

RESEARCH ARTICLE

Inference of Ancestral Recombination Graphs through Topological Data Analysis

Pablo G. Cámara^{1*}, Arnold J. Levine², Raúl Rabadán^{1*}

1 Department of Systems Biology and Department of Biomedical Informatics, Columbia University College of Physicians and Surgeons, New York, New York, United States of America, **2** The Simons Center for Systems Biology, Institute for Advanced Study, Princeton, New Jersey, United States of America

* pg2495@cumc.columbia.edu (PGC); rr2579@cumc.columbia.edu (RR)



OPEN ACCESS

Citation: Cámara PG, Levine AJ, Rabadán R (2016) Inference of Ancestral Recombination Graphs through Topological Data Analysis. *PLoS Comput Biol* 12(8): e1005071. doi:10.1371/journal.pcbi.1005071

Editor: Sergei L. Kosakovsky Pond, Temple University, UNITED STATES

Received: January 10, 2016

Accepted: July 20, 2016

Published: August 17, 2016

Copyright: © 2016 G. Cámara et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper and its Supporting Information files.

Funding: The work of the authors is supported by National Institutes of Health (NIH) grants U54 CA193313 and NIGMS R01-GM117591 (<http://www.cancer.gov/> and <https://www.nigms.nih.gov/Pages/default.aspx>). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

Abstract

The recent explosion of genomic data has underscored the need for interpretable and comprehensive analyses that can capture complex phylogenetic relationships within and across species. Recombination, reassortment and horizontal gene transfer constitute examples of pervasive biological phenomena that cannot be captured by tree-like representations. Starting from hundreds of genomes, we are interested in the reconstruction of potential evolutionary histories leading to the observed data. Ancestral recombination graphs represent potential histories that explicitly accommodate recombination and mutation events across orthologous genomes. However, they are computationally costly to reconstruct, usually being infeasible for more than few tens of genomes. Recently, Topological Data Analysis (TDA) methods have been proposed as robust and scalable methods that can capture the genetic scale and frequency of recombination. We build upon previous TDA developments for detecting and quantifying recombination, and present a novel framework that can be applied to hundreds of genomes and can be interpreted in terms of minimal histories of mutation and recombination events, quantifying the scales and identifying the genomic locations of recombinations. We implement this framework in a software package, called TARGeT, and apply it to several examples, including small migration between different populations, human recombination, and horizontal evolution in finches inhabiting the Galápagos Islands.

Author Summary

Evolution occurs through different mechanisms, including point mutations, gene duplication, horizontal gene transfer, and recombinations. Some of these mechanisms cannot be captured by tree graphs. We present a framework, based on the mathematical tools of computational topology, that can explicitly accommodate both recombination and mutation events across the evolutionary history of a sample of genomic sequences. This approach generates a new type of summary graph and algebraic structures that provide quantitative information on the evolutionary scale and frequency of recombination events. The accompanying software, TARGeT, is applied to several examples, including migration

between sexually-reproducing populations, human recombination, and recombination in Darwin's finches.

This is a *PLOS Computational Biology* Methods paper.

Introduction

Since the publication of the first draft of the human genome [1, 2], there has been an explosion in genomic data. The genomes of thousands of different human individuals have been sequenced [3], several hundreds of eukaryotic genomes have been characterized, and new viral, bacterial and archaeal species are being sequenced on an almost daily basis [4, 5]. Darwin provided a historical dimension to the taxonomical enterprise, proposing that closely related species in the hierarchical taxonomy share ancestors. Since then, tree-like structures have been proposed to represent the evolutionary/historical relationship between organisms. In the last few years, however, the richer and more comprehensive genomic characterization of many organisms have underscored the need of representations that are not strictly tree-like. Phenomena such as horizontal gene transfer in bacteria [6], the ability of viruses to borrow and lend genes across species, and hybridization in metazoa (in plants, in particular [7, 8]) are exposing some of the limitations imposed by tree-like phylogenetic structures. The definition of species itself becomes cumbersome in bacteria and viruses [9]. Within many species, including humans, genetic recombination is so pervasive that tree-like representations are useless. It is then natural to wonder what other frameworks could be used to capture phylogenetic relationships without losing the interpretability and simplicity of trees [10–12]. Of particular interest are representations that reduce to trees when evolution is tree-like; that capture genetic relations between ancestors, and identify genomic regions originating from different ancestral lineages; and, more generally, that allow for an interpretation of the observed data in terms of a chronological sequence of events.

Several such frameworks have been proposed in the last two decades. The study of phylogenetic networks has been an area particularly active [13–15]. Phylogenetic networks provide representations that extend trees to graphs (networks), generating loops when the data does not fit into a tree. Some of those methods can easily be applied to more than one hundred genomes [16–21] providing the opportunity for large-scale representations. However, the biological interpretation of these representations is limited, as loops represent inconsistencies with trees, but it is unclear how these inconsistencies arose historically, what genomic regions were involved, or how frequently an exchange happened. Other types of representations, sometimes named *explicit networks* [13, 22], do aim to provide a historical account in terms of a chronology of events. Ancestral recombination graphs (ARGs) provide potential explanations of the observed data in terms of a progression of recombination and mutation events. As in trees, mutations are represented as events along the branches. Recombinations, however, appear as the fusion of two parental branches into one offspring branch. ARGs provide simple histories that can be used in association mapping [23–25], SNP genotyping [26] or inference of the frequency and scale of recombination [27]. However, these applications are hindered by the computational infeasibility of constructing ARGs that explain hundreds of sequences. The construction of minimal ARGs, containing the minimum number of recombination events

required to explain the sample in absence of convergent evolution and back-mutation, is an NP-hard problem [28–30]. Several approximations have been developed in the last few years, including galled trees [31, 32], branch and bound [33], heuristic [23] and sequentially Markov coalescent approaches [34].

Recently, a new framework to study genomic relationships has been proposed [35–37], based on topological data analysis [38–40]. Topology is the area of mathematics that aims to characterize properties of spaces up to continuous deformations, for instance the number of disconnected components, loops and holes of a space. TDA extends the concepts and tools of topology to finite metric spaces, that is, finite sets of points and distances between them. Taking the premise that a set of points has been sampled from an unknown underlying space, TDA attempts to infer the topological features of the space (Fig 1A). Stability results [35, 41, 42] guarantee that small fluctuations in the data only create small changes in the inferred topological features, providing robust characterizations of the data.

In a TDA framework, genomes are characterized by points in a high dimensional space where pairwise distances are genetic distances between sequences. Assuming that each genomic site mutates at most once across the evolutionary history of the sample, the genetic distance between two genomes can only increase with the acquisition of novel mutations. The only way of “closing” a loop (a close path) in this space is therefore by means of a recombination event [35]. Hence, an approach to studying recombination in the sample of genetic sequences is to study the loops that those sequences generate when represented in the above way.

A valuable attribute of TDA methods is that they are informative about the scale or size of the inferred topological features. Given a finite set of data points, there is an infinite number of spaces that are compatible with the points. TDA structures this spectrum of possibilities by introducing a notion of scale (Fig 1B): at a given scale ϵ , two points are connected in the underlying space if their distance is smaller than ϵ . Topological features compatible with the data can be then summarized in terms of sets of intervals, named *barcodes* [43] (Fig 1C). Each interval in a barcode represents the range of scales across which a particular topological feature (e.g. a loop) is present in the inferred topological space. In the genomic context introduced above, barcodes of loops summarize the frequency and scale (mutational distance between recombining sequences) of recombination events, and provide a basic structure on which statistics of genomic exchange can be built [37].

TDA methods are particularly well suited for large datasets. In the context of molecular phylogenetics and evolution, they have been applied to the study of viral recombination and reassortment [35], bacterial species [36] and point estimators in population genetics [37]. However, these implementations of TDA have limitations, as they are not tailored for the biological problem they try to address. Specifically, traditional TDA methods only use information about genetic distances between sequences, and so they discard the full structure of segregating characters, missing numerous recombination events that are required to explain the data. Relatedly, it is unclear which specific evolutionary histories explaining the data TDA informs about, and what is the precise relation between barcodes and these histories.

Here we address these two important aspects, improving on the scalable capabilities of TDA to extract robust information on the possible evolutionary histories of a sample of genetic sequences. In particular, we show that by systematically sampling subsets of segregating sites and performing TDA, we are able to identify most of the necessary recombination events identified by bound methods [33, 44, 45], providing a significant improvement of past methods [35–37] in terms of interpretation and sensitivity. Moreover, we introduce a novel type of graph (*topological ARG* or *tARG*), closely related to minimal ARGs, that captures ensembles of minimal recombination histories; and we show that TDA informs about the topological features and genetic scales of these graphs. Like minimal ARGs [22, 23], tARGs can be considered

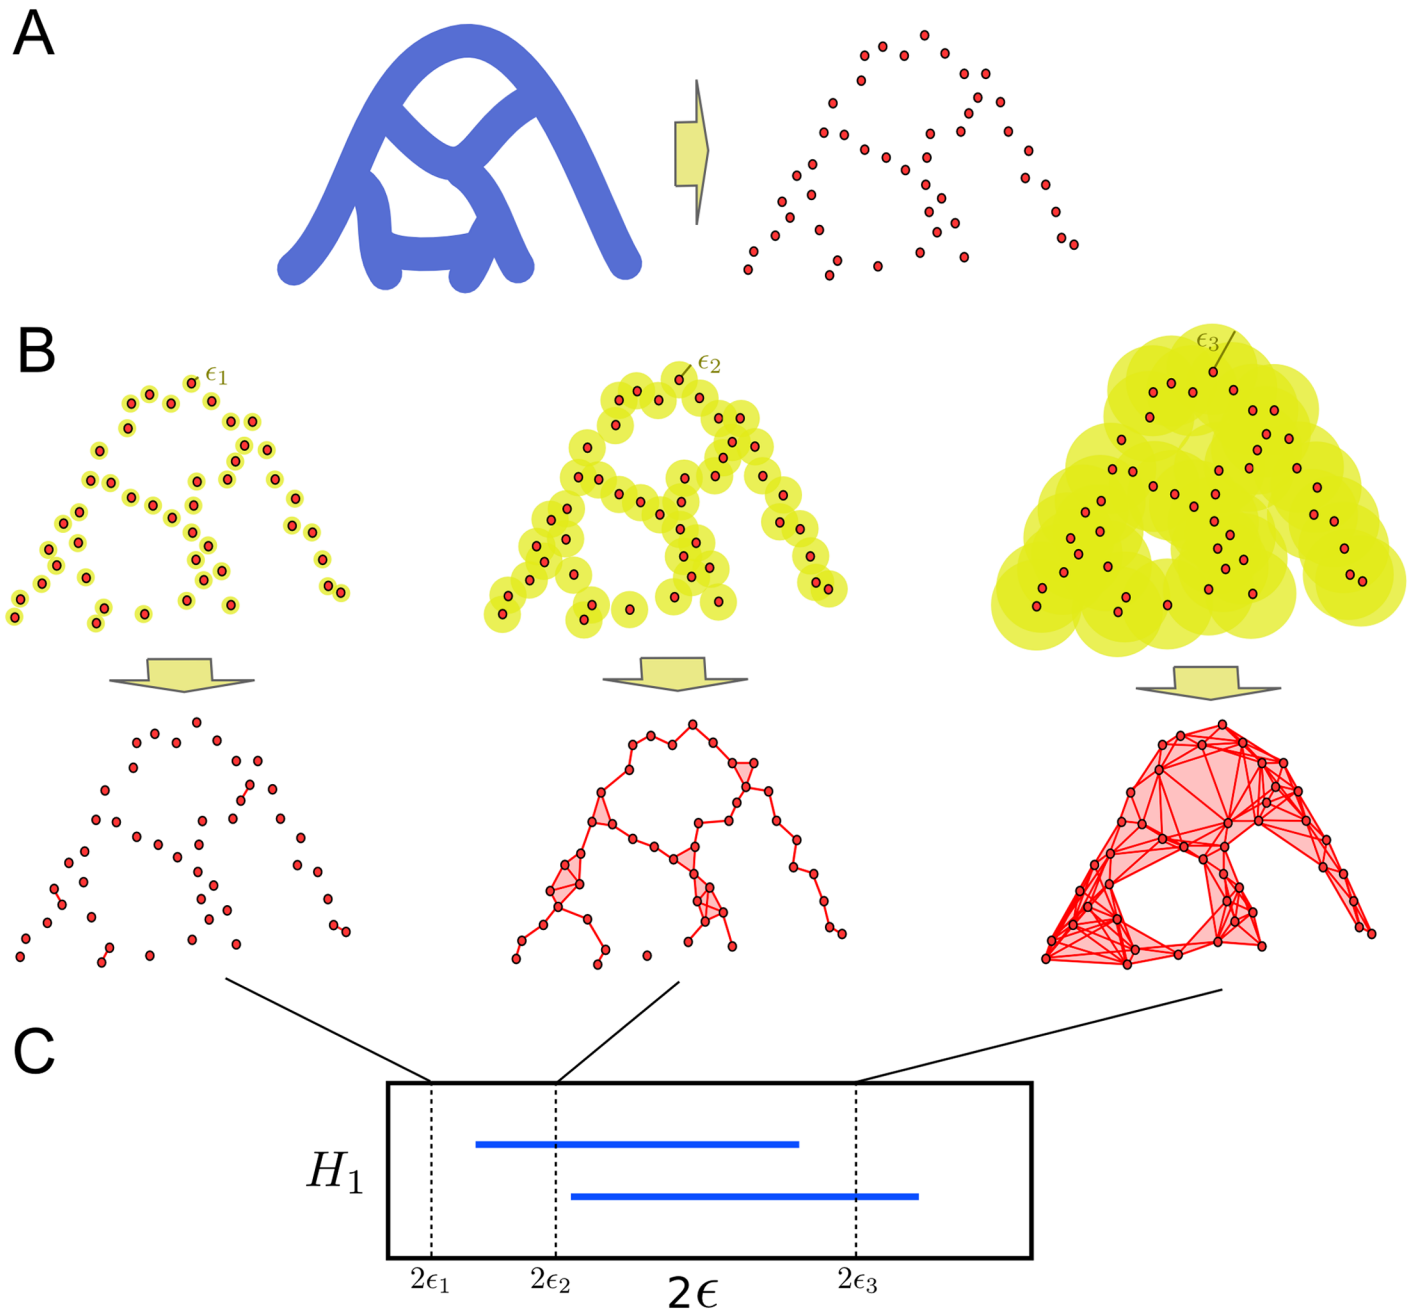


Fig 1. Topological data analysis. (A) Topological data analysis aims to infer the topological features (e.g. loops, voids, etc.) of an unknown space from a finite set of sampled points. (B) Persistent homology, a tool of TDA, builds simplicial complexes (generalizations of networks that include higher dimensional elements like triangles and tetrahedra), by taking balls of radius ϵ centred on the sampled points. Points are connected in the simplicial complex if the corresponding balls intersect. This construction is known as Vietoris-Rips complex. Persistent homology tracks how the topological features of Vietoris-Rips complexes change with ϵ . (C) Barcodes are suitable representations of persistent homology. Each interval in the barcode represents the range of ϵ across which a particular topological feature (for instance, a loop) is present in the inferred topology. In this figure, the barcode of the first persistent homology, that tracks the presence of loops, is shown. The two intervals in the barcode correspond to the two loops present in the original space.

doi:10.1371/journal.pcbi.1005071.g001

as explicit, parsimonious, interpretable phylogenetic representations. The main advantage of tARGs and barcodes versus minimal ARGs is, however, the possibility of obtaining such phylogenetic information in polynomial time, which allows us to deal with hundreds of sequences. We have implemented this method in a software, called TARGet, and have illustrated it with several examples, including small migration between diverging populations, human recombination, and horizontal evolution of finches inhabiting the Galápagos archipelago. The software, instructions and example files used in the manuscript can be obtained from <https://github.com/RabadanLab/TARGet>.

Results

Topological ARGs

An ARG is an explicit phylogenetic network representing a possible evolutionary history of a sample of genetic sequences, where only mutation and recombination events are present and convergent evolution is not considered and so never occurs [22, 46, 47]. ARGs are very useful constructs in population genetics and phylogenetics. However, the problem of building a minimal ARG from a set of genetic sequences is known to be NP-hard [28–30]. The use of ARGs has therefore been traditionally limited to small samples, consisting of a handful of sequences.

In this section, we introduce a particular class of minimal ARGs and a set of related graphs. Then, using computational algebraic topology, in the next section we show that it is possible to extract, in polynomial time, phylogenetic information from this class of minimal ARGs, without having to explicitly construct them. Thus, by restricting to this specific class of graphs, we are able to extend the realm of ARGs to large samples of sequences.

To be specific, we consider a sample \mathcal{S} consisting of n distinct genetic sequences with m binary segregating characters. The latter can be single nucleotide polymorphisms (SNPs), indels, gene duplications or any other genetic trait that takes one of two possible states, 0 or 1, in each sequence. An ARG is then formally defined as a directed acyclic graph \mathcal{N} with n leaf nodes and a unique root node, where every node other than the root has in-degree one (tree node) or two (recombination node), every segregating character labels a unique edge in \mathcal{N} (infinite sites assumption), and every sequence in \mathcal{S} labels a unique leaf in \mathcal{N} . Moreover, each node in \mathcal{N} is labelled by a m -length binary sequence, such that the sequence labelling a tree node differs from the sequence of the parent node only at the character labelling the edge that connects the two nodes; and the sequence labelling a recombination node is a combination of the sequences labelling the two parent nodes. Single-crossover recombinant sequences are formed by taking the first k sites from the sequence of one of the parent nodes (*prefix*) and appending the last $m - k$ sites from the sequence of the other parent node (*suffix*), for $k \in [1, m - 1]$.

There is an infinite number of ARGs that can explain a given sample \mathcal{S} [22]. A stochastic model, such as the coalescent model with recombination [46, 48], would assign probabilities to each possible ARG. Here, however, we adopt a parsimony approach and consider ARGs that are minimal (in a sense defined below), without assuming an underlying probabilistic model. Such a model-independent approach has proven useful in summarizing genetic sequences into evolutionary histories where all events are required.

Specifically, we consider ARGs that contain exactly the minimum number R_{\min} of single-crossover recombinations required to explain the sample, and that minimize the function

$$D(\mathcal{N}) = \sum_{r=0}^{R_{\min}} d_r \tag{1}$$

where the sum runs over all recombination events in \mathcal{N} , and d_r is the Hamming distance

between the two parental sequences involved in the r -th recombination. This is a more restricted definition of minimal ARG than the one that usually appears in population genetics literature [22], where the condition on $D(\mathcal{N})$ is generally not required. We use the term *ultra-minimal ARG* to refer to this restricted type of minimal ARG. Ultra-minimal ARGs are thus minimal ARGs where recombination events involve parental sequences that are as genetically close as possible. They introduce a higher level of parsimony than minimal ARGs, being informative not only about the minimum number of recombination events, but also about the minimum genetic distance between the recombining sequences that took part in those events. By construction, an ultra-minimal ARG explaining any given sample always exists. Examples are shown in Figs 2 and 3.

A minimal ARG can be condensed by collapsing all unlabelled edges, so that the resulting graph can be embedded into an m -dimensional hypercube and its diagonals (that is, the line segments joining non-consecutive vertices) (Fig 2). The number of edges and vertices of such a condensed representation is $m + 2R_{\min}$ and $m + R_{\min} + 1$, respectively, whereas the number of independent loops is R_{\min} , where a loop is said to be independent if it cannot be embedded in the union of other loops. In this representation, the distance between two nodes is defined as the number of edges in the shortest path connecting the nodes, and is equal to the Hamming distance between the corresponding sequences.

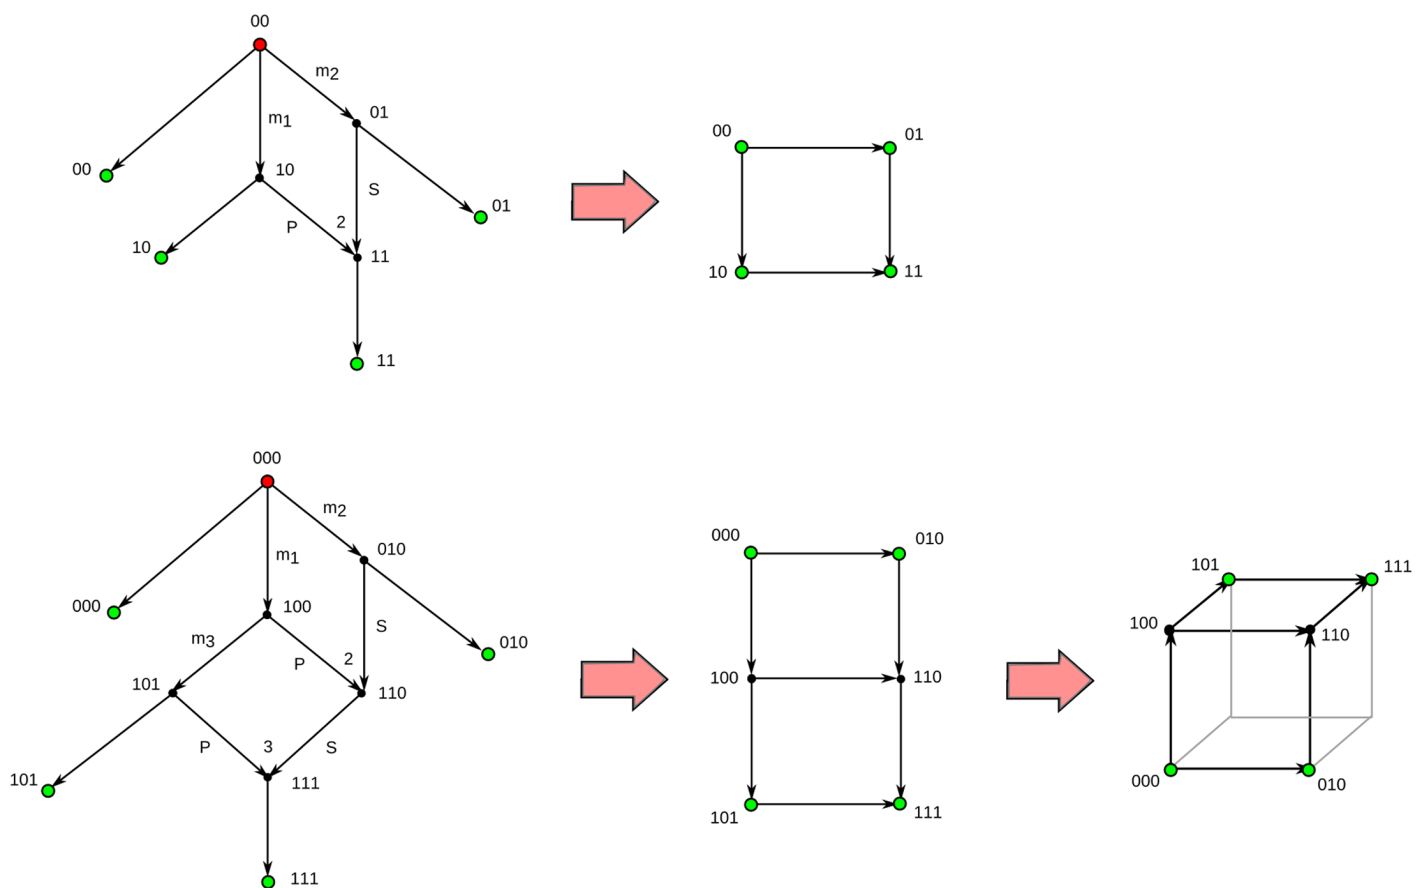


Fig 2. ARGs and condensed graphs. Two examples of ultra-minimal ARGs and the condensed graphs that result from collapsing their unlabelled edges. The root node is marked red whereas sampled nodes are marked green. Mutations in the r -th character are indicated by m_r . Edges pointing to a recombination node are labelled with the letter P or S, depending on whether they contribute to the prefix or suffix of the recombinant sequence. Recombinant nodes are marked with the position of the recombination breakpoint. All nodes are labelled by their sequence of characters. Condensed graphs of ARGs can be embedded into m -dimensional hypercubes and their diagonals.

doi:10.1371/journal.pcbi.1005071.g002

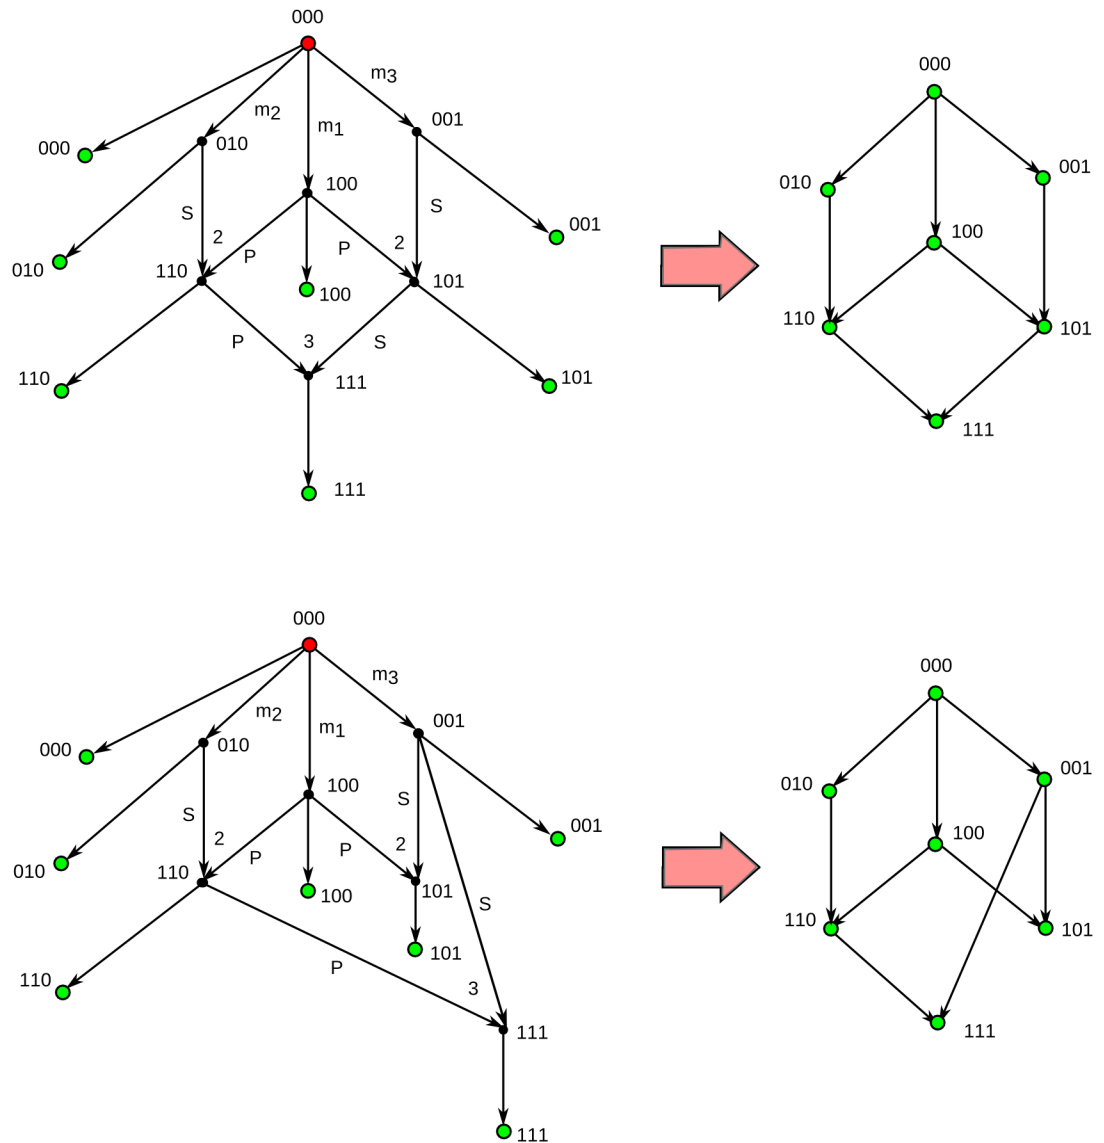


Fig 3. Ultra-minimal ARGs. Two examples of ARGs containing the minimum number of recombination events, $R_{\min} = 3$, required to explain a sample of $n = 7$ sequences with $m = 3$ segregating sites. Both ARGs are minimal ARGs. However, only the minimal ARG at the bottom is an ultra-minimal ARG.

doi:10.1371/journal.pcbi.1005071.g003

Given a sample S of genetic sequences, we would like to obtain information about the ultra-minimal ARGs that explain S , without explicitly constructing them. To that end, we consider the undirected graph $\mathcal{G} = (V, E)$, with vertices V and edges $E = E_1 \cup \dots \cup E_b$, that results from the union of all condensed ultra-minimal ARGs $\mathcal{G}_i = (V, E_i)$ explaining S and having the same set of vertices V (Fig 4). We call this construction *topological ARG* (tARG). A tARG therefore summarizes the collection of most parsimonious histories associated to a sample of genetic sequences. However, unlike minimal ARGs, tARGs are completely determined by their vertices. By considering tARGs instead of minimal ARGs, we are able to reduce an NP-hard problem into a much simpler (but still very informative) topological problem, as we describe in next section.

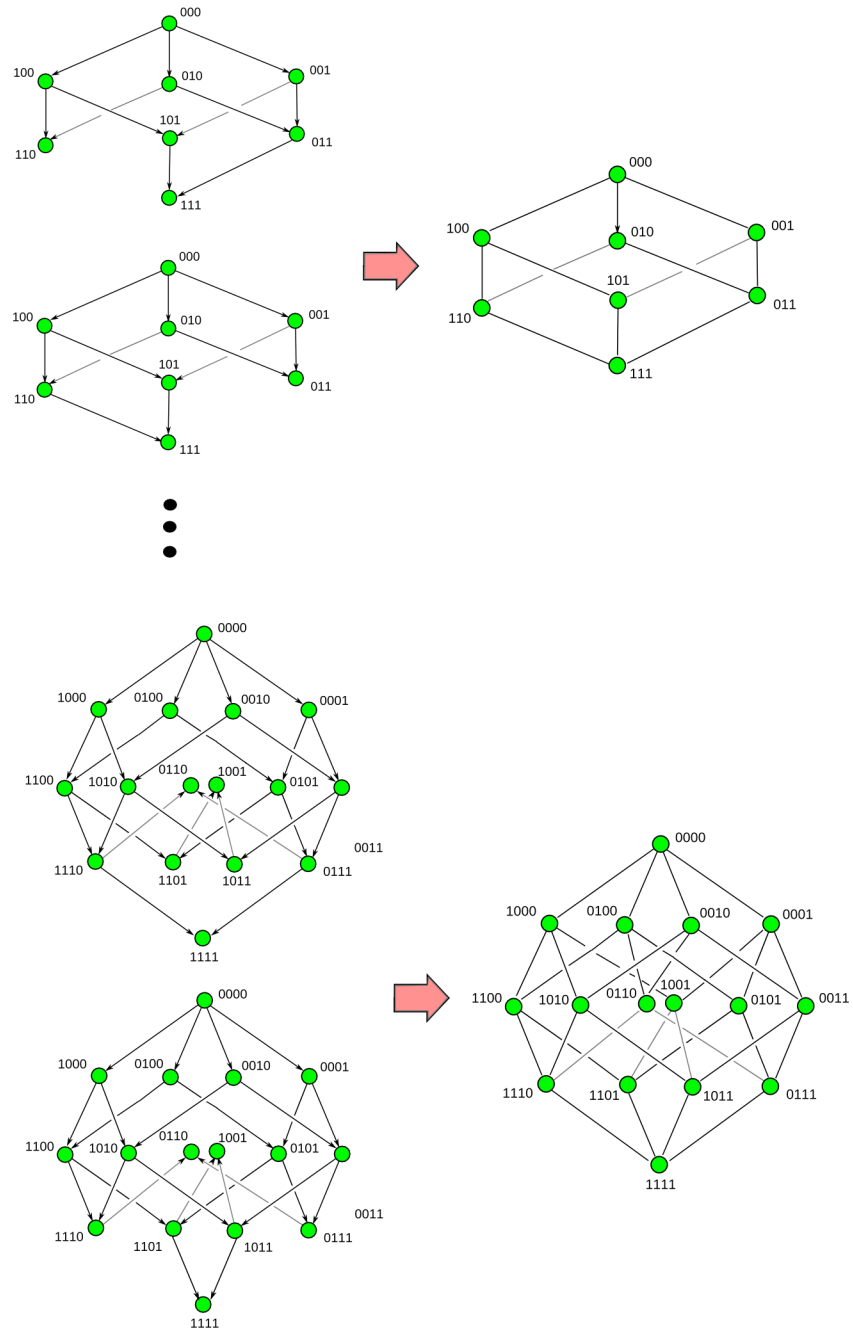


Fig 4. Topological ARGs. Examples of condensed ultra-minimal ARGs (left) and their corresponding tARGs (right). In a tARG the edges are completely determined by the vertices. The topology of the resulting tARG can differ from that of the original condensed ultra-minimal ARGs.

doi:10.1371/journal.pcbi.1005071.g004

Persistent homology and recombination inference

Topological data analysis has emerged during the last decade as a branch of applied topology that attempts to infer topological features of spaces (such as the number of loops and holes) from sets of sampled points [38]. The topological features of a space are preserved under continuous deformations of the space and can be arranged in mathematical structures called

homology groups [49]. We refer the reader to refs. [49, 50] for formal definitions and basic introductions to algebraic topology. In brief, the n^{th} homology group of a space is an algebraic structure that encompasses all $(n + 1)$ -dimensional holes of the space. Of special interest to us is the first homology group, whose elements correspond to loops.

Homology groups can be computed by replacing the original space with a simpler one, known as *simplicial complex*, which has the same topological features as the original space but consists of a finite set of elements (Fig 1B). A simplicial complex is a generalization of a network that, in addition to nodes and vertices, includes higher dimensional elements like triangles and tetrahedra. Simplicial complexes are powerful because they allow the implementation of algebraic operations to extract the topological features of the space.

When only a finite set of points of the space is given, there is still a well-defined notion of homology groups, known as *persistent homology* [39, 40], which capture the topological features of the underlying space. At each value of a scale parameter ϵ , a simplicial complex (known as Vietoris-Rips complex) can be constructed by considering the intersections of balls of radius ϵ centred at the sampled points (Fig 1B). Points are joined if their corresponding balls intersect. This process produces a sequence of simplicial complexes parametrized by ϵ , from which persistent homology can be computed using available algorithms [39, 40]. Remarkably, the computation time of persistent homology is polynomial in the number of points [39, 40].

Persistent homology can be represented using barcodes [43]. These are graphical representations where each element of persistent homology is represented by a segment spanning the interval $[\epsilon_b, \epsilon_d]$, where ϵ_b and ϵ_d are the values of the parameter ϵ at which the corresponding feature is respectively formed and destroyed in the sequence of simplicial complexes (Fig 1C). Thus, each segment in a barcode represents a topological feature inferred from the data, and the position and length of the segment are informative of the size of the topological feature. The values ϵ_b and ϵ_d are referred as *birth* and *death time* of the topological feature, respectively.

In the current context, we exploit the use of persistent homology to infer topological features of an unknown tARG, given a set of sampled nodes (Fig 5). The use of persistent homology to detect the presence of recombination in genetic samples was proposed in [35]. However, the relation between persistent homology and explicit evolutionary histories incorporating recombination events was not studied. Our aim is inferring information about the loops of the tARG, as they correspond to recombination events present in the collection of most parsimonious histories explaining the sample. To that end, we consider the Hamming distance matrix of the sample and compute persistent homology using the algorithm developed in ref. [39, 40]. Since computing the distance matrix and persistent homology requires respectively $\mathcal{O}(n^2m)$ and $\mathcal{O}(n^3)$ operations [39, 40], the running time grows at most cubically with the number of genetic sequences. An advantage of using persistent homology instead of just counting loops in a nearest neighbour graph is that we also obtain valuable information about the genetic distances between recombining sequences.

The barcode that results from this computation contains information about the number and size of the loops in the tARG underlying the sample (Fig 5). Each segment in the barcode represents a loop in the tARG, and therefore a recombination event in an ultra-minimal ARG explaining the sample. The position of each segment provides information about the genetic scales involved in the corresponding recombination event. Specifically, $2\epsilon_d$ sets an upper bound to the mutational distance between the two recombining sequences, since all pairwise distances between nodes in the loop are smaller than $2\epsilon_d$. The number of segments in the barcode (namely, the dimension of the first persistent homology group) or *persistent first Betti number*, b_1 , is hence a lower bound of the number of recombination events in the tARG, \bar{R}_{\min} . Note that, since a tARG is the union of multiple minimal histories, \bar{R}_{\min} can be larger than

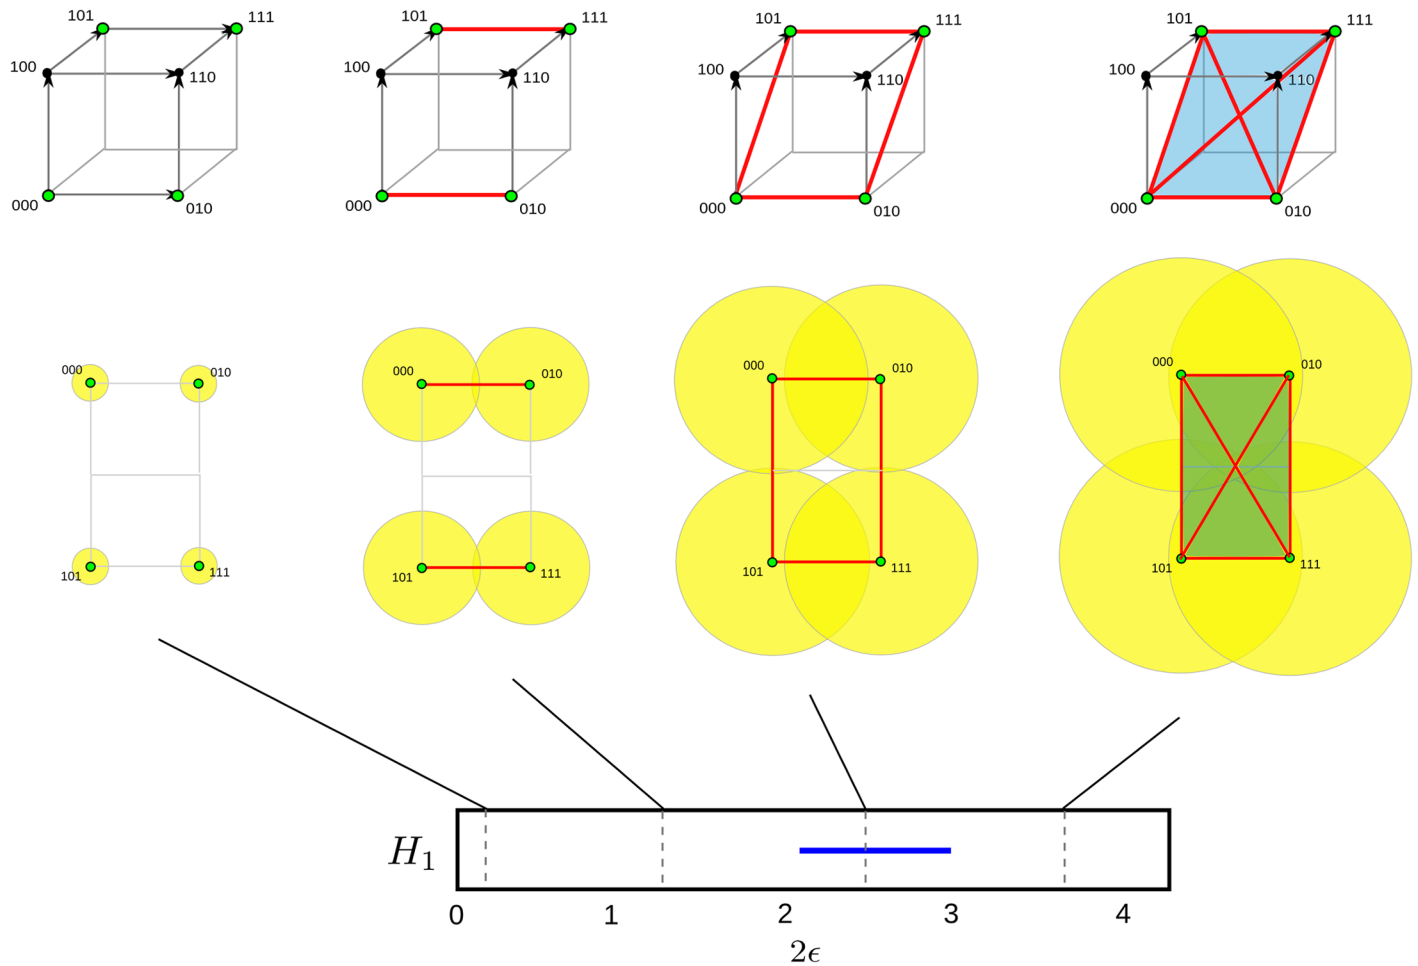


Fig 5. Persistent homology of a sample of genetic sequences. Barcode and Vietoris-Rips complexes at several values of the parameter ϵ , for the sample of sequences $S = \{000, 010, 101, 111\}$. Only the first homology group (H_1) is shown. At small ϵ the four sampled points are disconnected. Increasing ϵ leads to a loop, that appears as a single element of H_1 . Further increasing ϵ fills in the loop, leading to a single connected surface. An ultra-minimal ARG explaining S , and the corresponding tARG are shown in Fig 2 (bottom). The barcode only captures one of the $R_{\min} = 2$ recombination events.

doi:10.1371/journal.pcbi.1005071.g005

R_{\min} . In particular, $\bar{R}_{\min} > R_{\min}$ when there are three characters for which all eight possible allele combinations appear in the sample. In general, this can only happen at very large recombination rates.

The barcode ensemble of a sample

The sensitivity of persistent homology to detect recombination decreases as the number m of segregating characters increases. Indeed, in that case the dimensionality of the ambient space is larger and the sample becomes sparser. For this reason, b_1 is in general a loose lower bound of \bar{R}_{\min} . To address a similar problem, Myers and Griffiths introduced the idea of combining the local bounds that result from partitioning the sequence, building a more stringent global bound [45]. In this way, information about the ordering of characters is incorporated and the location of recombination breakpoints is constrained in the sequence. This general idea was applied in [45] to the haplotype bound, $n - m - 1 \leq R_{\min}$, to build a stronger lower bound of R_{\min} , denoted R_{MG} .

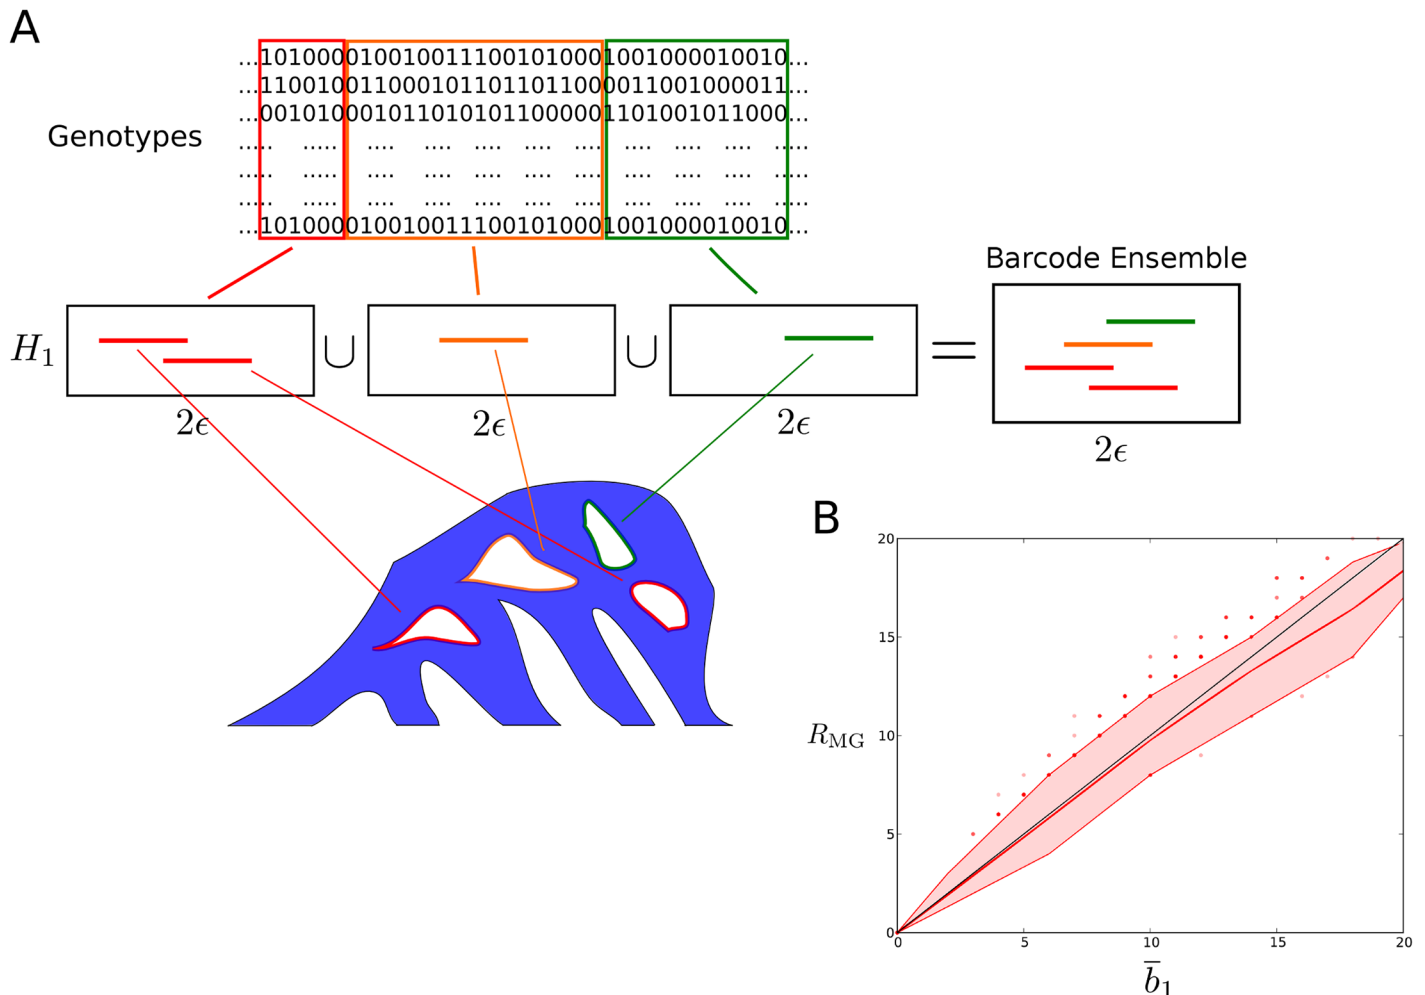


Fig 6. Barcode ensemble of a sample. (A) Schematic representation of the barcode ensemble of a genomic sample. Persistent homology is computed for each genomic interval of a partition of the sequence. Barcodes associated to different genomic intervals capture different recombination events. The union of all barcodes is the barcode ensemble. The total number of intervals in the barcode ensemble is denoted as \bar{b}_1 . The partition is chosen such that \bar{b}_1 is maximized. (B) Comparison between lower bounds $\bar{b}_1 \leq \bar{R}_{\min}$ and $R_{MG} \leq R_{\min}$ in coalescent simulations. Values of \bar{b}_1 and R_{MG} for simulated samples of 40 sequences with 12 segregating sites, sampled from a population under the coalescent model with recombination. 4,000 samples were simulated in total. The colored band represents the interdecile range, whereas the central line represents the mean. The values of \bar{b}_1 and R_{MG} are strongly correlated (Pearson's $r = 0.98, p < 10^{-100}$). At very high recombination rates, \bar{b}_1 tends to be larger than R_{MG} , as cases where $\bar{R}_{\min} > R_{\min}$ occur more frequently.

doi:10.1371/journal.pcbi.1005071.g006

A similar idea can be applied in the context of barcodes to build a *barcode ensemble*, given by the disjoint union of the persistent first-homology barcodes of a set of optimally chosen, non-overlapping intervals within the sequence alignment (Fig 6A). Given a partition of a genetic sequence, the barcode associated to each interval captures information about recombination events with breakpoint in that interval. Due to the curse of dimensionality mentioned in the previous paragraph, the union of the barcodes associated to two contiguous genomic intervals often captures more recombination events than the barcode associated to the union of the two genomic intervals. Therefore, by systematically exploring all possible partitions of the genetic sequence, it is possible to find a partition that maximizes the total number of bars in the barcodes. The solution is often not unique, as different partitions may lead to the same total number of bars. One may reduce this degeneration by considering additional criteria, such as also maximizing the total length of the bars (so that they are more informative about

genetic distances). The formal details of the barcode ensemble construction are presented in the Methods section.

The barcode ensemble incorporates information about the full structure of characters in the sample, largely increasing the sensitivity of persistent homology to recombination and providing information on the location of the recombination breakpoints in the sequence. The number of bars in the barcode ensemble, \bar{b}_1 , is an improved lower bound of \bar{R}_{\min} , in the same way as R_{MG} is an improved lower bound of R_{\min} :

$$\begin{array}{ccc}
 \text{tARG} & \longrightarrow & \bar{R}_{\min} \geq \bar{b}_1 \geq b_1 \\
 \uparrow & & \\
 \text{(ultra) minimal ARG} & \longrightarrow & R_{\min} \geq R_{MG} \geq n - m - 1
 \end{array}$$

In biological data, \bar{b}_1 and R_{MG} are in general very close to each other (Fig 6B), as tARGs with $\bar{R}_{\min} > R_{\min}$ occur very rarely. However, unlike R_{MG} , barcode ensembles provide additional phylogenetic information, such as bounds on the mutational distances between recombining sequences (note that birth and death times in barcode ensembles refer to local genetic distances, namely mutational distances across the genomic interval associated to the particular bar). These features put barcode ensembles at the very interesting interface between the fast, but phylogenetically limited, existing lower bounds to R_{\min} ; and the slow, but phylogenetically rich methods for reconstructing minimal ARGs. We have implemented the computation of barcode ensembles in publicly available software, called TARGeT.

Examples

We consider five examples that illustrate how the formal developments presented in previous sections can be used to extract useful phylogenetic information from samples of genetic sequences. The first example is a simple toy model where an explicit minimal ARG can be easily constructed. It displays how the information contained in the barcode ensemble of the sample directly maps to features of ultra-minimal ARGs. The second example, based on simulated data of two sexually reproducing populations exchanging genetic material at low rate, shows the applicability of persistent homology to large datasets, consisting of several hundreds of sequences. It also demonstrates the use of phylogenetic information contained in the barcode ensemble to distinguish among various biological settings with similar recombination rates. The third and fourth examples consist respectively of 250 and 100 kilobase regions in the HLA and MS32 loci of ~ 100 humans, where several meiotic recombination hotspots localize. The fifth example consists of a 9 megabase scaffold in the genome of 112 Darwin’s finches [51]. These last three examples serve to illustrate the applicability of barcode ensembles to real datasets.

A simple example. We illustrate the use and interpretation of barcode ensembles with a simple example, consisting of a sample of 4 genetic sequences with 7 binary characters: 1111001, 1111111, 0000110 and 0000000. Minimal ARGs explaining this sample require two single-crossover recombination events. An ultra-minimal ARG is presented in Fig 7A. The most ancestral recombination event involves genetically distant parental gametes, leading to a large loop in the ARG. To the contrary, the most recent recombination event involves genetically close parental gametes, leading to a second small loop in the ARG. These features are captured by the barcode ensemble of the 4 sequences (Fig 7B), which consists of two bars, corresponding to the two recombination events. The position of the bars represent the genetic scales associated to the recombination events, with the $2\epsilon_d = 5$ (3) bar corresponding respectively to the large (small) recombination loop. These death times are good upper bounds for

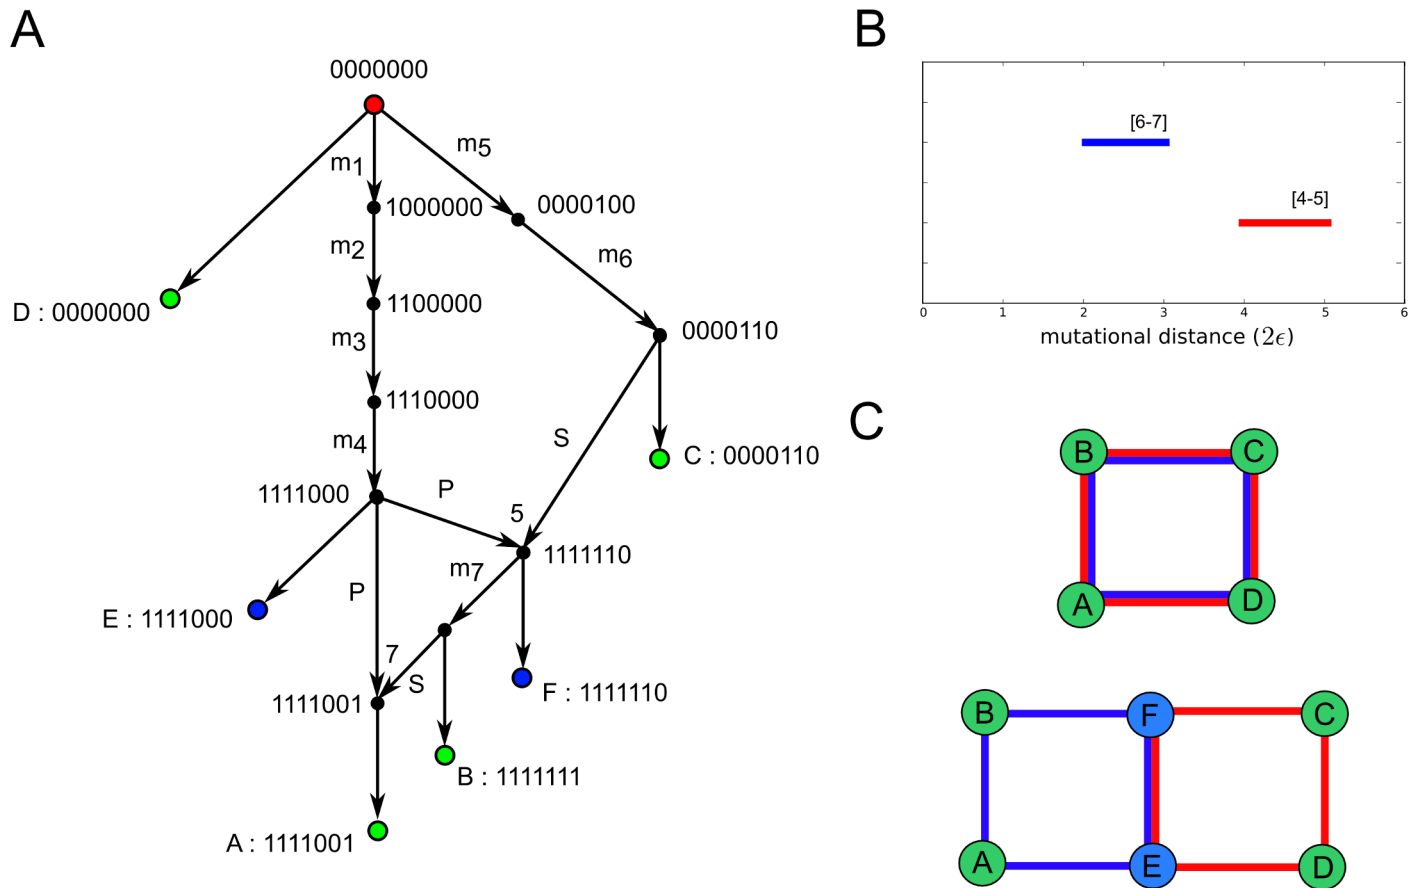


Fig 7. Ultra-minimal ARG, first-homology barcode ensemble and reconstructed tARG of a sample of 4 sequences. The four sampled sequences are represented by green leaf nodes in the ultra-minimal ARG depicted in (A). The ARG involves two single-crossover recombination events. Both recombination events and their genetic scales (mutational distance between recombining sequences) are correctly captured by the barcode ensemble of the samples, shown in (B). Intervals containing the location of recombination breakpoints are indicated over each bar. Persistent homology generators can be used to reconstruct the topology of the tARG, as depicted in (C). Without adding any extra sequences to the sample, the two bars are associated to the same four generators, allowing only to reconstruct the large envelope of the two loops in the tARG. Adding sequences E and F to the sample (represented by blue leaf nodes in (A)) disentangles the generators of the two loops, fully reconstructing the topology of the tARG.

doi:10.1371/journal.pcbi.1005071.g007

the mutational distance between recombinant sequences in the two genomic intervals associated to each bar (characters 1 to 5, and 6 to 7, respectively). The position of the crossover breakpoints associated to these recombination events is also correctly reproduced. Hence, taking as input the 4 sequences, the barcode ensemble extracts phylogenetic information from ultra-minimal ARGs that explain the sample, without requiring complete reconstruction of the ARGs.

We note here the importance of using the barcode ensemble instead of the ordinary barcode, used in previous phylogenetic applications of persistent homology [35]. In this simple example $b_1 = 1$, and only one of the two recombination events would have been detected if the ordinary first-homology barcode had been used. The barcode ensemble largely increases the sensitivity to detect recombination events.

We can attempt to reconstruct the tARG of the sample by using persistent homology generators (Fig 7C). Whereas there are theorems ensuring the stability of barcodes against small perturbations [41, 42], the generators of persistent homology identified by TDA strongly depend on the sample, and multiple choices of basis are possible. Hence, the use of generators to

reconstruct the tARG is usually limited to small datasets. In this simple example, both bars in the barcode ensemble are generated by the four sampled sequences. Therefore, the reconstructed loop enclosing each recombination event is the same in both cases and corresponds to the large enveloping loop in the ultra-minimal ARG (Fig 7A). Adding the internal nodes 1111000 and 1111110 to the sample permits disentangling the generators of the two loops (Fig 7C), fully reconstructing the topology of the underlying tARG.

Genetic exchange between two divergent populations. We now consider a more involved example consisting of two sexually-reproducing populations, simulated under the coalescent model with recombination. The two populations diverged $24N$ generations before present. Their effective population sizes are taken to be constant and given by N and $N/5$. We consider two different cases, depicted in Fig 8. In the first case (Fig 8A), the two populations are completely isolated from each other. In the second case (Fig 8B), to the contrary, there is a small migration rate between the two populations. The recombination rate is the same in both cases. Alternatively, in a phylogenetic context, this setting describes the incomplete lineage sorting of two species, with or without the presence of gene flow.

We randomly sampled 250 sequences from the large population and 50 sequences from the small population. The full sample consisted of 300 sequences with 300 segregating sites. We present in Fig 8 the barcode ensemble for simulated samples without and with migration. The computation took approximately 33 minutes (wall-clock time) in a modern 8-cores desktop computer. Whereas the number of detected recombination events in the tARG, counted by the number bars, is similar in both cases, their genetic scales are very different. Specifically, in the presence of migration the size of some of the loops in the tARG is large, corresponding to migration events followed by a recombination event (Fig 8B). This is indicated by the presence of bars with large death time ϵ_d in the barcode ensemble of the case with migration, corresponding to recombination events with large mutational distances between recombining sequences.

Hence, in this example the barcode ensemble provides rich phylogenetic information that could be hardly obtained by other methods. Methods that attempt to construct a minimal (or nearly-minimal) ARG are computationally inefficient for such large sample sizes [33, 52]. Fast bound methods [33, 44, 45], on the other hand, do not provide enough phylogenetic information to distinguish between the cases with and without migration, as the total recombination rate is the same in both situations. Sequentially Markov coalescent approaches [34] produce an ARG that is far from being minimal but is a good approximation to the maximum likelihood. However, these methods require an underlying coalescent model, with mutation, recombination and population structure parameters given as priors. Finally, algorithms for constructing phylogenetic split networks [21] are fast and provide very different outputs in each of the above two cases. However, the interpretation of the output in terms of recombination and migration events is obscure.

Human leukocyte antigen (HLA) locus. The previous examples serve to illustrate the relation between features of the barcode ensemble of a genetic sample and those of the ultra-minimal ARGs explaining the sample. However, both examples are based on simulated data. We now consider a more realistic example, consisting of 180 phased genotypes from a ~ 250 kilobase region of the HLA locus of 90 individuals belonging to the Luhya in Webuye, Kenya (LWK) population, sequenced as part of the International HapMap Project [53]. In total, the region contains 471 SNPs. Recombination hotspots in this part of the HLA locus have been studied in detail in the past through sperm typing [54] and other high-resolution methods [55]. This example therefore serves to illustrate the capacity of the barcode ensemble to localize recombination events in realistic situations. With this aim, we also considered 194 phased genotypes from a smaller region (40 kilobase) within the same HLA locus of 97 individuals from the same population, sequenced by the 1,000 Genomes Project Consortium [56]. This

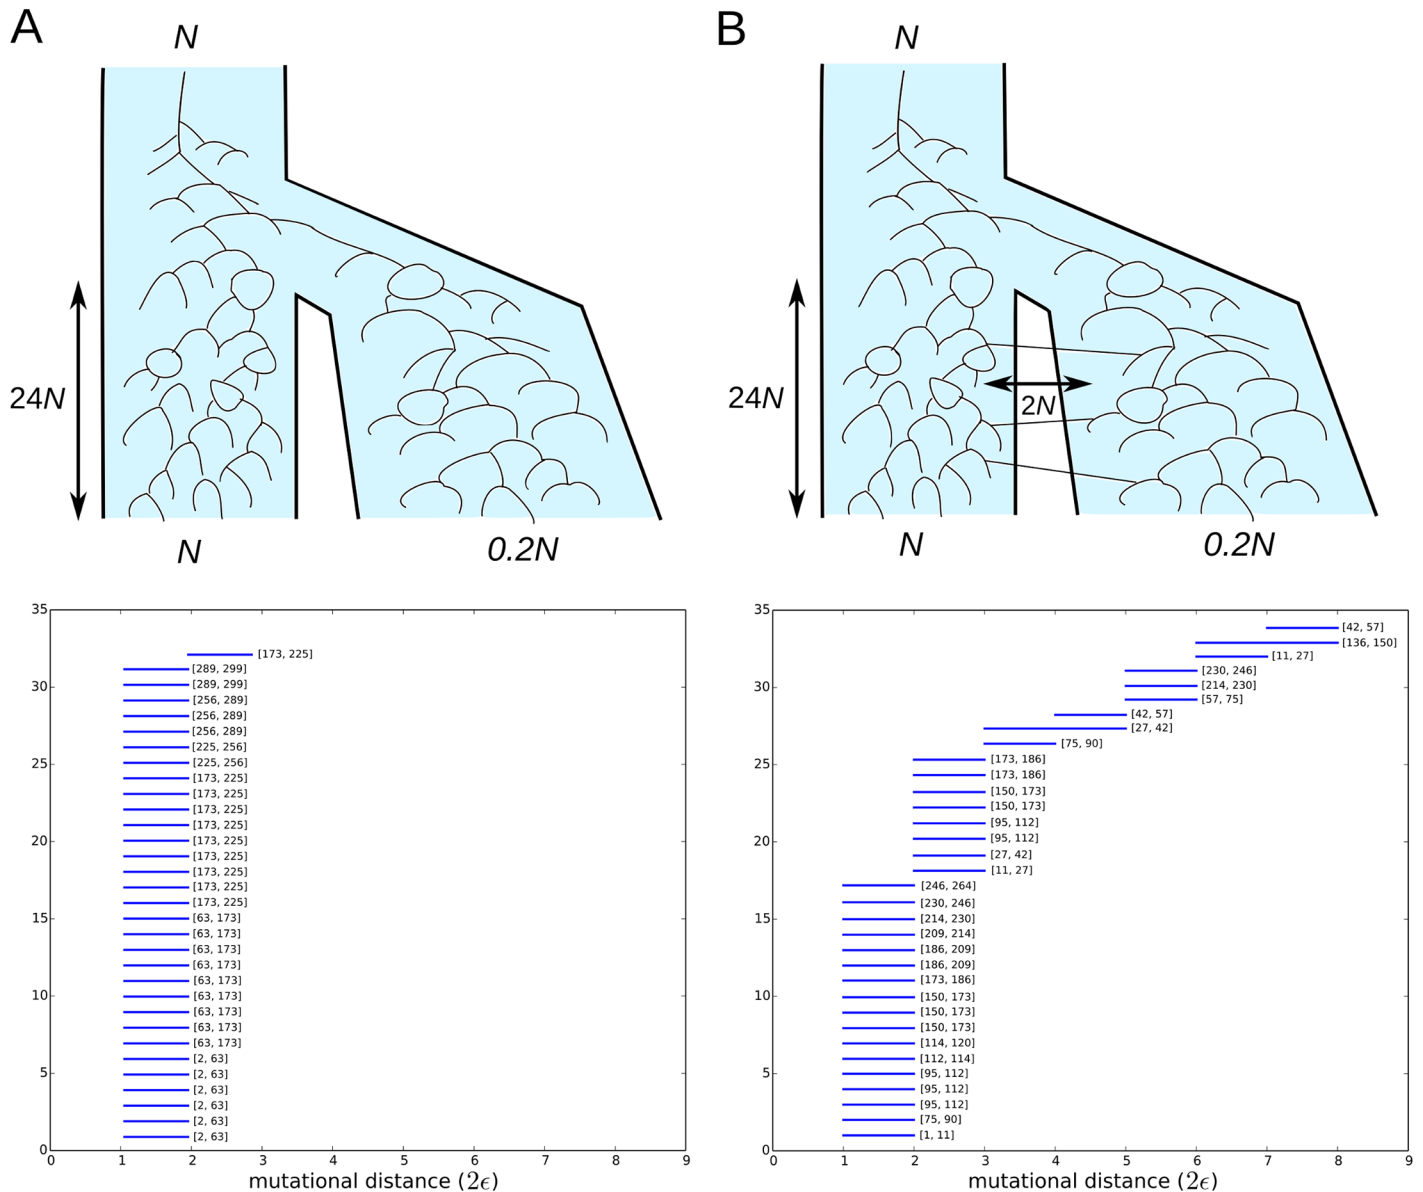


Fig 8. Barcode ensemble of two divergent sexually-reproducing populations. The case in (A) assumes the two populations are completely isolated. All recombination events present in the barcode ensemble involve genetically close parental gametes. The case in (B) considers a small migration rate between the two populations. Some of the recombination events present in the barcode ensemble involve genetically distant parental strains, leading to larger death times ϵ_d in the barcode ensemble. The total number of detected recombination events is similar in both cases and uniform across the entire genome. Intervals with the location of the recombination breakpoints are indicated for each recombination event, where positions refer to segregating sites.

doi:10.1371/journal.pcbi.1005071.g008

additional dataset contained a higher density of SNPs (482 SNPs in total), allowing for a higher resolution in the localization of recombination events.

We used TARGeT to compute the first-homology barcode ensembles of the two datasets and analysed the distribution of bars across the HLA locus (Fig 9A and 9B). The computation took 19 and 14 minutes (wall-clock time) in a modern 8-cores desktop computer, respectively for the HapMap and 1,000 Genomes datasets. Comparison with the African-American recombination map [55], based on more than 2 million crossovers in 30,000 unrelated African-

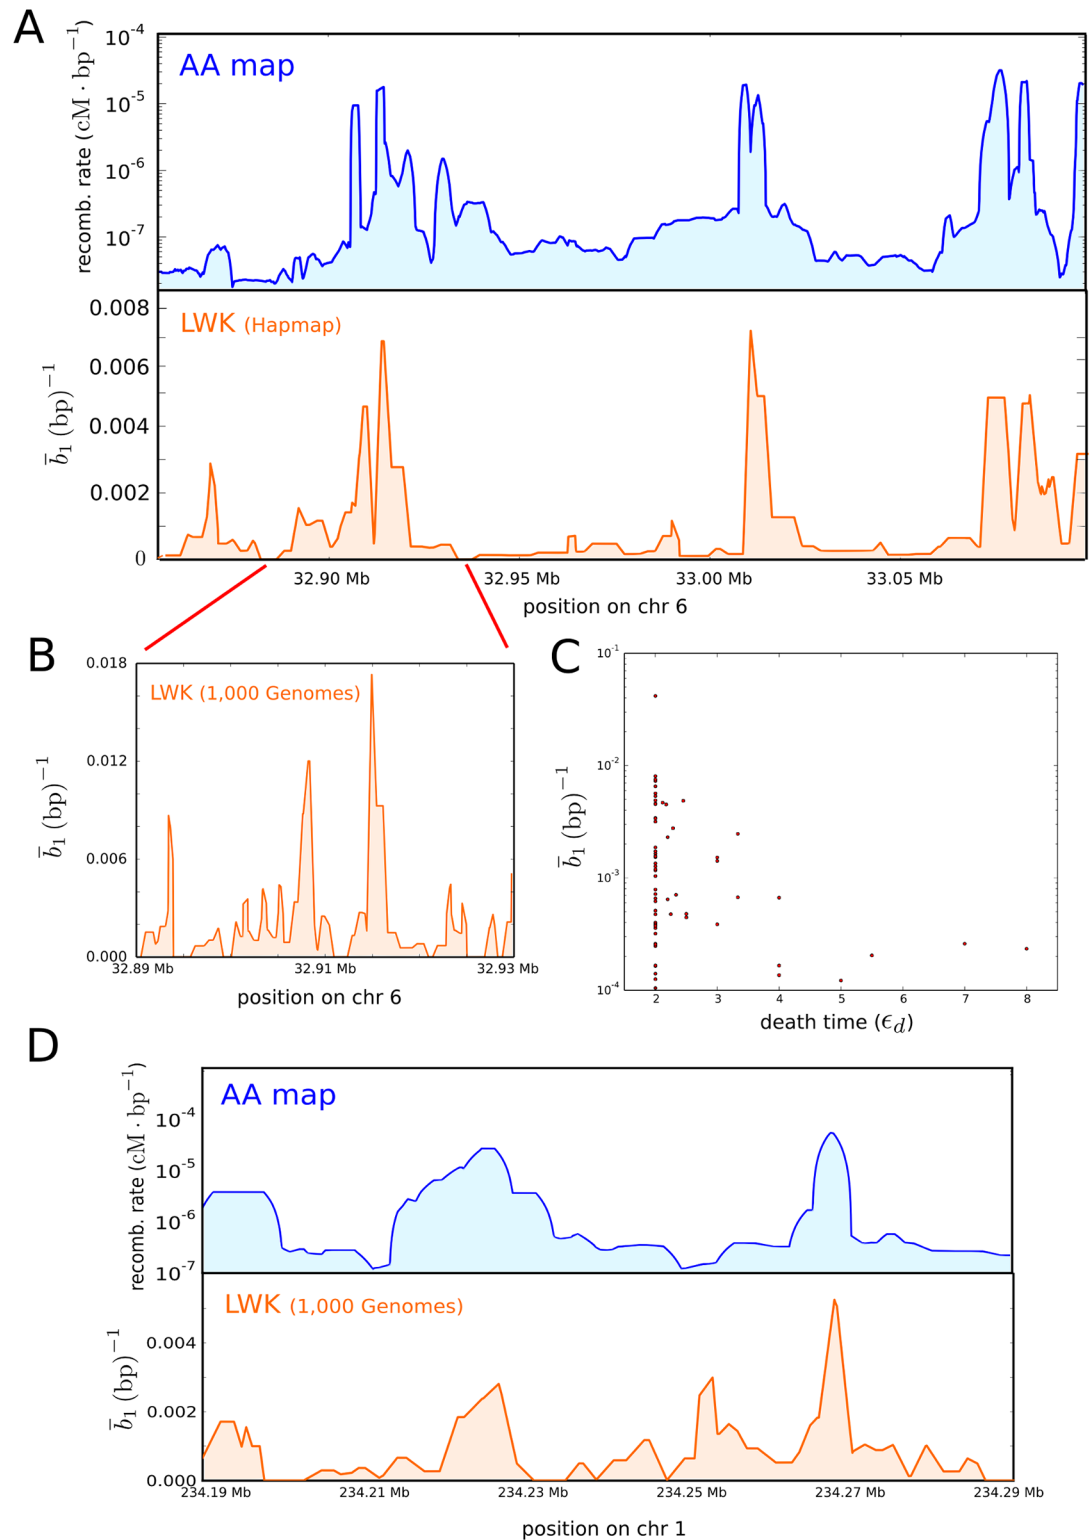


Fig 9. Barcode ensemble across the HLA and MS32 mini-satellite loci of the LWK population. (A) Recombination rates (top) across a 250 kilobase region of the HLA locus according to the African-American recombination map, based on 30,000 individuals [55]. The vertical axis is in logarithmic scale. The distribution of recombination events (bottom) detected by the barcode ensemble of a sample of 90 individuals from the LWK population sequenced by the International HapMap Consortium [53] is consistent with the observed recombination

rates. Note that in neutral models of evolution the number of recombination events in minimal ARGs is roughly expected to grow logarithmically with the recombination rate of the population [57]. (B) Distribution of recombination events detected by the barcode ensemble of a sample of 97 individuals from the LWK population, sequenced by the 1,000 Genomes Project Consortium [56]. The higher density of SNPs in this dataset allows for a higher resolution in the localization of recombination events as well as a higher sensitivity. (C) Density of recombination events per nucleotide against their average death time $2\epsilon_d$, for recombination events captured by the barcode ensemble in (A). Each point represents a genomic position for which the barcode ensemble detects recombination. The horizontal axis represents the average death time of the bars in the barcode ensemble that are associated to that genomic position. Events with large ϵ_d , corresponding to recombination events with a large mutational distance between recombining sequences, are mostly associated to regions with low number of recombinations, as expected from neutral models of evolution [57]. (D) Recombination rates (top) across a 100 kilobase region near the MS32 mini-satellite locus according to the African-American recombination map [55]. The vertical axis is in logarithmic scale. The distribution of recombination events (bottom) detected by the barcode ensemble of a sample of 97 individuals from the LWK populations sequenced by the 1,000 Genomes Project Consortium [56] is consistent with the observed recombination rates.

doi:10.1371/journal.pcbi.1005071.g009

Americans, shows a large degree of consistency between the recombination rates and the genomic position of recombination events detected by the barcode ensembles (Fig 9A). The distribution of mutational distances associated to recombination events is qualitatively consistent with coalescent arguments (Fig 9C). In particular, bars with large death time ϵ_d , corresponding to recombination events with a large mutational distance between recombining sequences, are mostly associated to regions of low recombination rate, consistently with the longer coalescence time for these regions [57].

Human MS32 mini-satellite locus. Similarly, we also considered 194 phased genotypes from a ~ 100 kilobase region (416 segregating sites) near the MS32 mini-satellite locus of 97 individuals from the LWK population, sequenced by the 1,000 Genomes Project Consortium [56]. We used TARGeT to compute the barcode ensemble and studied the distribution of bars across this genomic region (Fig 9D). The computation took 10 minutes (wall-clock time) in a modern 8-cores desktop computer. As in the previous example, the genomic position of the recombination events detected by the barcode ensemble (Fig 9D) was consistent with the recombination rate across this region, as determined by the African-American recombination map [55].

Darwin's finches. Our last example consists of the genetic sequences of 112 Darwin's finches, belonging to 15 different species inhabiting the Galápagos archipelago and Cocos Island [51]. We aligned and genotyped a 9 megabase scaffold of their genome and, after filtering for high-quality variants, we focussed on a set of 140 SNPs that were homozygous across the 112 samples, thus avoiding potential phasing artefacts. By considering this set, we mostly restrict to very ancestral recombination/gene flow events, close to the origin of radiation from a common ancestor 1.5 million years ago [58]. We used TARGeT to obtain the first-homology barcode ensemble of the sample, as well as the partially reconstructed tARG. The computation took 9 minutes (wall-clock time) in a modern 8-cores desktop computer.

The first-homology barcode ensemble (Fig 10A) contains 13 recombination events, mostly involving samples from multiple species and usually including samples from the genus *Certhiidea* (Fig 10B), the most ancestral lineage among the genera present in the sample [51]. These results add support to the evidence for genetic introgression found in [51]. Our analysis also reveals that the crossover breakpoints of these events localize at four different genomic regions within the 9 megabase scaffold that we have considered in this example (Fig 10C).

Parameter estimation

The examples above illustrate the use and interpretation of barcode ensembles in molecular phylogenetics. As we have discussed, an important feature of topological approaches to phylogenetics is that they inform about most parsimonious evolutionary histories. Being model-independent

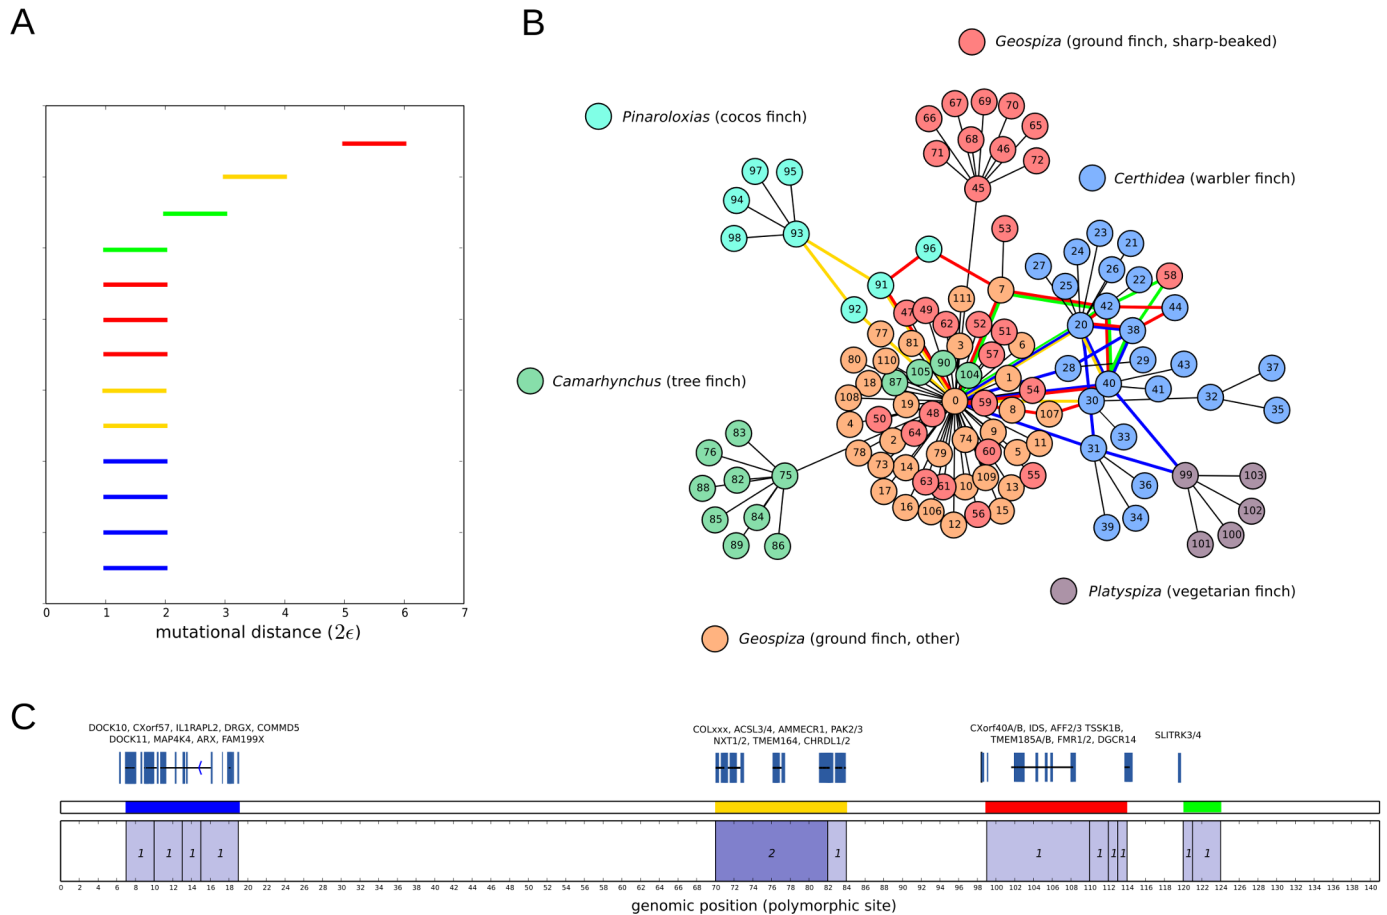


Fig 10. Barcode ensemble and partially reconstructed tARG of a sample of 112 Darwin's finches. The barcode ensemble is shown in (A), based on 140 homozygous SNPs present in a 9 megabase scaffold. In total, 13 recombination/gene flow events are captured in the barcode ensemble, with different genetic scales. Bars are colored according to the position of the corresponding recombination breakpoint in the genome, as depicted in (C). We also indicate the number of recombination events detected at each genomic interval, as well as some of the orthologous genes present at regions where recombination events are detected. The reconstructed tARG is presented in (B). Loops in the reconstructed tARG are outlined using the same code of colors. We have also included leaf nodes that do not participate in any recombination event, using a nearest neighbour algorithm based on genetic distance. Edge lengths are arbitrary.

doi:10.1371/journal.pcbi.1005071.g010

approaches, they describe minimal sets of events required to explain a sample of sequences, without assuming any probabilistic model of evolution. In some situations, however, we are interested in estimating the parameters of a specific evolutionary model from the observed data (e.g. the recombination rate in a coalescent model with recombination). To that end, barcode ensembles can be taken as summary statistics from which to build parameter estimators. For instance, in Fig 11A we show the dependence of \bar{b}_1 on the recombination rate for a set of 1,000 coalescent model simulations. The expected \bar{b}_1 of the barcode ensemble is informative of the recombination rate, growing monotonically with the later. Compared to sequentially Markov coalescent (SMC) approaches for ARG inference [34], \bar{b}_1 is strongly correlated with the number of recombinations in SMC ARGs derived from the same set of sequences (Pearson's $r = 0.93, p < 10^{-100}$, S1 Fig). Although the coefficient of variation is $\sim 35\%$ larger for \bar{b}_1 (S1 Fig), its computing time is substantially lower (> 9 times faster after parallelizing in a modern 8-cores desktop computer, S1 Fig), being a robust approach to coalescent-model recombination rate estimation in large datasets. Furthermore, unlike the number of recombinations in SMC ARGs, \bar{b}_1 is unbiased at small recombination rates, vanishing when the recombination rate is zero (Fig 11A).

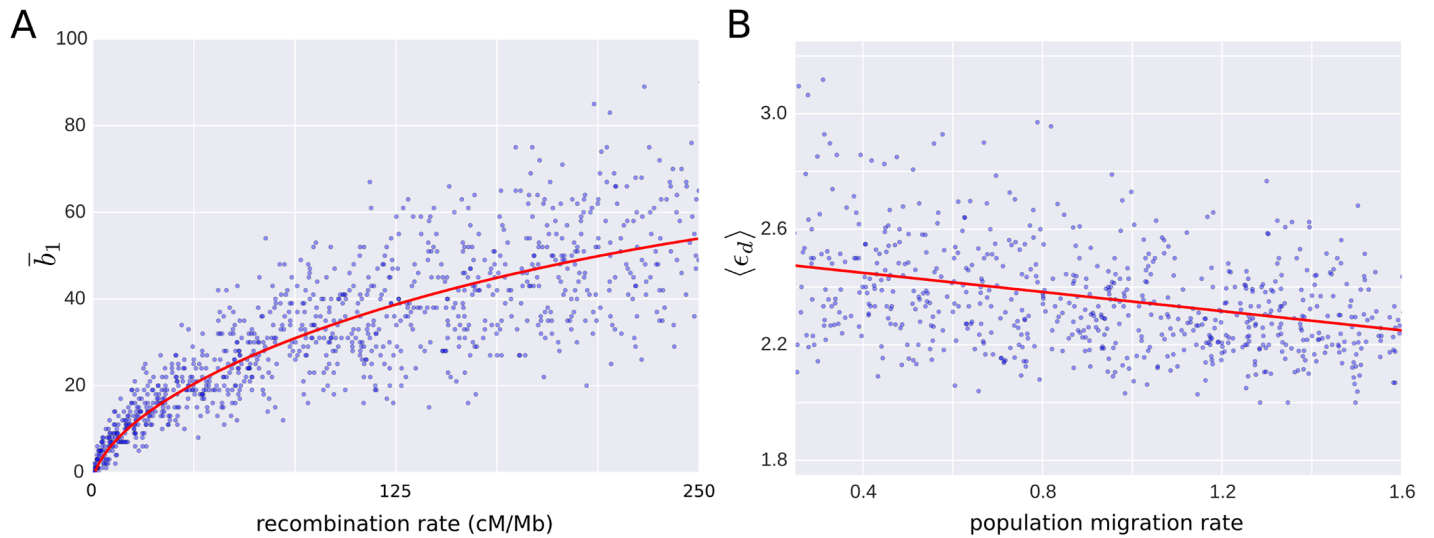


Fig 11. Parameter estimation in models of evolution. (A) Dependence of \bar{b}_1 on the recombination rate parameter for a set of 1,000 simulations of a basic coalescent model. Each simulation consists of 200 sequences, 30 kilobase long. The expected \bar{b}_1 of the barcode ensemble grows monotonically with the recombination rate, providing a good measure of the later. The smoothed average is shown in red. (B) Dependence of the average death time, $\langle \epsilon_d \rangle$, on the migration rate of two divergent populations with fixed recombination and variable migration rates, based on 900 simulations. Each simulation consists of 150 sampled sequences, 10 kilobase long. The same structure as in Fig 8 was considered for the two populations. The expected value of $\langle \epsilon_d \rangle$ decreases monotonically with the migration rate, being informative of the later.

doi:10.1371/journal.pcbi.1005071.g011

Although recombination rate estimation is a very direct example, the barcode ensemble of a sample of genetic sequences contains other rich phylogenetic information apart from \bar{b}_1 , which can be used for more complex parameter estimation in structured models of evolution. Consider, for instance, the case of two divergent populations with migration and recombination discussed above. In this model, the average genetic distance between recombining sequences is expected to decrease with the migration rate, as the average time to the most recent common ancestor between foreign and local gametes in a population is shorter. In Fig 11B we show the dependence of the average death time ($\langle \epsilon_d \rangle$) on the migration rate parameter, for the barcode ensembles of a set of 900 coalescent model simulations with fixed recombination and variable migration rates. As expected, $\langle \epsilon_d \rangle$ is informative of the migration rate, decreasing monotonically with the later. It is therefore a good measure for estimating migration rates. Consistently, $\langle \epsilon_d \rangle$ correlates with time to the most recent common ancestor of recombining sequences in SMC ARGs obtained from the same data (Pearson's $r = 0.55$, $p < 10^{-72}$, S1 Fig). Although the coefficient of variation of $\langle \epsilon_d \rangle$ is $\sim 60\%$ larger (S1 Fig), extracting this type of information from SMC ARGs requires the implementation of a greedy algorithm, substantially increasing the running time (~ 8 times slower in a single core of modern desktop computer, S1 Fig) and therefore limiting its applicability to large datasets.

These two simple examples illustrate the utility of barcode ensembles for building parameter estimators in specific models of evolution. Importantly, being model-independent, they are robust and flexible tools which can be applied in an infinitely large number of possible evolutionary models.

Discussion

As the famous title of the essay by Dobzhansky “Nothing in Biology Makes Sense Except in the Light of Evolution” underscores, evolutionary processes are central orchestrating themes in biology. Mutations, recombinations and other evolutionary processes get imprinted into

genomes through selection, reflecting the accumulated history giving rise to an organism. Phylogenetics try to reconstruct the evolutionary history through the comparison of genomes of related organisms. In addition to reporting relationships and elucidating particular histories, one would like to understand and quantify how different evolutionary processes have occurred. The identification and quantification of evolutionary processes can be challenging due to the lack of a well-established universal framework to capture evolutionary relationships beyond trees. In addition, robust statistical inference needs to exploit the large number of genomes that are now becoming available, aggravating the computational burden and obscuring interpretations. Ideally, we would like to have a biologically interpretable framework able to quantify different evolutionary processes by analyzing large numbers of genomes.

In this paper we have proposed a few steps in this direction. We have extended the notion of barcodes in persistent homology to identify the genetic scale and number of recombination events. We have shown that, by correctly studying persistent homology in subsets of segregating sites, it is possible to characterize the genomic regions where recombination takes place and identify the gametes involved in particular recombination events. The persistent homology barcodes derived from each of these sets can be structured as a “barcode ensemble” where each bar captures a recombination event. Barcode ensembles can be interpreted as counting and quantifying the scale of recombination events in a variation of Ancestral Recombination Graphs (ARGs). Topological ARGs represent a summary of potential recombination histories that can explain the data. The method proposed, *TARGeT*, is scalable to hundreds of genomes. As an alternative to some phylogenetic networks, barcode ensembles provide robust quantification of events, the distribution of genetic scales, computational scalability and interpretative graphs.

Barcode ensembles are versatile in that they do not assume any specific model of evolution, providing explicit, interpretable summaries of the minimal set of recombination events required to explain a sample of genetic sequences. Here we have illustrated their use in several practical cases. However, the range of possible applications is unlimited. In some cases, it may be convenient to perform minor modifications to the approach described here. For instance, although in our exposition we have only made use of Hamming distance and binary sequences, the main concepts we have presented extend straightforwardly to other genetic distances. The use of these metrics can be particularly useful in cases with rapidly diverging samples or substantial mutational biases. In other cases, information about the ancestral and derived alleles for each character in the sample may be available. Although *t*ARGs have no natural directionality, the inclusion of the ancestral sequence in the original sample may lead in those cases to more stringent bounds on \bar{R}_{\min} , similarly to what occurs with other approaches to recombination inference [22]. Finally, more efficient integer linear programming algorithms, like the one of [33], could in principle be also generalized to the computation of barcode ensembles.

Methods

First-homology barcode ensemble

We extended the construction of ref. [45] to persistent homology barcodes. From a geometric perspective, this corresponds to projecting the original space on sets of mutually orthogonal hyperplanes in the ambient hypercube, and computing persistent homology in each of those projections. For that aim, we need to establish an ordering relation on barcodes. Being sets of intervals, it is natural to take the maximum of two barcodes to be given by the one with largest L^0 -norm, namely largest b_1 . If both barcodes have the same L^0 -norm, we may successively compare other norms (e.g. other L^p -norms), until the tie is broken or, otherwise, one of the two

barcodes is arbitrarily chosen. The algorithm of [45] is then generalized to persistent homology barcodes as follows:

1. Let \mathcal{B}_{ik} be the first-homology barcode of the sequences that result from the i -th to k -th characters in \mathcal{S} . Set $\mathcal{R}_{ij} = 0$ and $k = 2$.
2. For $j = 1, \dots, k - 1$, set $\mathcal{R}_{jk} = \max\{\mathcal{R}_{ji} \cup \mathcal{B}_{ik} : i = j, \dots, k - 1\}$
3. If $k < m$, increment k by 1 and go to step 2.

The barcode ensemble of \mathcal{S} is the union barcode \mathcal{R}_{1m} that results from this algorithm.

We implemented the algorithm in a publicly available multi-threaded software, `TARGET`, which is distributed under the GNU General Public License (GPL v3). The application is fully written in Python 2.7, and relies on Dionysus C++ library for persistent homology computations (<http://www.mrzv.org/software/dionysus>). Since considering all possible sequence partitions is unnecessary and computationally infeasible in most cases, we follow the strategy of ref. [45] and allow the user to limit the number of partitions by the maximum number of segregating characters within each subset of \mathcal{S} (specified by the command line option `-s`), and by the maximum distance between segregating characters in the subset (specified by the command line option `-w`). In addition, we also allow the user to exclude from \mathcal{S} segregating characters that are compatible (namely, that satisfy the Hudson-Kaplan four-gamete test [44]) with all the other characters in \mathcal{S} (specified by the command line option `-e`). For each genomic interval, a filtration of Vietoris-Rips complexes is constructed using Hamming distance and the persistent first-homology group is computed over \mathbf{Z}_2 .

Population genetics simulations

We performed 4,000 simulations of a sample of 40 sequences with 12 segregating sites, using the software `ARGweaver` [34]. The population was simulated using a coalescent infinite sites model with recombination. The population-scaled recombination rate, ρ , was randomly generated in each simulation, taking values from a uniform distribution between 0 and 110. For each simulated sample, Myers and Griffiths lower bound $R_{MG} \leq R_{min}$ was computed using the software `RecMin` [45], with parameters `-s 12 -w 12`. Lower bounds $\bar{b}_1 \leq \bar{R}_{min}$ were computed using our application `TARGET`, with parameters `-s 12 -w 12`.

To study the dependence of \bar{b}_1 on the recombination rate parameter in coalescent models, we performed 1,000 simulations of a sample of 200 sequences. The population-scaled recombination rate, ρ , was randomly generated in each simulation, taking values between 0 and 216. For each simulated sample, `TARGET` was run with parameters `-s 11 -w 11`, and `ARGweaver`'s tool `arg-sample` was run with parameters `-m 7e-9 -n 400 --sample-step 10`, discarding the first 200 iterations.

Samples of genetic exchange between two divergent populations were simulated using the software `ms` [59], using the commands

```
ms 300 1 -s 300 -r 40 10000 -I 2 250 50 -ej 6.0 1 2 -n 2 0.2 -m 1 2 0.5
and,
```

```
ms 300 1 -s 300 -r 40 10000 -I 2 250 50 -ej 6.0 1 2 -n 2 0.2
```

respectively for the cases with and without migration. The barcode ensemble of each sample was computed using `TARGET` with parameters `-s 12 -w 14 -e`.

To study the dependence of $\langle \epsilon_d \rangle$ on the migration rate in this scenario, we performed 900 simulations using the software `ms` [59] and `seq-gen` [60], with the commands

```
ms 150 1 -T -r 60 10000 -I 2 125 25 -ej 6.0 1 2 -n 2 0.2 -m 1 2 X
and
```



```
seq-gen -mHKY -l 10000 -s 0.004 -p 50000
```

where the migration rate X in the first command takes random values from a uniform distribution between 0 and 2. For each simulated sample, `TARGet` was run with parameters `-w 8 -s 8`, and `arg-sample` was run with parameters `-m 1e-7 -n 400 -r 1.5e-7`. We extracted from SMC ARGs the time to the most recent common ancestor of recombining sequences using a greedy algorithm that searches for the shortest non-zero path connecting the two sequences.

HLA and MS32 loci

We downloaded phased genotype data from HapMap phase III [53], corresponding to all SNPs of LWK population between rs6457661 and rs3129301 in chromosome 6. We also downloaded phased genotype data from 1,000 Genomes Project [3], corresponding to all SNPs of LWK population between positions 32,887,978 and 32,927,978 of chromosome 6, and half of the SNPs of LWK population between positions 234,190,031 and 234,291,193 of chromosome 1. All coordinates refer to human assembly hg18. The barcode ensemble of each dataset was computed using `TARGet` with parameters `-s 12 -w 12`.

Darwin's finches genotyping

Raw paired-end reads from 112 Darwin finches [51] were obtained from SRA archive (accession number PRJNA263122) and aligned against the consensus sequence of *Geospiza Fortis*, version `GeoFor_1.0/geoFor1`, scaffold JH739904. We followed essentially the same procedure than that of ref. [51] for the alignment, SNP calling, genotyping and filtering. In short, the alignment was performed with Burrows-Wheeler aligner (BWA) [61], version 0.7.5, using BWA-MEM algorithm and default parameters. PCR duplicates were marked using Picard tools (<http://picard.sourceforge.net/>). Indel realignment, SNP discovery and simultaneous genotyping across the 112 samples was performed using Genome Analysis Toolkit (GATK) [62], following GATK best practice recommendations [63]. SNP calls were filtered by keeping variants with SNP quality > 100 , total depth of coverage > 117 and < 1750 , ratio between SNP quality and depth of coverage > 2 , Fisher strand bias < 60 , mapping quality > 50 , mapping quality rank > -4 and read position rank sum > -2 . In total, 13,980 variant positions passed these filters. To avoid phasing errors, we only considered SNPs that were homozygous across the 120 samples. The resulting genotypes were processed with `TARGet` for barcode ensemble computation, using the options `-s 14 -w 14`.

Supporting Information

S1 Fig. Comparison to SMC approaches to ARG inference. (A) The number of recombination events in SMC ARGs [34], $R_{\text{argweaver}}$, plotted against the number of bars in the barcode ensemble, \bar{b}_1 . Both quantities are strongly correlated (Pearson's $r = 0.93$, $p < 10^{-100}$). Plot based on 1,000 coalescent model simulations of a sample of 200 sequences. (B) Coefficient of variation of $R_{\text{argweaver}}$ (blue) and \bar{b}_1 (red) as a function of the recombination rate. (C) Distribution of wall-clock running times for the simulations in (A). (D) Average time to the most recent common ancestor of recombining sequences in SMC ARGs, $\langle T_{\text{argweaver}} \rangle$, plotted against the average death time of bars in the barcode ensemble, $\langle \epsilon_d \rangle$, for two divergent populations with recombination and migration. Both quantities are largely correlated (Pearson's $r = 0.55$, $p < 10^{-72}$). Plot based on 900 simulations of a sample of 150 sequences. (E) Coefficient of variation of $\langle T_{\text{argweaver}} \rangle$ (blue) and $\langle \epsilon_d \rangle$ (red) as a function of the migration rate. (F) Distribution of wall-clock running times for the simulations in (D). ARGweaver running times also include

the time required to extract $\langle T_{\text{argweaver}} \rangle$ from SMC ARGs. (TIFF)

Acknowledgments

We thank C. Solis-Lemus, C. Ane, D. S. Rosenbloom, K. Emmett, M. Lesnick, S. Zairis and V. Aleksandrov for useful discussions and comments.

Author Contributions

Conceived and designed the experiments: PGC RR.

Analyzed the data: PGC.

Wrote the paper: PGC AJL RR.

Developed and implemented the algorithm: PGC.

References

1. Lander ES, Linton LM, Birren B, Nusbaum C, Zody MC, Baldwin J, et al. Initial sequencing and analysis of the human genome. *Nature*. 2001; 409(6822):860–921. doi: [10.1038/35057062](https://doi.org/10.1038/35057062) PMID: [11237011](https://pubmed.ncbi.nlm.nih.gov/11237011/)
2. Venter JC, Adams MD, Myers EW, Li PW, Mural RJ, Sutton GG, et al. The sequence of the human genome. *science*. 2001; 291(5507):1304–1351. doi: [10.1126/science.1058040](https://doi.org/10.1126/science.1058040) PMID: [11181995](https://pubmed.ncbi.nlm.nih.gov/11181995/)
3. 1000 Genomes Project Consortium et al. An integrated map of genetic variation from 1,092 human genomes. *Nature*. 2012; 491(7422):56–65. doi: [10.1038/nature11632](https://doi.org/10.1038/nature11632) PMID: [23128226](https://pubmed.ncbi.nlm.nih.gov/23128226/)
4. Edwards RA, Rohwer F. Viral metagenomics. *Nature Reviews Microbiology*. 2005; 3(6):504–510. doi: [10.1038/nrmicro1163](https://doi.org/10.1038/nrmicro1163) PMID: [15886693](https://pubmed.ncbi.nlm.nih.gov/15886693/)
5. Albertsen M, Hugenholtz P, Skarshewski A, Nielsen KL, Tyson GW, Nielsen PH. Genome sequences of rare, uncultured bacteria obtained by differential coverage binning of multiple metagenomes. *Nature biotechnology*. 2013; 31(6):533–538. doi: [10.1038/nbt.2579](https://doi.org/10.1038/nbt.2579) PMID: [23707974](https://pubmed.ncbi.nlm.nih.gov/23707974/)
6. Ochman H, Lawrence JG, Groisman EA. Lateral gene transfer and the nature of bacterial innovation. *Nature*. 2000; 405(6784):299–304. doi: [10.1038/35012500](https://doi.org/10.1038/35012500) PMID: [10830951](https://pubmed.ncbi.nlm.nih.gov/10830951/)
7. Burke JM, Arnold ML. Genetics and the fitness of hybrids. *Annual review of genetics*. 2001; 35(1):31–52. doi: [10.1146/annurev.genet.35.102401.085719](https://doi.org/10.1146/annurev.genet.35.102401.085719) PMID: [11700276](https://pubmed.ncbi.nlm.nih.gov/11700276/)
8. Adams KL, Wendel JF. Polyploidy and genome evolution in plants. *Current opinion in plant biology*. 2005; 8(2):135–141. doi: [10.1016/j.pbi.2005.01.001](https://doi.org/10.1016/j.pbi.2005.01.001) PMID: [15752992](https://pubmed.ncbi.nlm.nih.gov/15752992/)
9. Doolittle WF, Papke RT. Genomics and the bacterial species problem. *Genome biology*. 2006; 7(9):116. doi: [10.1186/gb-2006-7-9-116](https://doi.org/10.1186/gb-2006-7-9-116) PMID: [17020593](https://pubmed.ncbi.nlm.nih.gov/17020593/)
10. Doolittle WF. Phylogenetic classification and the universal tree. *Science*. 1999; 284(5423):2124–2128. doi: [10.1126/science.284.5423.2124](https://doi.org/10.1126/science.284.5423.2124) PMID: [10381871](https://pubmed.ncbi.nlm.nih.gov/10381871/)
11. Woese CR. A new biology for a new century. *Microbiology and Molecular Biology Reviews*. 2004; 68(2):173–186. doi: [10.1128/MMBR.68.2.173-186.2004](https://doi.org/10.1128/MMBR.68.2.173-186.2004) PMID: [15187180](https://pubmed.ncbi.nlm.nih.gov/15187180/)
12. O'Malley MA, Koonin EV. How stands the Tree of Life a century and a half after The Origin. *Biol Direct*. 2011; 6:32. doi: [10.1186/1745-6150-6-32](https://doi.org/10.1186/1745-6150-6-32)
13. Huson DH, Rupp R, Scornavacca C. *Phylogenetic networks: concepts, algorithms and applications*. Cambridge University Press; 2010.
14. Huson DH, Scornavacca C. A survey of combinatorial methods for phylogenetic networks. *Genome biology and evolution*. 2011; 3:23–35. doi: [10.1093/gbe/evq077](https://doi.org/10.1093/gbe/evq077) PMID: [21081312](https://pubmed.ncbi.nlm.nih.gov/21081312/)
15. Morrison DA. *Introduction to phylogenetic networks*. RJR Productions. 2011;.
16. Bandelt HJ, Dress AW. Split decomposition: a new and useful approach to phylogenetic analysis of distance data. *Molecular phylogenetics and evolution*. 1992; 1(3):242–252. doi: [10.1016/1055-7903\(92\)90021-8](https://doi.org/10.1016/1055-7903(92)90021-8) PMID: [1342941](https://pubmed.ncbi.nlm.nih.gov/1342941/)
17. Bandelt HJ, Forster P, Sykes BC, Richards MB. Mitochondrial portraits of human populations using median networks. *Genetics*. 1995; 141(2):743–753. PMID: [8647407](https://pubmed.ncbi.nlm.nih.gov/8647407/)
18. Huson DH. SplitsTree: analyzing and visualizing evolutionary data. *Bioinformatics*. 1998; 14(1):68–73. doi: [10.1093/bioinformatics/14.1.68](https://doi.org/10.1093/bioinformatics/14.1.68) PMID: [9520503](https://pubmed.ncbi.nlm.nih.gov/9520503/)

19. Bandelt HJ, Forster P, Röhl A. Median-joining networks for inferring intraspecific phylogenies. *Molecular biology and evolution*. 1999; 16(1):37–48. doi: [10.1093/oxfordjournals.molbev.a026036](https://doi.org/10.1093/oxfordjournals.molbev.a026036) PMID: [10331250](https://pubmed.ncbi.nlm.nih.gov/10331250/)
20. Bandelt HJ, Macaulay V, Richards M. Median networks: speedy construction and greedy reduction, one simulation, and two case studies from human mtDNA. *Molecular phylogenetics and evolution*. 2000; 16(1):8–28. doi: [10.1006/mpev.2000.0792](https://doi.org/10.1006/mpev.2000.0792) PMID: [10877936](https://pubmed.ncbi.nlm.nih.gov/10877936/)
21. Huson DH, Bryant D. Application of phylogenetic networks in evolutionary studies. *Molecular biology and evolution*. 2006; 23(2):254–267. doi: [10.1093/molbev/msj030](https://doi.org/10.1093/molbev/msj030) PMID: [16221896](https://pubmed.ncbi.nlm.nih.gov/16221896/)
22. Gusfield D. *ReCombinatorics: The Algorithmics of Ancestral Recombination Graphs and Explicit Phylogenetic Networks*. MIT Press; 2014.
23. Minichiello MJ, Durbin R. Mapping trait loci by use of inferred ancestral recombination graphs. *The American Journal of Human Genetics*. 2006; 79(5):910–922. doi: [10.1086/508901](https://doi.org/10.1086/508901) PMID: [17033967](https://pubmed.ncbi.nlm.nih.gov/17033967/)
24. Sutter NB, Bustamante CD, Chase K, Gray MM, Zhao K, Zhu L, et al. A single IGF1 allele is a major determinant of small size in dogs. *Science*. 2007; 316(5821):112–115. doi: [10.1126/science.1137045](https://doi.org/10.1126/science.1137045) PMID: [17412960](https://pubmed.ncbi.nlm.nih.gov/17412960/)
25. Wu Y. Association mapping of complex diseases with ancestral recombination graphs: models and efficient algorithms. *Journal of Computational Biology*. 2008; 15(7):667–684. doi: [10.1089/cmb.2007.0116](https://doi.org/10.1089/cmb.2007.0116) PMID: [18651799](https://pubmed.ncbi.nlm.nih.gov/18651799/)
26. Le SQ, Durbin R. SNP detection and genotyping from low-coverage sequencing data on multiple diploid samples. *Genome research*. 2011; 21(6):952–960. doi: [10.1101/gr.113084.110](https://doi.org/10.1101/gr.113084.110) PMID: [20980557](https://pubmed.ncbi.nlm.nih.gov/20980557/)
27. Wall JD. A comparison of estimators of the population recombination rate. *Molecular Biology and Evolution*. 2000; 17(1):156–163. doi: [10.1093/oxfordjournals.molbev.a026228](https://doi.org/10.1093/oxfordjournals.molbev.a026228) PMID: [10666715](https://pubmed.ncbi.nlm.nih.gov/10666715/)
28. Wang L, Zhang K, Zhang L. Perfect phylogenetic networks with recombination. *Journal of Computational Biology*. 2001; 8(1):69–78. doi: [10.1089/106652701300099119](https://doi.org/10.1089/106652701300099119) PMID: [11339907](https://pubmed.ncbi.nlm.nih.gov/11339907/)
29. Bordewich M, Semple C. On the computational complexity of the rooted subtree prune and regraft distance. *Annals of combinatorics*. 2005; 8(4):409–423. doi: [10.1007/s00026-004-0229-z](https://doi.org/10.1007/s00026-004-0229-z)
30. Bordewich M, Semple C. Computing the minimum number of hybridization events for a consistent evolutionary history. *Discrete Applied Mathematics*. 2007; 155(8):914–928. doi: [10.1016/j.dam.2006.08.008](https://doi.org/10.1016/j.dam.2006.08.008)
31. Gusfield D, Eddhu S, Langley C. Optimal, efficient reconstruction of phylogenetic networks with constrained recombination. *Journal of bioinformatics and computational biology*. 2004; 2(01):173–213. doi: [10.1142/S0219720004000521](https://doi.org/10.1142/S0219720004000521) PMID: [15272438](https://pubmed.ncbi.nlm.nih.gov/15272438/)
32. Gusfield D. Optimal, efficient reconstruction of root-unknown phylogenetic networks with constrained and structured recombination. *Journal of Computer and System Sciences*. 2005; 70(3):381–398. doi: [10.1016/j.jcss.2004.12.009](https://doi.org/10.1016/j.jcss.2004.12.009)
33. Song YS, Wu Y, Gusfield D. Efficient computation of close lower and upper bounds on the minimum number of recombinations in biological sequence evolution. *Bioinformatics*. 2005; 21(suppl 1):i413–i422. doi: [10.1093/bioinformatics/bti1033](https://doi.org/10.1093/bioinformatics/bti1033) PMID: [15961486](https://pubmed.ncbi.nlm.nih.gov/15961486/)
34. Rasmussen MD, Hubisz MJ, Gronau I, Siepel A. Genome-wide inference of ancestral recombination graphs. *PLoS genetics*. 2014; 10(5):e1004342. doi: [10.1371/journal.pgen.1004342](https://doi.org/10.1371/journal.pgen.1004342) PMID: [24831947](https://pubmed.ncbi.nlm.nih.gov/24831947/)
35. Chan JM, Carlsson G, Rabadan R. Topology of viral evolution. *Proceedings of the National Academy of Sciences*. 2013; 110(46):18566–18571. doi: [10.1073/pnas.1313480110](https://doi.org/10.1073/pnas.1313480110)
36. Emmett KJ, Rabadan R. Characterizing Scales of Genetic Recombination and Antibiotic Resistance in Pathogenic Bacteria Using Topological Data Analysis. In: *Brain Informatics and Health*. Springer; 2014. p. 540–551.
37. Emmett K, Rosenbloom D, Camara P, Rabadan R. Parametric inference using persistence diagrams: A case study in population genetics. *arXiv preprint arXiv:14064582*. 2014;.
38. Carlsson G. Topology and data. *Bulletin of the American Mathematical Society*. 2009; 46(2):255–308. doi: [10.1090/S0273-0979-09-01249-X](https://doi.org/10.1090/S0273-0979-09-01249-X)
39. Edelsbrunner H, Letscher D, Zomorodian A. Topological persistence and simplification. *Discrete and Computational Geometry*. 2002; 28(4):511–533. doi: [10.1007/s00454-002-2885-2](https://doi.org/10.1007/s00454-002-2885-2)
40. Zomorodian A, Carlsson G. Computing persistent homology. *Discrete & Computational Geometry*. 2005; 33(2):249–274. doi: [10.1007/s00454-004-1146-y](https://doi.org/10.1007/s00454-004-1146-y)
41. Cohen-Steiner D, Edelsbrunner H, Harer J. Stability of persistence diagrams. *Discrete & Computational Geometry*. 2007; 37(1):103–120. doi: [10.1007/s00454-006-1276-5](https://doi.org/10.1007/s00454-006-1276-5)
42. Chazal F, Cohen-Steiner D, Guibas LJ, Mémoli F, Oudot SY. Gromov-Hausdorff Stable Signatures for Shapes using Persistence. In: *Computer Graphics Forum*. vol. 28. Wiley Online Library; 2009. p. 1393–1403.

43. Ghrist R. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*. 2008; 45(1):61–75. doi: [10.1090/S0273-0979-07-01191-3](https://doi.org/10.1090/S0273-0979-07-01191-3)
44. Hudson RR, Kaplan NL. Statistical properties of the number of recombination events in the history of a sample of DNA sequences. *Genetics*. 1985; 111(1):147–164. PMID: [4029609](https://pubmed.ncbi.nlm.nih.gov/4029609/)
45. Myers SR, Griffiths RC. Bounds on the minimum number of recombination events in a sample history. *Genetics*. 2003; 163(1):375–394. PMID: [12586723](https://pubmed.ncbi.nlm.nih.gov/12586723/)
46. Griffiths RC, Marjoram P. Ancestral inference from samples of DNA sequences with recombination. *Journal of Computational Biology*. 1996; 3(4):479–502. doi: [10.1089/cmb.1996.3.479](https://doi.org/10.1089/cmb.1996.3.479) PMID: [9018600](https://pubmed.ncbi.nlm.nih.gov/9018600/)
47. Griffiths RC, Marjoram P. An ancestral recombination graph. *Institute for Mathematics and its Applications*. 1997; 87:257.
48. Hudson RR. Properties of a neutral allele model with intragenic recombination. *Theoretical population biology*. 1983; 23(2):183–201. doi: [10.1016/0040-5809\(83\)90013-8](https://doi.org/10.1016/0040-5809(83)90013-8) PMID: [6612631](https://pubmed.ncbi.nlm.nih.gov/6612631/)
49. Hatcher A. *Algebraic Topology*. Cambridge University Press; 2002.
50. Ghrist R. *Elementary Applied Topology*. Self-published, printed by CreateSpace; 2014.
51. Lamichhaney S, Berglund J, Almén MS, Maqbool K, Grabherr M, Martinez-Barrio A, et al. Evolution of Darwin's finches and their beaks revealed by genome sequencing. *Nature*. 2015; 518(7539):371–375. doi: [10.1038/nature14181](https://doi.org/10.1038/nature14181) PMID: [25686609](https://pubmed.ncbi.nlm.nih.gov/25686609/)
52. Song YS, Hein J. Parsimonious reconstruction of sequence evolution and haplotype blocks. In: *Algorithms in Bioinformatics*. Springer; 2003. p. 287–302.
53. International HapMap 3 Consortium et al. Integrating common and rare genetic variation in diverse human populations. *Nature*. 2010; 467(7311):52–58. doi: [10.1038/nature09298](https://doi.org/10.1038/nature09298) PMID: [20811451](https://pubmed.ncbi.nlm.nih.gov/20811451/)
54. Jeffreys AJ, Kauppi L, Neumann R. Intensely punctate meiotic recombination in the class II region of the major histocompatibility complex. *Nature genetics*. 2001; 29(2):217–222. doi: [10.1038/ng1001-217](https://doi.org/10.1038/ng1001-217) PMID: [11586303](https://pubmed.ncbi.nlm.nih.gov/11586303/)
55. Hinch AG, Tandon A, Patterson N, Song Y, Rohland N, Palmer CD, et al. The landscape of recombination in African Americans. *Nature*. 2011; 476(7359):170–175. doi: [10.1038/nature10336](https://doi.org/10.1038/nature10336) PMID: [21775986](https://pubmed.ncbi.nlm.nih.gov/21775986/)
56. 1000 Genomes Project Consortium et al. A global reference for human genetic variation. *Nature*. 2015; 526(7571):68–74. doi: [10.1038/nature15393](https://doi.org/10.1038/nature15393) PMID: [26432245](https://pubmed.ncbi.nlm.nih.gov/26432245/)
57. Hein J, Schierup M, Wiuf C. *Gene genealogies, variation and evolution: a primer in coalescent theory*. Oxford university press; 2004.
58. Petren K, Grant P, Grant B, Keller L. Comparative landscape genetics and the adaptive radiation of Darwin's finches: the role of peripheral isolation. *Molecular Ecology*. 2005; 14(10):2943–2957. doi: [10.1111/j.1365-294X.2005.02632.x](https://doi.org/10.1111/j.1365-294X.2005.02632.x) PMID: [16101765](https://pubmed.ncbi.nlm.nih.gov/16101765/)
59. Hudson RR. Generating samples under a Wright–Fisher neutral model of genetic variation. *Bioinformatics*. 2002; 18(2):337–338. doi: [10.1093/bioinformatics/18.2.337](https://doi.org/10.1093/bioinformatics/18.2.337) PMID: [11847089](https://pubmed.ncbi.nlm.nih.gov/11847089/)
60. Rambaut A, Grass NC. Seq-Gen: an application for the Monte Carlo simulation of DNA sequence evolution along phylogenetic trees. *Computer applications in the biosciences: CABIOS*. 1997; 13(3):235–238. PMID: [9183526](https://pubmed.ncbi.nlm.nih.gov/9183526/)
61. Li H, Durbin R. Fast and accurate short read alignment with Burrows–Wheeler transform. *Bioinformatics*. 2009; 25(14):1754–1760. doi: [10.1093/bioinformatics/btp324](https://doi.org/10.1093/bioinformatics/btp324) PMID: [19451168](https://pubmed.ncbi.nlm.nih.gov/19451168/)
62. McKenna A, Hanna M, Banks E, Sivachenko A, Cibulskis K, Kernysky A, et al. The Genome Analysis Toolkit: a MapReduce framework for analyzing next-generation DNA sequencing data. *Genome research*. 2010; 20(9):1297–1303. doi: [10.1101/gr.107524.110](https://doi.org/10.1101/gr.107524.110) PMID: [20644199](https://pubmed.ncbi.nlm.nih.gov/20644199/)
63. DePristo MA, Banks E, Poplin R, Garimella KV, Maguire JR, Hartl C, et al. A framework for variation discovery and genotyping using next-generation DNA sequencing data. *Nature genetics*. 2011; 43(5):491–498. doi: [10.1038/ng.806](https://doi.org/10.1038/ng.806) PMID: [21478889](https://pubmed.ncbi.nlm.nih.gov/21478889/)