

Exploring Evolutionary Relationships Across the Genome Using Topology Weighting

Simon H. Martin^{*.1} and Steven M. Van Belleghem^{*.1.†}

^{*}Department of Zoology, University of Cambridge, CB2 3EJ, United Kingdom, [†]Department of Biology, Center for Applied Tropical Ecology and Conservation, University of Puerto Rico, Rio Piedras 00931, Puerto Rico, and [‡]Smithsonian Tropical Research Institute, Panamá City 0843-03092, Panama

ORCID IDs: 0000-0002-0747-7456 (S.H.M.); 0000-0002-5083-8114 (S.M.V.B.)

ABSTRACT We introduce the concept of topology weighting, a method for quantifying relationships between taxa that are not necessarily monophyletic, and visualizing how these relationships change across the genome. A given set of taxa can be related in a limited number of ways, but if each taxon is represented by multiple sequences, the number of possible topologies becomes very large. Topology weighting reduces this complexity by quantifying the contribution of each taxon topology to the full tree. We describe our method for topology weighting by iterative sampling of subtrees (*Twisst*), and test it on both simulated and real genomic data. Overall, we show that this is an informative and versatile approach, suitable for exploring relationships in almost any genomic dataset. Scripts to implement the method described are available at <http://github.com/simonhmartin/twisst>.

KEYWORDS phylogenomics; introgression; incomplete lineage sorting; population genomics

THE relationship (or genealogy) among recombining DNA sequences from closely related taxa often varies across the genome, due to both variation in lineage sorting and introgression (Maddison 1997). Numerous methods focus on inferring, from this genealogical variation, the underlying population branching pattern (the species tree) (e.g., Heled and Drummond 2010) or demographic history (e.g., Lohse *et al.* 2012). It is now also possible to characterize the complete genomic landscape of relatedness using whole genome sequences. For example, Hobolth *et al.* (2007) and Dutheil *et al.* (2009) developed an approach that uses whole genome sequences to infer not only the population history, but also how and where in the genome the genealogy changes. More recent studies have attempted to characterize patterns of relatedness along larger numbers of whole

genomes, either by simply inferring phylogenetic trees for predefined windows (Martin *et al.* 2013; Fontaine *et al.* 2015), or by attempting to infer both the trees and the likely breakpoints that separate them (Gante *et al.* 2016). An emerging challenge with increasing numbers of sequences, is that both the inference and interpretation of genealogies becomes difficult, due to the rapid escalation of topological complexity. For example, for five haploid sequences, there are 15 possible unrooted, bifurcating tree topologies, whereas for 10 sequences there are >2 million. Here we address this challenge of characterizing and summarizing the genomic landscape of relatedness in large datasets with multiple genomes from multiple taxa.

One way to deal with the problem of increasing tree complexity is to focus specifically on the relationships among broader predefined taxa (hereafter taxon topologies), and not among all of the sequences. This is straightforward if each taxon is completely resolved into a monophyletic clade, so that the branching patterns within each taxon can simply be ignored, but it becomes challenging when the taxa are not reciprocally monophyletic (*i.e.*, when lineages are not completely sorted). This is often the case for closely related taxa, in which lineages from the same population coalesce (share a most recent common ancestor) in the ancestral population, and may therefore be more closely related to lineages

Copyright © 2017 Martin *et al.*

doi: <https://doi.org/10.1534/genetics.116.194720>

Manuscript received August 11, 2016; accepted for publication March 14, 2017; published Early Online March 23, 2017.

Available freely online through the author-supported open access option.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Supplemental material is available online at www.genetics.org/lookup/suppl/doi:10.1534/genetics.116.194720/-/DC1.

¹Corresponding author: Department of Zoology, University of Cambridge, Downing St., Cambridge, CB2 3EJ, United Kingdom. E-mail: shm45@cam.ac.uk

from other taxa than from their own (Figure 1A). We note that tree inference using large genomic windows, entire chromosomes, or whole genomes often yields completely sorted monophyletic taxa, but this may be artificial as such methods are usually forced to infer a single best-supported tree, even though the region may represent multiple distinct incompletely sorted ancestries. Population genetic statistics based on allele frequencies, such as F_{ST} , have the advantage of quantifying relatedness rather than simply qualitatively describing the topology. However, by ignoring topology entirely, they become difficult to interpret when the number of taxa is more than two. Therefore, there is a need for descriptive methods that incorporate both the qualitative tree-like structure of relationships and the quantitative variation in taxon relationships across the genome. Examples of such methods include the so-called ABBA-BABA test (Green *et al.* 2010; Durand *et al.* 2011; Martin *et al.* 2015) and f statistics (Reich *et al.* 2009, 2012; Patterson *et al.* 2012), both of which evaluate the support for alternative taxon topologies using allele frequencies at single nucleotide polymorphisms (SNPs). However, because individual SNPs are only informative for defining two separate groupings, these methods do not scale to more than four taxa.

Here we introduce the concept of topology weighting, which offers a simple and general solution to the problem of quantifying relationships among taxa that are not necessarily monophyletic. Given a tree of relationships for a set of taxa, each represented by an arbitrary number of sequences, topology weighting quantifies the contribution of each individual taxon topology to the full tree. We describe our approach to compute topology weightings, which we call *Twisst* (topology weighting by iterative sampling of subtrees), and explore the utility of this approach using simulated data as well as two different genomic datasets from butterflies and fungi. Overall, we show that this concept provides a useful means to explore relationships using genomic data, both to test hypotheses and generate new ones.

Materials and Methods

Topology weighting by iterative sampling of subtrees

A given set of taxa can be related in a limited number of ways. For example, for four taxa labeled A, B, C, and D, there are three possible unrooted bifurcating topologies: [(A,B),C,D], [(A,C),B,D], and [(A,D),B,C] (Figure 1A). Given a tree with any number of tips (or leaves), each assigned to a particular taxon, we define the weighting of a particular taxon topology, τ , as the fraction of all unique subtrees (in which each taxon is represented by a single tip) that match the particular taxon topology

$$w_{\tau} = \frac{K}{N}, \quad (1)$$

where N is the number of possible unique sample sets in which each taxon is represented by a single sample. This corresponds to the product of the number of samples in each

taxon. K is the number of these unique sample sets for which the corresponding subtree topology matches τ . Thus,

$$N = \prod_{j=1}^n s_j, K = \sum_{i=1}^N [x_i = \tau], \quad (2)$$

in which n is the number of defined taxa (groups), s_j is the number of samples in taxon j , and x_i is the subtree topology corresponding to subset i .

Topology weighting therefore reduces the complexity of the full tree to a number of values, each giving the proportionate contribution of a particular taxon topology (Figure 1A). This method has conceptual similarities to the quartet sampling approach for comparing topologies (Estabrook *et al.* 1985), except that here we do not consider all subtrees, but only those in which each tip represents a different taxon. Moreover, topology weighting can be applied to any number of taxa. The taxa can be defined arbitrarily, for example by phenotype or geography, as with the operational taxonomic units used in biogeography. Because the number of taxon topologies is limited, the weightings can be normalized to sum to 1, making them easily comparable between different parts of the genome.

We computed topology weightings using our *Twisst* approach, implemented by Python scripts (available for download at <https://github.com/simonhmartin/twisst>). The *Twisst* algorithm first computes all possible unrooted, bifurcating taxon topologies, and then determines the number of unique subtrees that match each topology by iteratively sampling a single individual from each taxon and “pruning” away all other branches and nodes. Our implementation makes use of the Environment for Tree Exploration toolkit version 3 (Huerta-Cepas *et al.* 2016).

The number of unique sample combinations (and corresponding subtrees) can be very large. However, the iterative process is sped up considerably by first collapsing monophyletic groups of samples from the same taxon and weighting these nodes proportionately (Supplemental Material, Figure S1 in File S2). Nevertheless, if the taxa are highly unsorted and the tree is large, it may not be possible to consider all possible sample combinations in a reasonable amount of time. In such cases, approximate weightings can be computed by randomly sampling a subset of sample combinations (Figure S2 in File S2). Random sampling for approximate weighting is performed with replacement, so that the errors fit a binomial distribution, allowing for the computation of a confidence interval (Figure S2 in File S2). Our implementation of *Twisst* also allows a threshold-based sampling procedure, in which sampling is repeated until a particular level of confidence is achieved. This allows further speed-ups since highly sorted trees, in which weightings are biased toward one or a few topologies, require less sampling for high confidence than unsorted or star-like trees, where most weightings are intermediate. For all analyses in this study, dataset sizes allowed for computation of complete weightings.

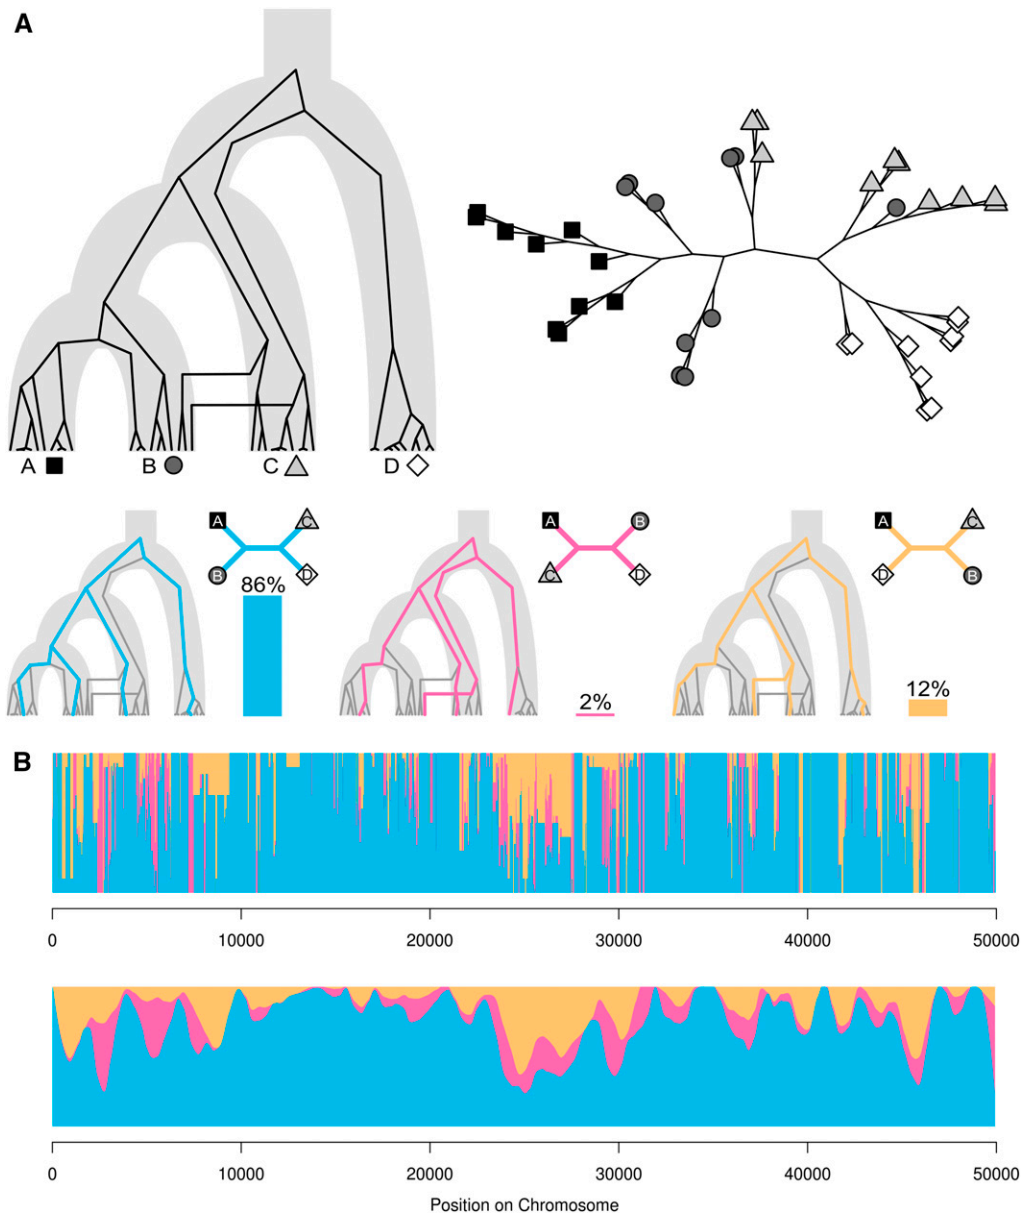


Figure 1 Topology weighting. (A) An example genealogy for four taxa (A, B, C and D), plotted as an unrooted tree on the top right. Taxa B and C are not monophyletic, due to both deep coalescences causing incomplete lineage sorting and gene flow. The three possible taxon topologies are shown, along with a single example subtree that matches each topology. The percentage of all subtrees matching each taxon topology (*i.e.*, the weightings) are shown by vertical bars. (B) Topology weightings plotted across a 50 kb region of a simulated recombining chromosome. Weightings for the three topologies are stacked (they always sum to 1, as they are proportions). Changes in the weightings along the chromosome indicate regions of distinct genealogical history separated by recombinations. Below, the same data are plotted with loess smoothing (span = 2.5 kb).

Analysis of simulated chromosomes

In order to test our method on realistic data for which the complete genealogical history is known, we simulated the evolution of recombining chromosomes using the coalescent simulator *msms* (Ewing and Hermisson 2010). Simulations involved four taxa, each represented by 10 haploid sequences (40 in total) that split in the order $\{(A,B),C\},D\}$ (Figure 2A), or five taxa, each represented by six sequences, that split in the order $\{(A,B),(C,D)\},E\}$ (Figure S7A in File S2), with split times of 0.5, 1, and 1.5 (in units of $4N$ generations). Population size was constant throughout. In all scenarios, unidirectional migration from C to B was simulated. The simulation was performed for a 1 Mb chromosome, with a population recombination rate ($4Nr$) of 0.01 or 0.001. Genealogies for each unique ancestry block (separated by recombinations) were recorded. These were used to calculate the true weightings using *Twisst*.

Three distinct evolutionary scenarios were simulated (Figure 2A and Figure S7A in File S2). *msms* command options are provided in File S1. The first was a “Neutral” scenario, with no selection and a low migration rate from B to C of 0.1 (in units of $4Nm$, where m is the fraction of B made up of migrants from C each generation). The second was an “Adaptive Introgression” scenario, which is the same as above except that a beneficial allele at a locus in the center of the 1 Mb chromosome is allowed to move from population C into B at time 0.1. This was achieved by initiating selection at this time point on a dominant allele that was fixed in population C and absent from the other populations. A selection coefficient of 0.005 was used for both the homozygote and heterozygote, with a diploid population size of 100,000, giving a selection strength ($2Ns$) of 1000 for both genotypes. The third was a “Barrier Locus” scenario, where the rate of migration from

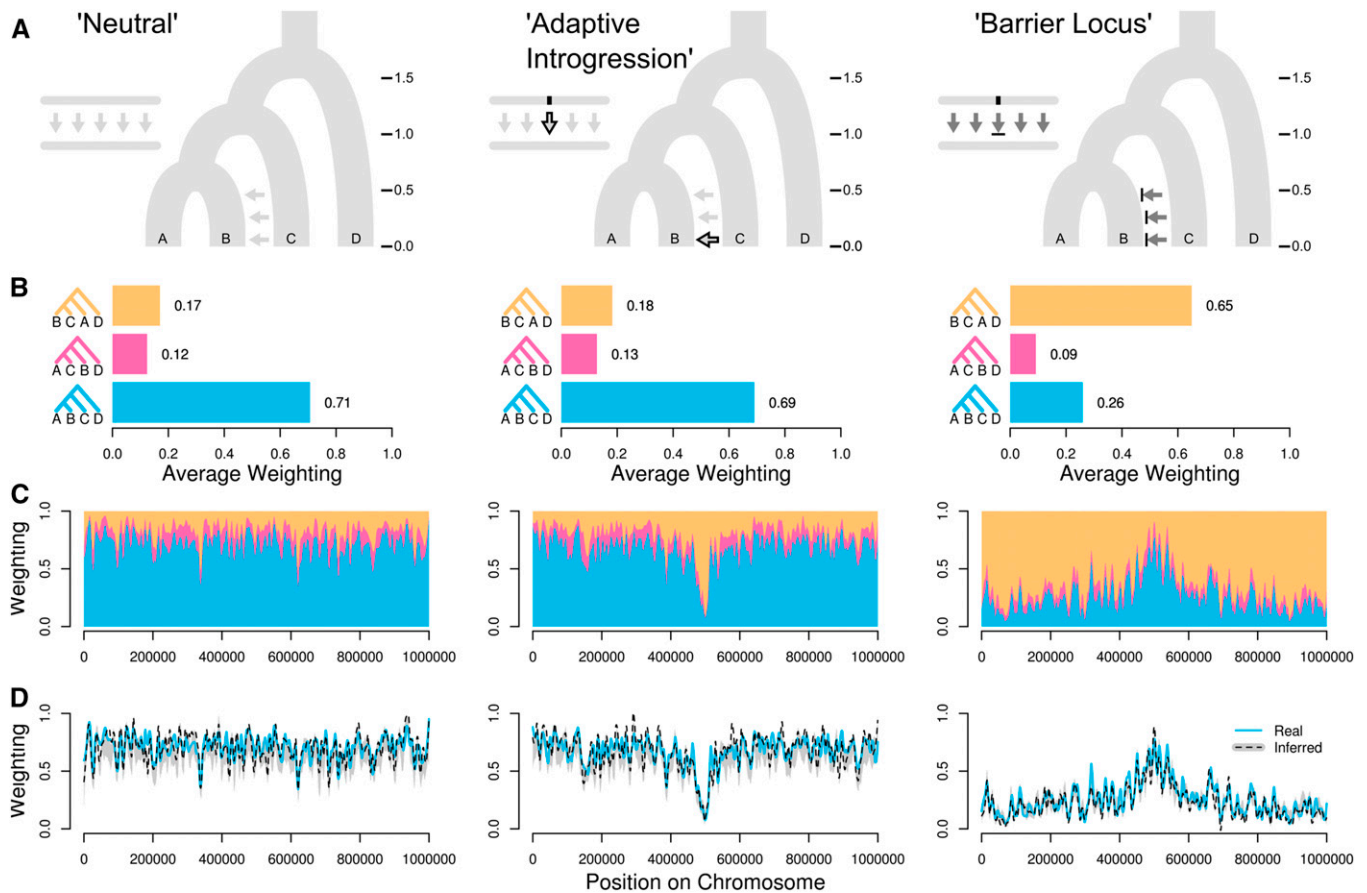


Figure 2 Tests on simulated chromosomes. (A) In all three demographic scenarios, populations split in the order $\{[(A,B),C],D\}$, at the split times indicated (in units of $4N$ generations), with migration from C to B (indicated by arrows). In the Neutral scenario, there is no selection and moderate migration. The Adaptive Introgression scenario is similar, except a beneficial allele at a locus in the center of the chromosome is allowed to move from population C into B at time 0.1. In the Barrier Locus scenario, the rate of migration is high, but an allele at the central locus that is fixed in C is selected against in population B. (B) Mean weightings for the three possible taxon topologies across the 1 Mb simulated chromosome. Note that we illustrate the topologies for the four taxa as rooted, with D as the outgroup for simplicity, but the rooting is not considered when computing the weightings. (C) Weightings for all three topologies plotted (stacked) across the chromosome, with loess smoothing (span = 20 kb). (D) Weightings for topology $\{[(A,B),C],D\}$ inferred from simulated sequence data using nonoverlapping 50 SNP windows and neighbor joining. Solid blue lines indicate the true values, and dashed black lines indicate the inferred values. Gray shading indicates the lower (5%) and upper (95%) quantiles based on 100 bootstrap replicates. Values are smoothed as above.

C to B was five (in units of $4Nm$, as above), and a dominant allele at the central locus that is fixed in C is selected against in population B. The same selection coefficient and population size as above were used.

We simulated sequences from the simulated genealogies using *seq-gen* (Rambaut and Grass 1997). Command options are provided in File S1. The branch scaling factor for mutation was 0.01. Since branches in the simulated genealogies are in units of $4N$ generations, taking N to be 100,000 gives μ (the per generation mutation rate) of 2.5×10^{-8} .

Inferring trees in sliding windows

We tested different methods for inferring trees in windows across the chromosome. The simplest approach used non-overlapping windows of a fixed number of SNPs. A range of window sizes were tested. Trees were then inferred for each

window using PhyML version 3.0 (Guindon *et al.* 2010), implementing either the BIONJ neighbor-joining algorithm (Gascuel 1997) or maximum likelihood optimization of the topology and branch lengths. To investigate the consistency of the tree inference and compute confidence thresholds for the weightings, we performed bootstrapping by randomly resampling the SNPs in each window with replacement and repeating the tree inference. We also tested an approach to infer likely window breakpoints from the data. Taking the topology weightings computed from 10 SNP windows, we used the R package *GenWin* (Beissinger *et al.* 2015) to fit a β -spline to the data and find likely inflection points, which we then used as window breakpoints and inferred a new set of trees for these. In addition, we tested the program *Saguaro*, which infers both the breakpoints and the distance matrix describing each region. Distance matrices were converted to trees using BIONJ, as above.

Power analyses

An important aspect of our approach is its dependence upon reliable trees, which may be inferred from relatively short sequence windows. To investigate the power available to infer topology weightings from short sequences, we simulated datasets under a range of sampling strategies and demographic scenarios, and then compared the true weightings to those computed using trees inferred from the simulated sequences.

Eight sampling strategies were compared, including four, five, six, or 10 sequences from either four or five populations (Figure S3 in File S2). For each sampling strategy, two different demographic scenarios were simulated. In the four-population scenarios, the populations split in the order [(1,2),(3,4)], with the basal split time ($t1$) at either 0.5 or $1 \times 4N$ generations in the past, and the splits between populations one and two and three and four both occurring at $0.1 \times 4N$ generations in the past ($t2$) (Figure S3 in File S2). In the five-population cases, the populations split in the order {[1,(2),(3,4)],5}. As above, $t1$ occurs at either 0.5 or $1 \times 4N$ generations in the past. The next split, between populations one and two and populations three and four, occurs at $0.2 \times 4N$ generations in the past ($t1b$), and the final two splits between populations one and two, and three and four both occur at $t2$ (Figure S3 in File S2).

In each run, we used *msms* (Ewing and Hermisson 2010) to simulate 500 genealogies for the given sampling design and demographic scenario. For each genealogy, we computed the topology weightings using *Twisst*, and then generated a simulated set of sequences using *seq-gen* (Rambaut and Grass 1997). The sequences were then truncated at different lengths to compare tree inference using 10, 25, 50, 100, 200, or 400 SNPs. Trees were inferred in Phyml (Guindon *et al.* 2010), using either BIONJ or with maximum likelihood optimization of the topology and branch lengths. Weightings were then computed from the inferred tree, and compared to the set of true weightings using a scaled Euclidean distance of

$$d = \frac{\sqrt{\sum_{i=1}^n \frac{(w_i - \hat{x}_i)^2}{m_i^2}}}{\sqrt{(n-1)}}, \quad (3)$$

where n is the number of weighting values (*i.e.*, the number of taxon topologies), w_i is the true weighting for topology i , \hat{x}_i is the inferred weighting for topology i . m_i is the absolute value of the maximum possible distance from w_i . This value therefore gives a distance between the true and inferred sets of weightings on a scale of 0–1, with 0 indicating identical values (*i.e.*, perfect inference) and 1 indicating a maximum possible discrepancy between the true and inferred values.

Not all SNPs are phylogenetically informative, and even those that are (those that are not singletons), are not necessarily informative about the relationships among the broader taxa, which is of primary interest for topology weighting. We therefore also tested the power of inference using a subclass of taxon informative sites (TISs), which we define as having at least two alleles present in at least two taxa. As above,

simulated sequences were truncated to contain the number of TISs.

Analysis of real genomic data

We tested *Twisst* on two published genomic datasets from *Neurospora* spp. (ascomycete fungi) and *Heliconius* spp. (butterflies), selected to represent different sampling strategies (four and five taxa, respectively), as well as different levels of evolutionary complexity. The *Neurospora* dataset (Corcoran *et al.* 2016) consisted of 22 aligned haploid genome sequences from *Neurospora tetrasperma* samples (10 of mating type A and 12 of mating type a), along with single genomes representing two related species: *Neurospora crassa* and *Neurospora hispaniola*. Whole genome alignments were obtained from <http://data-dryad.org/resource/doi:10.5061/dryad.162mh>. We used Lineage-10 (UK) samples of *N. tetrasperma*, as these had been shown to carry a strong signal of introgression from *N. hispaniola* (Corcoran *et al.* 2016). Trees were constructed for sliding windows of 50 SNPs using BIONJ as described above, with the requirement that each sample had to be genotyped at ≥ 40 of the 50 SNPs per window. Topology weightings were computed using *Twisst*, with four defined taxa: *N. tetrasperma mat a* (12 sequences), *N. tetrasperma mat A* (10 sequences), *N. crassa* (one sequence), and *N. hispaniola* (one sequence).

The *Heliconius* dataset consisted of 18 resequenced genomes (or 36 haploid genomes) from Martin *et al.* (2013). These samples comprised five populations: two geographically isolated races of *Heliconius melpomene*, from Panama (*H. m. rosina*, $n = 4$) and Peru (*H. m. amaryllis*, $n = 4$), and their respective sympatric relatives *Heliconius cydno chioneus* from Panama ($n = 4$) and *Heliconius timareta thelxinoe* from Peru ($n = 4$), with which they are known to hybridize; along with two additional samples of the more distant silvanifrom clade to serve as outgroups. We limited our analysis to two chromosomes: 18, which carries the gene *optix*, known to be associated with red wing pattern variation; and 21, the Z sex chromosome, which has been shown to experience reduced gene flow between these species, probably due to genetic incompatibilities (Martin *et al.* 2013). Fastq reads were downloaded from the European Nucleotide Archive, study accession no. ERP002440. Reads were mapped to the *H. melpomene* reference genome version 2 (Davey *et al.* 2016) using BWA-mem (Li and Durbin 2009; Li 2013), with default parameters. Genotyping was performed using the Genome Analysis Toolkit (DePristo *et al.* 2011) version 3 HaplotypeCaller and GenotypeGVCFs, with default parameters except that heterozygosity was set to 0.02. Phasing and imputation was performed using Beagle version 4 (Browning and Browning 2007). Trees were inferred as described above, and weightings were computed using *Twisst*, with the five taxa described above.

Data availability

All genotype files, window trees, and weightings for both simulated and real datasets are available from DataDryad at <http://datadryad.org/resource/10.5061/dryad.4jm83>. Scripts used to implement the *Twisst* method are available at <http://github.com/simonhmartin/twisst>.

Results

Analysis of simulated chromosomes

Topology weighting provides an informative summary of the genealogical data and highlights differences between the simulated scenarios (Figure 2). As described above, there are three possible unrooted topologies for the four taxa. In the Neutral scenario, the most prevalent topology, $\{(A,B),C\},D\}$, which reflects the population split times, has an average weighting of 71% across the chromosome. The other two topologies are both fairly rare, but $\{(B,C),A\},D\}$ is more common on average (17%) than $\{(A,C),B\},D\}$ (12%). This is because the former can result from both gene flow and incomplete lineage sorting (ILS), whereas the latter can only result from ILS, as there was no simulated migration between A and C or between B and D. In the Adaptive Introgression scenario, the weightings are very similar to the Neutral scenario on average, but in the center of the chromosome there is a strong excess of the topology $\{(B,C),A\},D\}$, created by the spread of a beneficial allele from population C into B. Finally, in the Barrier Locus scenario, high migration from C to B causes a swamping by the topology $\{(B,C),A\},D\}$, which has an average weighting of 65%. However, there is a broad peak at the center of the chromosome where the population branching topology $\{(A,B),C\},D\}$ had not been eroded, due to selection limiting introgression.

In the corresponding simulations with five taxa, there are 15 possible taxon topologies (Figure S7 in File S2). There is greater topological variation overall, as there are more ways that incomplete sorting can occur. Nonetheless, topology weights clearly detect the differences among the scenarios, highlighting the most abundant topologies as well as the location of the selected locus (Figure S7 in File S2).

Inferring weightings from simulated sequence data

Above, we computed the weightings directly from the simulated genealogies, but we are also able to show that topology weightings can be reliably estimated when the genealogies are inferred from simulated sequence data (Figure 2D and Figure S7D in File S2). Because neither the genealogies nor the recombination breakpoints at which genealogies switch are known, we tested several approaches for inferring genealogies for narrow intervals across the chromosome. First, we performed extensive power analyses, covering a range of demographic scenarios and sampling designs, to explore the relationship between the number of SNPs used for tree inference and the accuracy of topology weighting. Across the range of scenarios investigated, we find a consistent lower bound of 50 SNPs to achieve >90% accuracy (Figure S4, Figure S5, and Figure S6 in File S2). Focusing specifically on TISs (see above) makes no discernible difference, probably because most SNPs in our simulations are taxon informative. These tests also indicate that neighbor-joining trees provide more accurate weightings than maximum likelihood

trees, in addition to much faster computation (Figure S4, Figure S5, and Figure S6 in File S2).

We then analyzed trees inferred for nonoverlapping windows across our simulated recombining chromosomes. A fixed window size of 50 SNPs gives results that most closely approximate the true weightings (Figure 2D and Figure S7D in File S2). In agreement with our power analyses, with <50 SNPs the estimates are less accurate and tend to underestimate the weighting of the most prevalent topology (Figure S8 and Figure S9 in File S2). Weightings tending toward intermediate values are expected as the underlying trees become less well resolved. Interestingly, windows of ≥ 100 SNPs also result in reduced accuracy, but with a tendency to overestimate support for the most prevalent topology and underestimate support for others (Figure S8 and Figure S9 in File S2). This can be explained by the fact that large windows are forced to average over regions of distinct ancestry, therefore favoring the most widespread signal. To confirm this hypothesis, we repeated our neutral simulation using a 10-fold lower population recombination rate. In this new dataset, 100 SNP windows give the most accurate weightings, and even 200 SNP windows have high accuracy, while 50 SNP windows perform only marginally less well (Figure S10 and Figure S11 in File S2).

We tested whether bootstrapping over the SNPs in each window can be used to validate the accuracy of the observed weightings. Bootstrap weights tend to be similar but marginally more conservative, underestimating the weight of the most prevalent topology (Figure 2D). This is because the bootstrap trees tend to be slightly less well resolved, leading to more intermediate weightings. Bootstrapping is therefore a useful means to test the strength of support for an observed peak in the weighting of a particular topology. However, being inherently conservative, bootstrapping would not be able to determine whether an observed intermediate weighting was accurate or simply the result of a poorly resolved tree.

Because real recombination breakpoints are not evenly spaced, we also tested two approaches in which the window boundaries are inferred from the data itself. In our first approach, we used the R package *GenWin* (Beissinger *et al.* 2015) to fit a smooth spline to the weightings from 10 SNP windows and identify likely window boundaries as inflection points, and then inferred trees for the newly defined window regions. The resulting topology weightings match the true weightings fairly well, but not as well as for the fixed 50 SNP windows (Figure S12 and Figure S13 in File S2). As above, this appears to be due to poor tree inference in the smallest windows. The second approach used the method *Saguaro* (Zamani *et al.* 2013), which combines a hidden Markov model and a self-organizing map to infer both the trees and window boundaries. This approach poorly recapitulates the true weightings, greatly overestimating support for the most prevalent topology (Figure S12 and Figure S13 in File S2). We therefore used fixed windows of 50 SNPs for all further analyses.

Branch lengths differ among topology types

Topology weighting is primarily a descriptive method, but the weightings do carry information that can aid inferences about population history. The simulated Barrier Locus scenario (Figure 2) provides an interesting test case. Due to the overwhelming signal of introgression, it would be difficult to know which topology corresponds to the true population branching order (*i.e.*, the species tree) if this was not known. The topology $\{[(B,C),A],D\}$ is prevalent across much of the chromosome, but $\{[(A,B),C],D\}$ is prevalent around the chromosome center. It has been proposed that the original population branching order can be identified by considering branch lengths (Fontaine *et al.* 2015; Gante *et al.* 2016). Taxa that cluster together due to recent introgression tend to be separated by short branches, whereas those that cluster according to the population branching order should have deeper splits. Indeed, in trees inferred from 50 SNP windows, pairwise branch distances between the taxa suggest that subtrees matching $\{[(B,C),A],D\}$ tend to result from recent introgression between B and C (Figure S14 in File S2), thus implying that $\{[(A,B),C],D\}$ is the more likely population branching order.

Analysis of real genomic data

The *Neurospora* dataset consists of four taxa (three possible topologies) and is the simpler of the two real datasets analyzed (Figure 3, A and B). It was selected to test how well *Twist* is able to detect the signal of a previously described adaptive introgression event from *N. hispaniola* into *N. tetrasperma* individuals of the A mating type (Corcoran *et al.* 2016). This introgression covers the entire (~7 Mb) nonrecombining region of linkage group I (LGI). Indeed, we find a dramatic shift in the pattern of topology weightings in the central part of LGI (Figure 3C). The species-tree topology (topo1), which groups the two *N. tetrasperma* mating types as closest relatives, is prevalent across most of the genome but has very little weighting in the central part of LGI. Instead, it is replaced by topo3, which groups mating type A individuals of *N. tetrasperma* with *N. hispaniola*. Elsewhere, topo3 has limited weighting, nearly identical to that of topo2, and consistent with a low level of ILS throughout the genome. However, a region of linkage group IV also shows a weak shift in support toward topo3, potentially reflecting a separate introgression signal involving a small number of sequences.

The *Heliconius* dataset represents a more complex, five-taxon test case. The five taxa include an outgroup and two pairs of sympatric, nonsister taxa, between which gene flow is known to occur (Figure 4A). Of the 15 possible topologies (Figure 4B), the two most common across these chromosomes are topo3 and topo6. topo3 is consistent with the accepted species branching order, in which the allopatric *H. c. chioneus* and *H. t. thelxinoe* are sister taxa; whereas topo6 groups populations by geography, consistent with interspecific gene flow in both Panama and Peru. The former is by far

the most prevalent throughout the Z chromosome (Figure 4C). By contrast, the species topology has variable weighting across chromosome 18, and is outweighed in places by topologies consistent with gene flow (topo4, topo5, topo6, topo11, and topo14). In particular, there is a strong peak in the region of *optix* for topo11, which groups the taxa by wing pattern, and is consistent with the previously described adaptive introgression of the red-band allele between *H. m. amaryllis* and *H. t. thelxinoe* in Peru (Pardo-Diaz *et al.* 2012; The Heliconius Genome Consortium 2012). Zooming in on this peak shows a clear block of ~150 kb over which the introgression topology is weighted highly (Figure S15 in File S2). This block includes the regulatory region downstream of *optix* that is known to control wing pattern variation in these species (Baxter *et al.* 2010; Wallbank *et al.* 2016). Another four topologies that partially match the species branching order (topo1, topo2, topo10, and topo15) have moderate weightings throughout, whereas topologies consistent with neither the species tree nor gene flow (topo7, topo8, topo9, topo12, and topo13) have low weightings, especially across the Z chromosome, implying less ILS than on chromosome 18.

Discussion

Most statistics used in population genetics describe aspects of the underlying genealogy. For example, F_{ST} can be expressed as the relative rate of coalescence within subpopulations compared to the total population (Slatkin and Voelm 1991). Similarly, the *D* statistic of the ABBA-BABA test compares the relative rate of coalescence between two pairs of nonsister populations (Green *et al.* 2010; Durand *et al.* 2011). Topology weighting can be seen as a generalization of this principle, as it determines the relative frequency of all possible patterns of coalescence among samples from a set of defined taxa. Unlike the ABBA-BABA test, which is based on binary trees (samples either share the same allele or not) and therefore only four taxa, topology weighting uses a full genealogy and can, in principle, be applied to any number of taxa, each represented by any number of sequences. In practice, however, beyond six taxa, the number of possible taxon topologies becomes very large, making topology weighting less practical. Nevertheless, even when the number of taxon topologies is very large, there may be value in comparing the weightings of particular topologies that support specific hypotheses (Van Belleghem *et al.* 2017). Unlike other methods for comparing tree topologies (*e.g.*, Robinson and Foulds 1981; Estabrook *et al.* 1985), topology weighting reduces the problem to quantifying relationships among but not within defined taxa. There is also no real limitation on the number of samples included per taxon. Although computation of exact weightings may become infeasible for large trees, it remains possible to compute approximate weightings with small margins of error fairly rapidly. Our computational approach, *Twist*, is based on a simple counting procedure but we are confident that more efficient analytical solutions, or at least approximations, will be found.

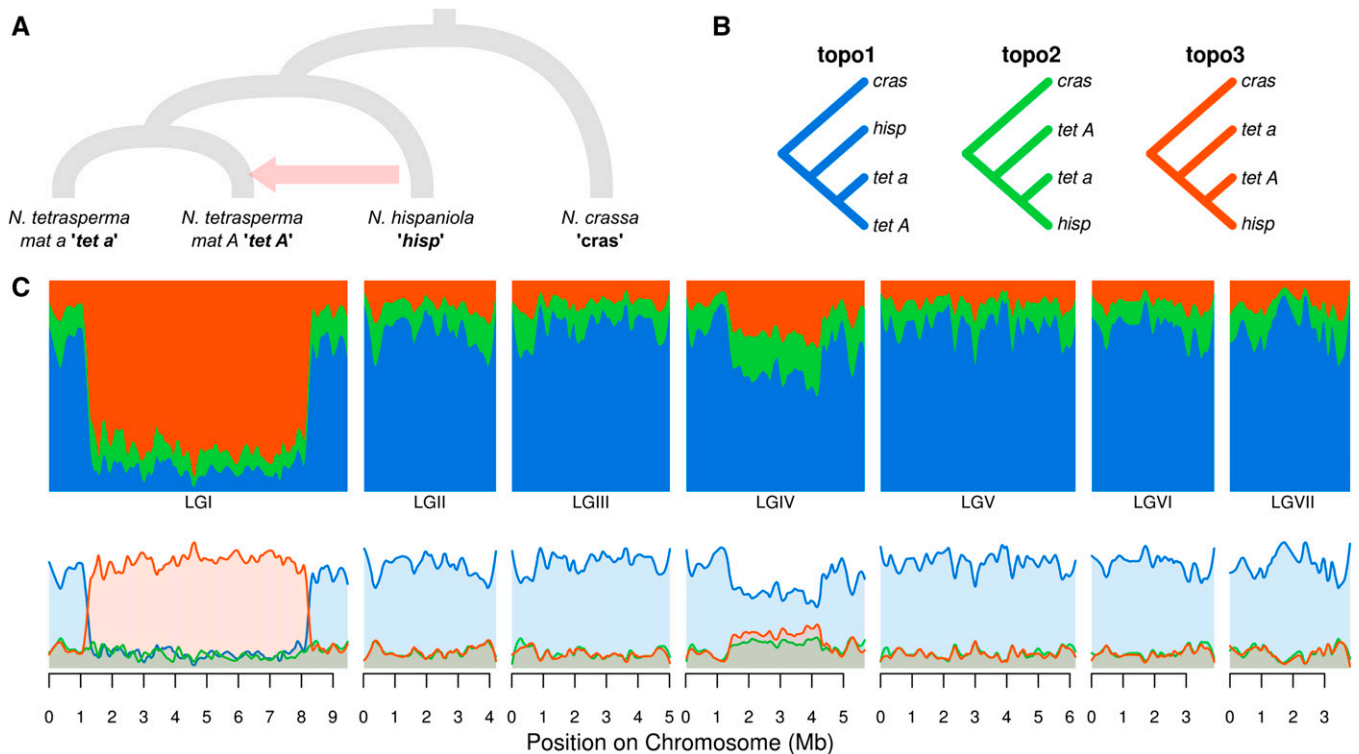


Figure 3 *Neurospora* analysis. (A) The putative species tree. Note that mating type *a* and *A* individuals of *N. tetrasperma* are shown as separate branches, while in reality, apart from the nonrecombining region of LGI, these samples represent a single recombining population. The putative introgression from *N. hispaniola* into *N. tetrasperma* mat *A* individuals (Corcoran *et al.* 2016) is indicated by an arrow. (B) The three possible taxon topologies for these four taxa. (C) Topology weightings for 50 SNP windows plotted across all seven linkage groups, with loess smoothing (span = 500 kb). The top and bottom plots show the same data, plotted as stacked or as separate lines, respectively.

An important consideration when applying this method is its dependence on the trees used. In most cases, the true genealogy for each distinct ancestry block is not known and must be inferred from the sequences. Accurate inference requires multiple informative SNPs. Our tests on simulated data highlight a central difficulty when analyzing recombining chromosomes: a trade-off between signal and resolution. Using larger numbers of SNPs increases our ability to infer the correct tree, but may average over genomic regions with different histories. This leads to a systematic overestimation of the weightings for more abundant topologies, whose signal tends to swamp out that of others. Using few SNPs per window allows for better resolution but can lead to inaccuracies in tree reconstruction from insufficient signal (*i.e.*, phylogenetic error), producing star-like trees and intermediate weightings for all topologies. Fortunately, this means that errors in tree inference are unlikely to result in spurious peaks in the weighting of a single topology, but instead may lead to underestimation of the height of a particular peak. A peak can be validated using bootstrapping, but an even better test is to demonstrate that it persists with decreasing window sizes. Our simulations, based on realistic recombination and mutation rates, indicate that a window size of 50 SNPs provides a good compromise between signal and resolution across a range of demographic scenarios, although larger windows may be acceptable if the population recombination rate is

known to be low. While in some cases with high recombination rates there may be too few mutations per recombination to accurately infer variation in genealogies across the genome, we expect that many cases will fall within a feasible range. Importantly, not all recombinations are relevant for topology weighting: only recombination events between lineages from distinct taxa (*i.e.*, effective recombinations or intertaxon recombinations) can alter taxon relationships, and hence alter the weightings. The ability to infer the patterns of topology weighting across the genome therefore depends on the relationship between the mutation rate and the rate of intertaxon recombination. Where possible, simulations tailored to the taxa being studied can be used to guide the choice of window size. In the future, improved methods to infer breakpoints from the data may further resolve this difficulty.

Another challenge is that diploid resequencing data should ideally be phased so that each tip in the tree represents a distinct haplotype. Phasing can be performed using probabilistic approaches informed by patterns of linkage disequilibrium (Browning and Browning 2007; Delaneau *et al.* 2013). Although such methods may be error-prone across large genomic distances, they have fairly high accuracy at short ranges (Bukowicki *et al.* 2016), making them suited to the narrow windows used for topology weighting. Moreover, the genomic regions involving intertaxon recombinations (*i.e.*, those relevant for topology weighting) are more likely to be

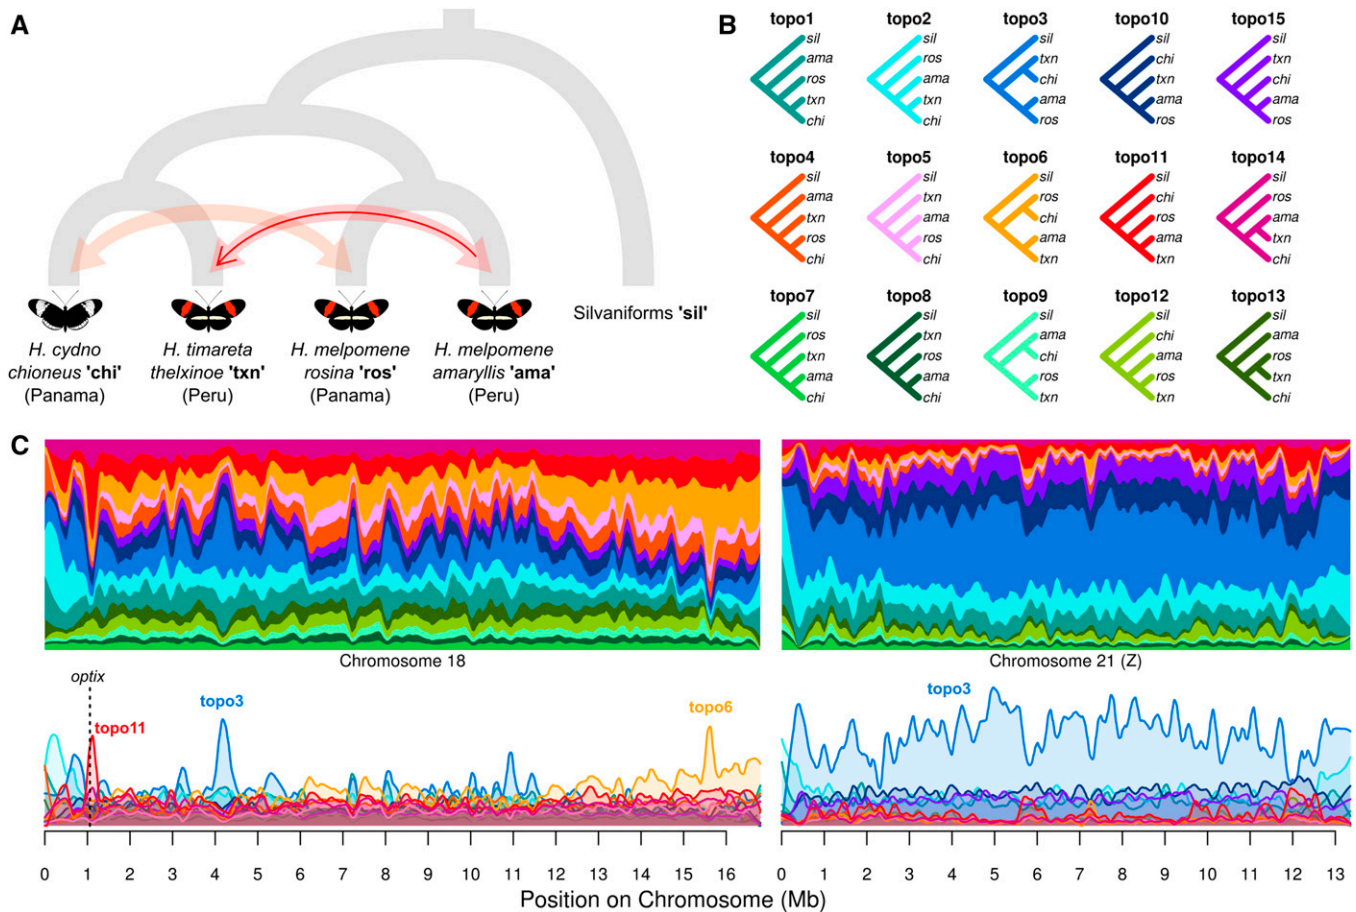


Figure 4 *Heliconius* analysis. (A) The putative species tree. Shaded arrows indicate ongoing gene flow between sympatric, nonister taxa in Panama and Peru, respectively (Martin *et al.* 2013). The solid red arrow indicates the putative adaptive introgression of the red wing-patterning allele near the gene *optix* (Pardo-Diaz *et al.* 2012; The *Heliconius* Genome Consortium 2012). (B) The 15 possible taxon topologies for these five taxa. (C) Topology weightings for 50 SNP windows plotted across chromosomes 18 and 21 (Z), with loess smoothing (span = 500 kb). The top and bottom plots show the same data, plotted as stacked or as separate lines, respectively. The location of *optix* on chromosome 18 is indicated by a dashed vertical line.

phased correctly because they tend to involve more divergent sequences.

Topology weighting is principally a descriptive method and can be applied with no prior knowledge of the studied samples, apart from some basis on which to define distinct groups, such as geography or phenotype. By capturing the tree-like nature of sequence evolution, it provides information that is not provided by descriptive statistics like F_{ST} (Figure S16, or clustering methods such as *Structure* (Pritchard *et al.* 2000). In addition to describing the taxon branching order, tree-based methods allow incorporation of additional parameters like a nucleotide substitution model and different rates of evolution in different parts of the tree. Unlike conventional phylogenomic methods, topology weighting captures information about fine-scale and quantitative variation across the genome. This power and resolution is highlighted in the *Heliconius* example studied here, where topologies supporting admixture are common across chromosome 18, but there is one narrow peak consistent with the adaptive introgression of a wing patterning allele near the gene *optix*, as described previously (Pardo-Diaz *et al.* 2012; The *Heliconius* Genome

Consortium 2012). We note that topology weighting simply describes the signal in the data, and does not explicitly test for introgression over other causes of discordant phylogenetic signal. For example, the elevated frequency of topologies consistent with introgression across *Heliconius* chromosome 18 compared to the Z chromosome is only partially due to an elevated rate of gene flow on autosomes (Martin *et al.* 2013), but also reflects increased ILS due to their larger effective population size relative to the sex chromosome. Weightings may be more difficult to interpret in cases with a less well understood evolutionary history. Nevertheless, it may be possible for example to differentiate between topologies representing the species tree and those reflecting recent introgression by comparing their branch lengths (which can be output by *Twisst*) (Fontaine *et al.* 2015; Gante *et al.* 2016). Finally, we have found that topology weighting provides a means to identify candidate loci underlying trait variation, based on clustering of taxa by phenotype [see also Van Belleghem *et al.* (2017), for a more extensive demonstration of this power]. In summary, topology weighting is a simple but versatile exploratory tool that is applicable to a diverse range of questions and datasets.

Acknowledgments

We thank Pasi Rastas, Markus Möst, and Joe Hanly for useful comments on the method. S.H.M. was funded by a research fellowship from St John's College, University of Cambridge. S.M.V.B. was funded by NSF (National Science Foundation) grant no. DEB-1257839. Additional funding was provided by European Research Council grant no. 339873 held by Chris Jiggins.

Literature Cited

- Baxter, S. W., N. J. Nadeau, L. S. Maroja, P. Wilkinson, B. A. Counterman *et al.*, 2010 Genomic hotspots for adaptation: the population genetics of müllerian mimicry in the *Heliconius melpomene* clade. *PLoS Genet.* 6: 12.
- Beissinger, T. M., G. J. Rosa, S. M. Kaeppler, D. Gianola, and N. de Leon, 2015 Defining window-boundaries for genomic analyses using smoothing spline techniques. *Genet. Sel. Evol.* 47: 30.
- Browning, S. R., and B. L. Browning, 2007 Rapid and accurate haplotype phasing and missing-data inference for whole-genome association studies by use of localized haplotype clustering. *Am. J. Hum. Genet.* 81: 1084–1097.
- Bukowicki, M., S. U. Franssen, and C. Schlötterer, 2016 High rates of phasing errors in highly polymorphic species with low levels of linkage disequilibrium. *Mol. Ecol. Resour.* 16: 874–882.
- Corcoran, P., J. L. Anderson, D. J. Jacobson, Y. Sun, P. Ni, M. Lascoux *et al.* 2016 Introgression maintains the genetic integrity of the sex-determining chromosome of the fungus *Neurospora tetrasperma*. *Genome Res.* 26: 486–498.
- Davey, J. W., M. Chouteau, S. L. Barker, L. Maroja, S. W. Baxter *et al.* 2016 Major improvements to the *Heliconius melpomene* genome assembly used to confirm 10 chromosome fusion events in 6 million years of butterfly evolution. *G3 (Bethesda)* 6: 695–708.
- Delaneau, O., J.-F. Zagury, and J. Marchini, 2013 Improved whole-chromosome phasing for disease and population genetic studies. *Nat. Methods* 10: 5–6.
- DePristo, M. A., E. Banks, R. Poplin, K. V. Garimella, J. R. Maguire *et al.*, 2011 A framework for variation discovery and genotyping using next-generation DNA sequencing data. *Nat. Genet.* 43: 491–498.
- Durand, E. Y., N. Patterson, D. Reich, and M. Slatkin, 2011 Testing for ancient admixture between closely related populations. *Mol. Biol. Evol.* 28: 2239–2252.
- Dutheil, J. Y., G. Ganapathy, A. Hobolth, T. Mailund, M. K. Uyenoyama *et al.*, 2009 Ancestral population genomics: the coalescent hidden Markov model approach. *Genetics* 183: 259–274.
- Estabrook, G. F., F. R. McMorris, and C. A. Meacham, 1985 Comparison of undirected phylogenetic trees based on subtrees of four evolutionary units. *Syst. Biol.* 34: 193–200.
- Ewing, G., and J. Hermisson, 2010 MSMS: a coalescent simulation program including recombination, demographic structure and selection at a single locus. *Bioinformatics* 26: 2064–2065.
- Fontaine, M. C., J. B. Pease, A. Steele, R. M. Waterhouse, D. E. Neafsey *et al.*, 2015 Mosquito genomics. Extensive introgression in a malaria vector species complex revealed by phylogenomics. *Science* 347: 1258524.
- Gante, H. F., M. Matschiner, M. Malmstrøm, K. S. Jakobsen, S. Jentoft *et al.*, 2016 Genomics of speciation and introgression in princess cichlid fishes from Lake Tanganyika. *Mol. Ecol.* 25: 6143–6161.
- Gascuel, O., 1997 BIONJ: an improved version of the NJ algorithm based on a simple model of sequence data. *Mol. Biol. Evol.* 14: 685–695.
- Green, R. E., J. Krause, A. W. Briggs, T. Maricic, U. Stenzel *et al.*, 2010 A draft sequence of the Neandertal genome. *Science* 328: 710–722.
- Guindon, S., J. -F. Dufayard, V. Lefort, M. Anisimova, W. Hordijk *et al.*, 2010 New algorithms and methods to estimate maximum-likelihood phylogenies: assessing the performance of PhyML 3.0. *Syst. Biol.* 59: 307–321.
- Heled, J., and A. J. Drummond, 2010 Bayesian inference of species trees from multilocus data. *Mol. Biol. Evol.* 27: 570–580.
- Hobolth, A., O. F. Christensen, T. Mailund, M. H. Schierup, W. Enard *et al.*, 2007 Genomic relationships and speciation times of human, chimpanzee, and gorilla inferred from a coalescent hidden Markov model. *PLoS Genet.* 3: e7.
- Huerta-Cepas, J., F. Serra, and P. Bork, 2016 ETE 3: reconstruction, analysis and visualization of phylogenomic data. *Mol. Biol. Evol.* 33: 1635–1638.
- Li, H., 2013 Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM. *arXiv* 1303.3997v2.
- Li, H., and R. Durbin, 2009 Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics* 25: 1754–1760.
- Lohse, K., N. H. Barton, G. Melika, and G. N. Stone, 2012 A likelihood-based comparison of population histories in a parasitoid guild. *Mol. Ecol.* 21: 4605–4617.
- Maddison, W., 1997 Gene trees in species trees. *Syst. Biol.* 46: 523–536.
- Martin, S. H., K. K. Dasmahapatra, N. J. Nadeau, C. Salazar, J. R. Walters *et al.*, 2013 Genome-wide evidence for speciation with gene flow in *Heliconius* butterflies. *Genome Res.* 23: 1817–1828.
- Martin, S. H., J. W. Davey, and C. D. Jiggins, 2015 Evaluating the use of ABBA-BABA statistics to locate introgressed loci. *Mol. Biol. Evol.* 32: 244–257.
- Pardo-Diaz, C., C. Salazar, S. W. Baxter, C. Merot, W. Figueiredo-Ready *et al.*, 2012 Adaptive introgression across species boundaries in *Heliconius* butterflies. *PLoS Genet.* 8: e1002752.
- Patterson, N., P. Moorjani, Y. Luo, S. Mallick, N. Rohland *et al.*, 2012 Ancient admixture in human history. *Genetics* 192: 1065–1093.
- Pritchard, J. K., M. Stephens, and P. Donnelly, 2000 Inference of population structure using multilocus genotype data. *Genetics* 155: 945–959.
- Rambaut, A., and N. C. Grass, 1997 Seq-Gen: an application for the Monte Carlo simulation of DNA sequence evolution along phylogenetic trees. *Bioinformatics* 13: 235–238.
- Reich, D., K. Thangaraj, N. Patterson, A. L. Price, and L. Singh, 2009 Reconstructing Indian population history. *Nature* 461: 489–494.
- Reich, D., N. Patterson, D. Campbell, A. Tandon, S. Mazieres *et al.*, 2012 Reconstructing native American population history. *Nature* 488: 370–374.
- Robinson, D. F., and L. R. Foulds, 1981 Comparison of phylogenetic trees. *Math. Biosci.* 53: 131–147.
- Slatkin, M., and L. Voelm, 1991 FST in a hierarchical island model. *Genetics* 127: 627–629.
- The *Heliconius* Genome Consortium, 2012 Butterfly genome reveals promiscuous exchange of mimicry adaptations among species. *Nature* 487: 94–98.
- Van Belleghem, S.M., P. Rastas, A. Papanicolaou, S.H. Martin, C. Arias *et al.* 2017 Complex modular architecture around a simple toolkit of wing pattern genes. *Nat. Ecol. Evol.* 1: 0052.
- Wallbank, R. W. R., S. W. Baxter, C. Pardo-Diaz, J. J. Hanly, S. H. Martin *et al.*, 2016 Evolutionary novelty in a butterfly wing pattern through enhancer shuffling. *PLoS Biol.* 14: 1–16.
- Zamani, N., P. Russell, H. Lantz, M. P. Hoepfner, J. R. Meadows *et al.*, 2013 Unsupervised genome-wide recognition of local relationship patterns. *BMC Genomics* 14: 347.

Communicating editor: J. D. Wall