one or two admixed populations onto this scaffold. With a small set of populations, this can sometimes be a useful approach, but it can largely be recapitulated within *qpGraph*, and the software does not support large models with more admixture events.

At the most automated end of the spectrum is *TreeMix* (Pickrell and Pritchard, 2012), which only asks the user to supply the list of populations and the number of admixture

events and then returns a single best-fitting model. The advantage of this strategy is that the program does all of the work of building the model, which is especially useful if one has limited prior knowledge about the populations. The main drawback, in my view, is that the way the program builds the model is by starting with an optimal mixture-free tree and then adding admixture events to account for deviations between the predictions of the tree model and the observed data. Depending on the true histories of the populations, this approach can be successful, but it can also increase the chances of falling into local optima imposed by the initial tree (especially if many populations are admixed). Additionally, the choice of how many admixture events to include, which can sometimes be difficult, is still left to the user.

In my experience, I have found f-statistics and admixture graphs to be very useful tools for learning about phylogeny and admixture. I hope that this guide will help others to get the most out of these tools in a wide range of real-world applications.

Acknowledgments

I would like to thank David Reich, Vagheesh Narasimhan, Nick Patterson, Robert Maier, Iosif Lazaridis, and Pavel Flegontov for helpful discussions.

References

Fan, S., Kelly, D. E., Beltrame, M. H., Hansen, M. E., Mallick, S., Ranciaro, A., Hirbo, J., Thompson, S., Beggs, W., Nyambo, T., et al. (2019). African evolutionary history inferred from whole genome sequence data of 44 indigenous African populations. *Genome Biol.*, 20(1):82.

Flegontov, P., Altınışık, N. E., Changmai, P., Rohland, N., Mallick, S., Adamski, N., Bolnick, D. A., Broomandkhoshbacht, N., Candilio, F., Culleton, B. J., et al. (2019). Palaeo-Eskimo genetic ancestry and the peopling of Chukotka and North America. *Nature*, 570:236–240.

Leppälä, K., Nielsen, S. V., and Mailund, T. (2017). admixturegraph: An R package for admixture graph manipulation and fitting. *Bioinformatics*, 33(11):1738–1740.

Lipson, M., Loh, P.-R., Levin, A., Reich, D., Patterson, N., and Berger, B. (2013). Efficient moment-based inference of admixture parameters and sources of gene flow. *Mol. Biol. Evol.*, 30(8):1788–1802.

Lipson, M. and Reich, D. (2017). A working model of the deep relationships of diverse modern human genetic lineages outside of Africa. *Mol. Biol. Evol.*, 34(4):889–902.

Lipson, M., Ribot, I., Mallick, S., Rohland, N., Olalde, I., Adamski, N., Broomandkhoshbacht, N., Lawson, A. M., López, S., Oppenheimer, J., et al. (2020). Ancient West African foragers in the context of African population history. *Nature*, 577:665–670.

Lipson, M., Szécsényi-Nagy, A., Mallick, S., Pósa, A., Stégmár, B., Keerl, V., Rohland, N., Stewardson, K., Ferry, M., Michel, M., et al. (2017). Parallel palaeogenomic transects reveal complex genetic history of early European farmers. *Nature*, 551(7680):368–372.

Mallick, S., Li, H., Lipson, M., Mathieson, I., Gymrek, M., Racimo, F., Zhao, M., Chennagiri, N., Nordenfelt, S., Tandon, A., et al. (2016). The Simons Genome Diversity Project: 300 genomes from 142 diverse populations. *Nature*, 538(7624):201–206.

Mathieson, I., Lazaridis, I., Rohland, N., Mallick, S., Patterson, N., Roodenberg, S. A., Harney, E., Stewardson, K., Fernandes, D., Novak, M., et al. (2015). Genome-wide patterns of selection in 230 ancient Eurasians. *Nature*, 528(7583):499–503.

Patterson, N., Moorjani, P., Luo, Y., Mallick, S., Rohland, N., Zhan, Y., Genschoreck, T., Webster, T., and Reich, D. (2012). Ancient admixture in human history. *Genetics*, 192(3):1065–1093.

Pease, J. B. and Hahn, M. W. (2015). Detection and polarization of introgression in a five-taxon phylogeny. *Syst. Biol.*, 64(4):651–662.

Pickrell, J. and Pritchard, J. (2012). Inference of population splits and mixtures from genome-wide allele frequency data. *PLoS Genet.*, 8(11):e1002967.

Pickrell, J. K., Patterson, N., Loh, P.-R., Lipson, M., Berger, B., Stoneking, M., Pakendorf, B., and Reich, D. (2014). Ancient west Eurasian ancestry in southern and eastern Africa. *Proc. Natl. Acad. Sci. U. S. A.*, 111(7):2632–2637.

Raghavan, M., Skoglund, P., Graf, K. E., Metspalu, M., Albrechtsen, A., Moltke, I., Rasmussen, S., Stafford Jr, T. W., Orlando, L., Metspalu, E., et al. (2014). Upper Palaeolithic Siberian genome reveals dual ancestry of Native Americans. *Nature*, 505(7481):87–91.

Reich, D., Thangaraj, K., Patterson, N., Price, A., and Singh, L. (2009). Reconstructing Indian population history. *Nature*, 461(7263):489–494.

Shinde, V., Narasimhan, V. M., Rohland, N., Mallick, S., Mah, M., Lipson, M., Nakatsuka, N., Adamski, N., Broomandkhoshbacht, N., Ferry, M., et al. (2019). An ancient Harappan genome lacks ancestry from Steppe pastoralists or Iranian farmers. *Cell*, 179(3):729–735.

Data Accessibility

The data that support the findings of this study are openly available through the European Nucleotide Archive (ENA), under accession numbers PRJEB9586 and ERP010710, and at the European Genome-phenome Archive (EGA), under accession number EGAS00001001959 (Mallick et al., 2016; Fan et al., 2019).