# QUANTIFYING GENETIC INNOVATION: MATHEMATICAL FOUNDATIONS FOR THE TOPOLOGICAL STUDY OF RETICULATE EVOLUTION 

MICHAEL LESNICK, RAÚL RABADÁN, AND DANIEL I. S. ROSENBLOOM


#### Abstract

A topological approach to the study of genetic recombination, based on persistent homology, was introduced by Chan, Carlsson, and Rabadán in 2013. This associates a sequence of signatures called barcodes to genomic data sampled from an evolutionary history. In this paper, we develop theoretical foundations for this approach.

First, we present a novel formulation of the underlying inference problem. Specifically, we introduce and study the novelty profile, a simple, stable statistic of an evolutionary history which not only counts recombination events but also quantifies how recombination creates genetic diversity. We propose that the (hitherto implicit) goal of the topological approach to recombination is the estimation of novelty profiles.

We then study the problem of obtaining a lower bound on the novelty profile using barcodes. We focus on a low-recombination regime, where the evolutionary history can be described by a directed acyclic graph called a galled tree, which differs from a tree only by isolated topological defects. We show that in this regime, under a complete sampling assumption, the $1^{\text {st }}$ barcode yields a lower bound on the novelty profile, and hence on the number of recombination events. For $i>1$, the $i^{\text {th }}$ barcode is empty. In addition, we use a stability principle to strengthen these results to ones which hold for any subsample of an arbitrary evolutionary history. To establish these results, we describe the topology of the Vietoris-Rips filtrations arising from evolutionary histories indexed by galled trees.

As a step towards a probabilistic theory, we also show that for a random history indexed by a fixed galled tree and satisfying biologically reasonable conditions, the intervals of the $1^{\text {st }}$ barcode are independent random variables. Using simulations, we explore the sensitivity of these intervals to recombination.


## Contents

1. Introduction ..... 2
1.1. Recombination ..... 2
1.2. Novelty Profiles ..... 3
1.3. Prior Topological Work on Recombination ..... 3
1.4. Barcodes as Lower Bounds of Novelty Profiles ..... 4
1.5. Remarks on the Practical Inference of Novelty Profiles ..... 5
1.6. Other Theoretical Work on the Topological Approach to Recombination ..... 6
1.7. Mathematical Contributions ..... 6
Acknowledgements ..... 7
2. Phylogenetic Graphs and Evolutionary Histories ..... 8
3. Novelty Profiles ..... 10
3.1. The Temporal Novelty Profile ..... 10
3.2. The Topological Novelty Profile ..... 11
4. Histories Indexed By Galled Trees ..... 13
5. Topological Preliminaries ..... 15
5.1. Persistent Homology ..... 15
5.2. Stability of the Topological Novelty Profile ..... 18
5.3. Discrete Morse Theory ..... 18
6. Barcodes of Histories Indexed by Galled Trees ..... 19
6.1. Barcodes of Histories indexed by Trees ..... 19
6.2. Metric Decomposition of an Evolutionary History ..... 20
6.3. Vietoris-Rips Filtrations of Almost Linear Metric Spaces ..... 21
6.4. Inference about Recombination from Barcodes ..... 28
7. Relaxing the Complete Sampling and Galled Tree Assumptions ..... 29
7.1. Relaxing the Complete Sampling Assumption ..... 29
7.2. Relaxing the Galled Tree Assumption ..... 29
8. Random Histories Indexed by Galled Trees ..... 30
8.1. Independence of Intervals in the First Barcode ..... 30
8.2. The Barcode of a Random History on a Source-Sink Loop: Numerical Results ..... 34
9. Discussion ..... 37
9.1. Potential Applications of the Novelty Profile ..... 37
9.2. On the Assumptions Underlying our Main Results ..... 37
9.3. Directions for Further Theoretical Work ..... 39
9.4. Practical Strategies for Estimating (Statistics of) Novelty Profiles ..... 39
Appendix A. Probability that the Coalescent with Recombination Generates a Galled Tree ..... 40
A.1. Overview of the Coalescent with Recombination ..... 40
A.2. Probability of Generating a Galled Tree as the Solution of a Linear System ..... 40
Appendix B. Subsamples Rarely Violate our Theoretical Bounds for Complete Samples ..... 43
References ..... 44

## 1. Introduction

1.1. Recombination. Recombination is a process by which the genomes of two parental organisms combine to form a new genome. Like genetic mutation, recombination gives rise to genetic diversity in evolving populations. But unlike mutation, recombination can unite advantageous traits which have arisen in separate lineages, or rescue an advantageous trait from an otherwise disadvantageous genetic background. In these ways, recombination hastens the pace at which adaptive genetic novelty arises.

Evolving populations can be studied by observing genetic sequences obtained from a sample of organisms. Several methods exist to estimate or bound the number of recombination events that have occurred in the ancestry of a sample and to identify the genomic locations where recombination may have occurred [31, 40, 48]. Yet these methods do not reveal how recombination generates genetic diversity: Recombination between two very distinct parents may create a genetically very novel offspring, contributing substantial diversity to the
population, but recombination between genetically similar parents can only create genetically similar offspring, contributing little diversity.
1.2. Novelty Profiles. In this work, we introduce a simple, stable statistic of an evolving population, the novelty profile, which quantifies how recombination contributes to genetic diversity. To define the novelty profile, we first need to select a formal model of an evolving population. We call the model we consider in this paper an evolutionary history. An evolutionary history $E$ is a directed acyclic graph $G$, together with a set $E_{v}$ at each vertex $v$ of $G$, satisfying certain conditions. We call $G$ a phylogenetic graph, and say that $G$ indexes $E$. Each vertex of $G$ represents an organism, each edge of $G$ represents a parental relationship, and each $E_{v}$ specifies the genome of the organism $v$. See Section 2 for the formal definition of an evolutionary history and an illustration.

The novelty profile of an evolutionary history is simply a list of $k$ monotonically decreasing numbers, where $k$ is the number of recombination events in the history. Roughly, each number measures the contribution to genetic diversity of one recombinant. We introduce two versions of this statistic, the temporal and topological novelty profiles. The definition of the temporal novelty profile is very elementary and intuitive, but depends on a specification of the time at which each organism is born. Moreover, the temporal novelty profile, while stable to perturbations (i.e., small changes) of the genomes, is unstable to perturbations of the birth times. In contrast, the topological novelty profile is defined in a way that does not depend on birth times. It is also stable to perturbations of the genomes. The topological novelty profile is a lower bound for the temporal novelty profile, in the sense that the $i^{\text {th }}$ element of the topological novelty profile is less than or equal to the $i^{\text {th }}$ element of the temporal novelty profile for all $i$.
1.3. Prior Topological Work on Recombination. The broader idea of quantifying the scale of recombination events, in addition to their number, is already present in earlier topological work on recombination [8, 9, 13, 24, 25, 43]; the recent textbook [44] provides an detailed introduction. Our definitions of novelty profiles are inspired by some of this previous work, and one of our main objectives here is to use novelty profiles to develop mathematical foundations for that work.

In the previous work, a popular topological data analysis method called persistent homology is used to associate a sequence $\mathcal{B}_{0}(S), \mathcal{B}_{1}(S), \mathcal{B}_{2}(S), \ldots$ of objects called barcodes to an arbitrary sample $S$ of an evolving population. Each barcode is a collection of intervals $[a, b)$ on the real line. In [13], it is shown that, under a standard infinite sites assumption ruling out multiple mutations at the same genetic site, if no recombination occurs in the population's history, then $\mathcal{B}_{i}(S)$ is empty for all $i \geq 1$; see Section 6.1. Hence a non-empty barcode $\mathcal{B}_{i}(S)$ for any $i \geq 1$ certifies that recombination has occurred at some point in the history. Within simulations of evolving populations, the number of intervals in the first barcode $\mathcal{B}_{1}(S)$ has been observed to increase with the simulated recombination rate [9, 13]. Moreover, it has been observed empirically that the endpoints of the intervals in the barcode $\mathcal{B}_{1}(S)$ depend on the genetic scale at which recombination events occur [8, 24, 43]. For instance, in studies of population admixture (i.e., interbreeding between distantly related subpopulations), intervals in $\mathcal{B}_{1}(S)$ with large values for the endpoints have been observed to appear in the barcode only in the presence of admixture $[8,43]$.

While these findings together suggest that the barcode encodes information about both the number of recombination events and the contribution of recombination to genetic diversity, the precise statistical nature of the relationship between barcodes and recombination has not been made clear. In this paper, we make progress towards understanding this relationship.
1.4. Barcodes as Lower Bounds of Novelty Profiles. We propose that the central inference problem implicit in the previous topological work on evolution is the estimation of novelty profiles. Given this, we are led to ask how the barcodes $\mathcal{B}_{i}(S)$ of genomic data studied in the previous work perform as estimators of the novelty profile. It is known that these barcodes can fail to detect recombination events, even in the simplest and most favorable circumstances [13], so one expects the barcodes to encode only partial information about the novelty profile.

In this paper, we study barcodes as lower bounds on the novelty profile. For context, we note that computable lower bounds on numbers of recombination events play a key role in the study of recombination [31, 40, 48]. Our lower bounds are in a similar spirit. Similarly, in topological data analysis, the idea of using barcodes to formulate lower bounds (e.g., on the Gromov-Hausdorff distance between compact metric spaces) is fundamental-it lies at the heart of the well-known stability theory for persistence [5, 16, 17]; see Section 5.

To formulate our bounds, we first restrict attention to a low-recombination regime, where the evolutionary histories are indexed by galled trees. Galled trees are directed acyclic graphs that are almost trees, in a sense: They may have cycles, but these cycles are topologically separated from one another; see Section 4 for the precise definition. Galled trees have received considerable attention in the phylogenetics literature as computationally convenient models of evolution with infrequent recombination [33], [27]. They have been of interest primarily because certain phylogenetic network reconstruction problems that are computationally hard in general admit polynomial-time solutions when restricted to galled trees. To clarify the biological relevance of galled tree models of evolution, in Appendix A we study the probability $P$ that a galled tree correctly models an evolutionary history. We work with a coalescent model of evolution, a standard model in population genetics. We show that for this model, $P$ can be computed by solving a linear system of equations, and we observe that for a fixed population size, $P$ tends to 1 as the recombination rate tends to 0 ; see also Remark 4.6.

We observe that for evolutionary histories indexed by galled trees, the temporal and topological novelty profiles are equal (Proposition 4.5). Our main result relating barcodes to recombination in the galled tree setting is the following (see Theorem 6.25 and the preceding definitions for the precise formulation):

Theorem. Let $\mathcal{E}$ be an evolutionary history indexed by a galled tree.
(i) The set of lengths of intervals in the barcode $\mathcal{B}_{1}(\mathcal{E})$ is a lower bound on the novelty profile. In particular, the number of intervals in $\mathcal{B}_{1}(\mathcal{E})$ is a lower bound on the number of recombination events in $\mathcal{E}$.
(ii) $\mathcal{B}_{i}(\mathcal{E})$ is empty for $i \geq 2$.

Part (i) of the theorem does not hold for barcodes $\mathcal{B}_{1}(S)$ of arbitrary samples $S \subseteq \mathcal{E}$ (Example 6.26). However, using a well-known stability property of persistent homology, we observe that the theorem extends to an approximate version which holds for an arbitrary sample $S$, even in the presence of noise (Corollary 7.1). The quality of the approximation
depends on the similarity of the geometries of $S$ and $\mathcal{E}$, as measured by the Gromov-Hausdorff distance. Along similar lines, the theorem further extends to an approximate version for histories indexed by arbitrary phylogenetic graphs (Corollary 7.3); here, the quality of approximation is controlled by the number of mutations which must be ignored to obtain a history indexed by a galled tree.

These results are deterministic; in cases where the history is sampled at random from a known distribution, one hopes to be able to obtain stronger probabilistic results. As a first step towards such results, we show in Section 8.1 that for a random history indexed by a fixed galled tree $G$ and satisfying a biologically reasonable independence condition, the intervals of the $1^{\text {st }}$ barcode are independent random variables indexed by the recombinants of $G$.

We then study the distributions of these random variables via simulation, for one class of random models of genetic sequence evolution. Our simulation results indicate that even when we have sampled all individuals in the evolutionary history, the barcode is usually a rather loose lower bound on the novelty profile. For example, in the most favorable circumstances, a recombinant of high novelty is detected in our simulations about a third of the time. Nevertheless, the barcodes provide partial information about the novelty profile. Notably, we observe in our simulations that when a recombination event of novelty $n$ is detected by the barcode, the average length of the corresponding interval is approximately $c+d \sqrt{n}$ for constants $c$ and $d$.
1.5. Remarks on the Practical Inference of Novelty Profiles. The primary motivation for our results relating barcodes to novelty profiles is to further our understanding of the prior topological work on recombination. Our results clarify what the barcodes of genomic data do and do not tell us about recombination, and they highlight the mathematical challenges involved in understanding the connection between barcodes and recombination more fully.

That said, we are hopeful that novelty profiles can find practical use in biology applications. For this, we need to be able to infer (statistics of) novelty profiles from real-world genomic data. It is thus natural to ask whether our bounds relating barcodes to novelty profiles can be applied in practice to such inference. In Section 9.2, we consider this question in detail. While we find that such inference may be possible in some circumstances, for typical genomic data the assumptions underlying our bounds seem too strong to apply in a useful way. Thus, we see our bounds as a first step towards a more applicable theory for inferring information about novelty profiles from topological statistics of genomic data. Some possible directions forward along these lines are discussed in Section 9.

While we are indeed hopeful that our bounds can be strengthened to obtain results that are more readily applied in practice, we expect that in the near term, it may be more fruitful for practical applications to pursue alternative approaches to the inference of (statistics of) novelty profiles. For smaller samples, it may be effective to estimate the novelty profile by first estimating the full evolutionary history of the sample; there is well-developed technology for this [27, 45]. For larger samples, where estimation of a full evolutionary history is computationally infeasible, a machine learning approach may be the most practical way forward; for this, it may be possible to adapt ideas from recent work on the learning of recombination rates from topological features [32]. While a full development of these ideas is beyond the scope of this paper, we discuss them briefly in Section 9.4.
1.6. Other Theoretical Work on the Topological Approach to Recombination. Theoretical foundations for the application of persistent homology to recombination have also been studied in recent work of Cámara et al. [8] and Parida et al. [43], though from a rather different angle than ours. [8] considers connections to the problem of constructing minimal ancestral recombination graphs (ARGs) for single-breakpoint models of recombination. (An ARG roughly corresponds to what we call an evolutionary history in this paper; a minimal ARG is a history of a given set of genome sequences with as few recombination events as possible.) In contrast, we do not consider ARG reconstruction or constrain recombination to a single-breakpoint model.

The theory developed in [43] concerns population admixture. The work models the evolution not only of individual organisms, but also of entire populations, and defines barcodes signatures both at the individual level and at the population level. It is shown that under natural assumptions on the inter-population and intra-population genetic distances, one can deduce information about barcodes at the population level from barcodes at the individual level. However, no direct theoretical relationship is established between the barcodes and recombination or admixture.

While our aim and technical approach differ from these previous works, we do share the common goal of understanding persistent homology as a signature of genetic recombination.
1.7. Mathematical Contributions. One key feature of the barcode signatures of recombination studied here is that they depend only on the metric structure on an evolutionary history, i.e., the genetic distances between organisms - in our formalism, the Hamming distance, or monotonic transformations thereof such as the Jukes-Cantor distance. In fact, these barcodes are given by a standard construction which associates barcodes to any finite metric space $M$. In this construction, one first builds a 1-parameter family of simplicial complexes $\mathcal{V}(M)$ called the Vietoris-Rips filtration (VRF).

The topological study of VRFs is a central theoretical problem in topological data analysis. While some fundamental results about VRFs are well known, including a stability theorem [7, 15, 17], relatively little is known about concrete computations of the topology of VRFs, outside of special cases; even for points distributed uniformly on a circle or ellipse, the problem is already non-trivial, and has been the subject of recent research [1, 3].

The result of Chan et al. [13], that $\mathcal{B}_{i}(S)=\emptyset$ for $i \geq 1$ when $S$ is a sample of a history with no recombination, amounts to a proof that the VRF of a tree-like metric space is topologically trivial, up to multiplicity of connected components; see Proposition 6.3. Analogously, the mathematical heart of our main results about recombination for galled trees is a topological description of $\mathcal{V}(\mathcal{E})$, for $\mathcal{E}$ an evolutionary history indexed by a galled tree, regarded as a metric space: We use discrete Morse theory [26] to show that each simplicial complex in $\mathcal{V}(\mathcal{E})$ is homotopy equivalent to a disjoint union of bouquets of circles, where each circle corresponds to a unique recombination event. Moreover, we completely describe the topological behavior of the inclusion maps in $\mathcal{V}(\mathcal{E})$ and give bounds on the number of intervals in $\mathcal{B}_{1}(\mathcal{E})$. For the precise statements, see Proposition 6.12 and Theorems 6.11 and 6.13.

Our topological study of $\mathcal{V}(\mathcal{E})$ hinges on the study of the VRFs of almost linear metric spaces; we say a metric space is almost linear if (up to isometry) it is obtained from a finite subspace of $\mathbb{R}$ by adding a single point. In brief, almost linear metric spaces enter into our analysis in the following way: We observe in Proposition 6.10 that, up to isometry, the metric
space $\mathcal{E}$ can be constructed by iteratively gluing together tree-like metric spaces and almost linear metric spaces using a coproduct construction. (The coproducts are taken in a category of based metric spaces, allowing the basepoint to change.) Moreover, letting $P \vee Q$ denote a coproduct of two based metric spaces $P$ and $Q$, we have that $\mathcal{V}(P \vee Q)$ is, up to homotopy, a wedge sum of $\mathcal{V}(P)$ and $\mathcal{V}(Q)$ (Proposition 6.7). Since the VRFs of tree-like metric spaces are topologically trivial, it follows that to describe the topology of $\mathcal{V}(\mathcal{E})$, it suffices to describe the topology of the VRF of an almost linear metric space; Theorem 6.13 gives such a description.

Outline. For some of the material of this paper, we must assume that the reader is familiar with elementary algebraic topology. However, much of our material on novelty profiles does not require a background in algebraic topology, and we believe that this material may be of independent interest. Thus, we have arranged the paper so that the material on topology appears as late as possible.

Section 2 introduces our mathematical formalism for working with evolving populations in the presence of recombination. Section 3 introduces novelty profiles. Section 4 reviews galled trees and establishes that in the special case of galled trees, the temporal and topological novelty profiles are equal. Section 5 reviews aspects of persistent homology and discrete Morse theory needed in the remainder of our paper, and observes that the topological novelty profile is stable. Section 6.1 briefly reviews the results of Chan et al. on barcodes of evolutionary histories indexed by trees. Section 6.2 studies the VRFs of coproducts of based metric spaces, and Section 6.3 presents our topological analysis of the VRFs of almost linear metric spaces. Using the results of Sections 6.1 to 6.3, Section 6.4 establishes our main deterministic result about the barcodes of evolutionary histories indexed by galled trees. Section 7 applies the stability of persistent homology to extend this result to subsamples of histories indexed by arbitrary phylogenetic graphs.

Section 8.1 establishes that for a suitably chosen random history indexed by a fixed galled tree, the intervals in the $1^{\text {st }}$ persistence barcode are independent random variables. With this as motivation, Section 8.2 uses simulation to study the statistical properties of the barcode of a random history with a single recombination event. Section 9 discusses the applicability of our results and ideas to real-world genomic data, and explores directions for future work.

Two appendices tie our results explicitly to coalescent theory. Appendix A studies the probability that an evolutionary history generated by the coalescent model is a galled tree. Appendix B observes in simulation that, although our main result for histories indexed by galled trees does not hold exactly for arbitrary subsamples, the lower bound on the number of recombination events implied by that result is only rarely violated under subsampling.

Acknowledgements. We thank Ulrich Bauer for helpful discussions about how to prove Theorem 6.13, our main result about the Vietoris-Rips filtrations of almost linear metric spaces. Ulrich provided valuable input about the use of the triangle inequality in that argument, and also suggested the use of the discrete gradient vector field of [35]. We also thank Pablo Cámara and Kevin Emmett for valuable discussions, Matthew Zaremsky for sharing the counterexample of Remark 6.15, and Greg Henselman for helpful feedback on our discussion of discrete Morse theory. Finally, we thank Peter Landweber and the anonymous reviewers for suggestions which helped improve the paper. Lesnick was partially supported by funding from the Institute for Mathematics and its Applications, NIH grants U54CA193313
and T32MH065214, and an award from the J. Insley Blair Pyne Fund. Rabadán and Rosenbloom were funded by NIH grants U54CA193313 and R01GM117591.

## 2. Phylogenetic Graphs and Evolutionary Histories

We now introduce our mathematical formalism for the topological study of reticulate evolution. The formalism is similar to that used elsewhere in the literature on reticulate evolution, though some of our terminology is non-standard; for context, see for example [33] and the references therein.

Definition 2.1 (Phylogenetic Graph). A phylogenetic graph is a finite directed acyclic graph $G$ such that

1. $G$ has a unique vertex $r$, the root, with in-degree 0 ,
2. Each vertex of $G$ has in-degree at most 2 .

We call a vertex in $G$ of in-degree 1 a clone, and a vertex of in-degree 2 a recombinant. If $(v, w)$ is a directed edge in $G$, we say $v$ is a parent of $w$. We define a rooted tree to be a phylogenetic graph with no recombinants.

Fig. 1 illustrates a simple phylogenetic graph.


Figure 1. A phylogenetic graph with 10 vertices: 7 clones (gray), 2 recombinants (black), and the root (white).

For $G$ a rooted directed acyclic graph with vertex set $V$ and $S \subseteq V$, we say $v \in S$ is the minimum of $S$ if for all $s \in S$, any directed path from $r$ to $s$ in $G$ contains $v$. $S$ may not have a minimum element, but if the minimum element exists, it is clearly unique.

Let Set denote the collection of all finite sets.
Definition 2.2 (Evolutionary History). For $G$ a phylogenetic graph with vertex set $V$, an (evolutionary) history indexed by $G$ is a map $\mathcal{E}: V \rightarrow$ Set with the following three properties:

1. If $w$ is a clone with parent $v$, then $\mathcal{E}_{v} \subseteq \mathcal{E}_{w}$.
2. For each $m \in \cup_{v \in V} \mathcal{E}_{v}$, the set $\left\{v \in V \mid m \in \mathcal{E}_{v}\right\}$ has a minimum element.
3. If $w$ is a recombinant with parents $u$ and $v$, then

$$
\mathcal{E}_{u} \cap \mathcal{E}_{v} \subseteq \mathcal{E}_{w} \subseteq \mathcal{E}_{u} \cup \mathcal{E}_{v}
$$

We call the elements of the sets $\mathcal{E}_{v}$ mutations.
Fig. 2 gives an example of a history indexed by the phylogenetic graph of Fig. 1.
Remark 2.3. The biological interpretation of the above definitions is this: A phylogenetic graph describes the ancestral relationships between organisms in a history, and each set $\mathcal{E}_{v}$ specifies the genome of organism $v$, in terms of the difference between that genome and
some fixed (unspecified) reference genome. Properties 1 and 2 are standard in phylogenetics; they specify that each mutation arises only once in the history, and that each clone inherits all the mutations of its parent. Together, these two properties are often referred to as the infinite sites assumption. Property 3 stipulates that if both parents of a recombinant carry a mutation, then the recombinant inherits that mutation, and moreover, any mutation carried by a recombinant is inherited from a parent.

Remark 2.4. As indicated earlier, evolutionary histories are often called ancestral recombination graphs (ARGs) in the literature [27, 45]. However, since we wish to make a clear distinction between the history and its underlying phylogenetic graph, we prefer the terminology presented here.


Figure 2. A history indexed by the phylogenetic graph of Fig. 1.

Definition 2.5 (Symmetric Difference Metric on an Evolutionary History). Define a metric $d$ on finite sets, the symmetric difference metric by taking

$$
d(A, B)=|(A \cup B) \backslash(A \cap B)|
$$

for any finite sets $A, B$. For any history $\mathcal{E}$ indexed by a phylogenetic graph $G$ with vertex set $V$, this restricts to a metric on the set $\left\{\mathcal{E}_{v} \mid v \in V\right\}$. We denote the resulting metric space as met $\mathcal{E}$, or when no confusion is likely, simply as $\mathcal{E}$.

Remark 2.6. It is common in the phylogenetics literature to model genomes as binary vectors, and to metrize a set of genomes using the Hamming distance. It is easy to see that the formalism we've introduced here is essentially equivalent. Under this equivalence, other common phylogenetic distances (e.g., Jukes-Cantor distance, Nei-Tamura distance) correspond to monotonic transformations of the symmetric difference metric $d$. In fact, all the results of this paper formulated in terms of $d$ extend immediately to such monotonic transformations.

Remark 2.7. In real-world evolving populations, the infinite sites assumption, described above, may not always hold. In other words, the same mutation may occur in different organisms despite being absent in their common ancestors. Such mutations, termed homoplasies, may be observed in sampled data either if the per-site mutation rate is high (which is typical
for species with short genomes, such as RNA viruses) or if the mutations confer high fitness. Homoplasies are typically rare for species with long genomes, as the probability of mutating twice at the same exact genetic site is small. If they do occur, homoplasies usually involve few sites, so that the metric space underlying the history differs only slightly from that of a history satisfying the infinite sites assumption.

## 3. Novelty Profiles

3.1. The Temporal Novelty Profile. For $G$ a phylogenetic graph with vertex set $V$, define a partial order on $V$ by taking $v \leq w$ if there is a directed path in $G$ from $v$ to $w$. We say $t: V \rightarrow \mathbb{R}$ is a time function if $t(v)<t(w)$ whenever $v<w$. We interpret $t(v)$ as the birth time of organism $v$.

Definition 3.1 (Temporal Novelty Profile). Given a history $\mathcal{E}$ indexed by $G$, a time function $t: V \rightarrow \mathbb{R}$, and a recombinant $r$ of $G$, we define $\mathcal{N}(r, t)$, the temporal novelty of $r$, by

$$
\mathcal{N}(r, t):=\min \left\{d\left(\mathcal{E}_{v}, \mathcal{E}_{r}\right) \mid t(v)<t(r)\right\} .
$$

We define $\mathcal{N}(\mathcal{E}, t)$, the temporal novelty profile of $\mathcal{E}$ (with respect to $t$ ) to be the list of temporal novelties $\mathcal{N}(r, t)$, for all recombinants $r$ of $G$, sorted in decreasing order.

Example 3.2. Fig. 3 illustrates novelty profiles for two histories indexed by the same simple phylogenetic graph. For any time function on the history shown in Fig. 3a, the unique recombinant has temporal novelty 1 , so the temporal novelty profile of this history is the single-element list (1). Similarly, for any time function on the history shown in Fig. 3b, the temporal novelty profile is the single-element list (6).


$$
\{a, b, c, d, e, f, G\}
$$

(a) The temporal novelty profile of the history shown is the single-element list (1), for any time function. The topological novelty profile is the same.

(b) The temporal novelty profile of the history shown is the single-element list (6), for any time function. The topological novelty profile is the same.

Figure 3

Example 3.3. Fig. 4 illustrates a history where two recombinants have the same parents. For the time function shown, the novelty profile is $(5,1)$. The small second entry reflects the fact the two recombinant genomes are genetically close to one another. Exchanging the time values of the bottom-most two vertices yields another time function for this history, for which the temporal novelty profile is $(6,1)$. If we take the time values of the bottom-most two vertices to both be 3 , then the temporal novelty profile is $(6,5)$.


Figure 4. The temporal novelty profile of the history and time function shown is $(5,1)$; the topological novelty profile is the same.

Remark 3.4 (Stability). Suppose we are given histories $\mathcal{E}$ and $\mathcal{E}^{\prime}$ indexed by the same phylogenetic graph $G$ with $\left|d\left(\mathcal{E}_{v}, \mathcal{E}_{w}\right)-d\left(\mathcal{E}_{v}^{\prime}, \mathcal{E}_{w}^{\prime}\right)\right| \leq \epsilon$ for all vertices $v, w$ of $G$. Note that this condition holds by the triangle inequality if $d\left(\mathcal{E}_{v}, \mathcal{E}_{v^{\prime}}\right) \leq \frac{\epsilon}{2}$ for all $v \in G$. Let $t$ be any time function on the vertices of $G$. We then have that

$$
d_{\infty}\left(\mathcal{N}(\mathcal{E}, t), \mathcal{N}\left(\mathcal{E}^{\prime}, t\right)\right) \leq \epsilon
$$

where for vectors $A$ and $B$ of the same length, $d_{\infty}(A, B):=\max _{i}\left|A_{i}-B_{i}\right|$. Thus, the temporal novelty profile is stable with respect to genomic perturbations.

However, the temporal novelty profile is unstable with respect to perturbations of the time function. For example, consider the history $\mathcal{E}$ of Example 3.3 (Fig. 4). For $\delta \in(-1, \infty)$, let $t_{\delta}$ be the time function obtained from the time function $t$ shown in Fig. 4 by changing the time value of the vertex on the bottom right from 4 to $3+\delta$. Then for all $\delta \in(0,1)$, we have

$$
\mathcal{N}\left(\mathcal{E}, t_{\delta}\right)=(5,1), \quad \mathcal{N}\left(\mathcal{E}, t_{-\delta}\right)=(6,1),
$$

and therefore

$$
d_{\infty}\left(\mathcal{N}\left(\mathcal{E}, t_{\delta}\right), \mathcal{N}\left(\mathcal{E}, t_{-\delta}\right)\right)=1
$$

whereas stability would require that this distance approach 0 as $\delta$ approaches 0 .

### 3.2. The Topological Novelty Profile.

Definition 3.5 (Relative Minimum Spanning Tree). Given a weighted graph $G$ and a forest $F \subseteq G$ (i.e., a vertex-disjoint collection of subtrees), we define a spanning tree of $G$ rel $F$ simply to be a spanning tree $T$ of $G$ containing $F$. We say $T$ is a minimum spanning tree of $G$ rel $F$ if the sum of the edge weights of $T$ is as small as possible, among all spanning trees of $G$ rel $F$.

Note that by collapsing each tree in $F$ to a point, the problem of finding a minimum spanning tree rel $F$ is equivalent to the standard problem of finding an ordinary minimum spanning tree on a multigraph. (A multigraph is a graph which is allowed to have multiple edges between pairs of vertices.) Thus, all the basic facts about minimum spanning trees have analogues for relative minimum spanning trees. For example, we have the following:

Proposition 3.6. A spanning tree $T$ rel $F$ is minimum if and only if for all $i$, the $i^{\text {th }}$ smallest edge weight is less than or equal to the $i^{\text {th }}$ smallest edge weight in any other spanning tree rel $F$.

Proof. By the remarks above, it suffices to establish the result for ordinary spanning trees, i.e., the case where $F$ is the empty forest. Let $T$ be a minimum spanning tree, and let $U$ be any other spanning tree. To arrive at a contradiction, assume that for some $i$, the $i^{\text {th }}$ smallest edge weight in $T$ is greater than the $i^{\text {th }}$ smallest edge weight in $U$. Let $w$ denote the latter weight. Consider the subforests $T_{w} \subset T$ and $U_{w} \subset U$ consisting of all vertices and just those edges of weight at most $w . U_{w}$ contains more edges than $T_{w}$, so there exists a pair of vertices $(u, v)$ that lie in the same component of $U_{w}$ but not in the same component of $T_{w}$. In fact, there must exist some edge $e=(x, y)$ along the path from $u$ to $v$ in $U_{w}$ such that $x$ and $y$ lie in different components of $T_{w}$. Clearly, the path from $x$ to $y$ in $T$ must contain at least one edge $e^{\prime}$ with weight greater than $w$. Replacing $e^{\prime}$ with $e$ in $T$ gives a new spanning tree with strictly smaller weight than $T$, contradicting that $T$ is a minimum spanning tree.

Remark 3.7. It follows from Proposition 3.6 that the collection of edge weights in a relative minimum spanning tree is independent of the choice of the tree.

Definition 3.8 (Topological Novelty Profile). For $\mathcal{E}$ a history indexed by a phylogenetic graph $G$, let $F^{G}$ be the forest in $G$ obtained by removing all edges pointing to recombinants. Let $\bar{G}$ denote the complete graph with same vertex set as $G$. Regard $\bar{G}$ as a weighted graph by taking the weight of edge $(u, v)$ to be $d\left(\mathcal{E}_{u}, \mathcal{E}_{v}\right)$.

Let $T$ be a minimum spanning tree of $\bar{G}$ rel $F^{G}$. We define $\mathcal{T}(\mathcal{E})$, the topological novelty profile of $\mathcal{E}$, to be the list of distances

$$
\left\{d\left(\mathcal{E}_{u}, \mathcal{E}_{v}\right) \mid(u, v) \in T \backslash F^{G}\right\}
$$

counted with multiplicity and sorted in descending order. By Remark 3.7, $\mathcal{T}(\mathcal{E})$ does not depend on the choice of $T$.

We will observe in Section 5.2 that the topological novelty profile has an interpretation in terms of persistent homology.

Given two lists of numbers $A$ and $B$, each sorted in decreasing order, we write $A \leq B$ if $|A| \leq|B|$ and for each $i \in\{1, \ldots,|A|\}, A_{i} \leq B_{i}$.

Proposition 3.9. For any history $\mathcal{E}$ with time function $t$,

$$
\mathcal{T}(\mathcal{E}) \leq \mathcal{N}(\mathcal{E}, t)
$$

That is, the topological novelty profile is a lower bound for the temporal novelty profile.
Proof. Suppose $\mathcal{E}$ is indexed by $G$. We construct a spanning tree $T$ of $\bar{G}$ rel $F^{G}$ such the weights of edges in $T \backslash F^{G}$ correspond to the temporal novelty profile. The result then follows from Proposition 3.6.

To construct $T$, for each recombinant $r \in G$, choose a vertex $v(r)$ in $G$ with $t(v(r))<t(r)$, such that $d\left(\mathcal{E}_{v(r)}, \mathcal{E}_{r}\right)$ is as small as possible among all such vertices. We take $T$ to be the graph obtained from $F^{G}$ by adding in the edge $(v(r), r)$ for each recombinant $r$. It is easy to check that $T$ is in fact a tree.

Example 3.10. For the histories of Fig. 3a and Fig. 3b, the topological novelty profile is equal to the temporal one for all time functions. For the history and time function of Fig. 4, the topological and temporal novelty profiles are also equal, but one can select a different time function so that the two novelty profiles are not equal.

Example 3.11. Fig. 5a illustrates a history for which the temporal and topological novelty profiles are unequal for any choice of time function. The topological novelty profile is $(1,1)$, whereas the temporal novelty profile is always $(2,1)$. This example is degenerate, in the sense that the same genome (the empty one) appears at multiple vertices; Fig. 5b shows a variant of the example without this degeneracy.

(a) A history for which the topological and temporal novelty profiles are different for any time function.

(b) A variant of example (a) with no duplicate genomes. Each boldface letter represents three mutations, while each letter in plain font represents a single mutation.

Figure 5

Like temporal novelty profiles, topological novelty profiles are stable with respect to perturbations of the genomic data; we show this in Proposition 5.7. Proposition 4.5 below tells us that when $G$ is a galled tree, the temporal and topological novelty profiles are in fact equal.

## 4. Histories Indexed By Galled Trees

Our main bounds on novelty profiles concern the special case that our phylogenetic graphs are galled trees.

The definition of galled tree we give is equivalent to the one given in [33, Definition 6.11.1]. As noted in [33], this is slightly more general than the original definition [28, 53], which requires the cycles in a galled tree to be node-disjoint.

Definition 4.1 (Source-Sink Loop). We say an undirected graph is a loop if its geometric realization is homeomorphic to a circle. We call a directed graph $G$ a source-sink loop if

1. The undirected graph underlying $G$ is a loop.
2. $G$ has a unique source and unique sink.

Definition 4.2 (Sum of Directed Graphs). For directed graphs $G$ and $H$, with $v$ a source in $G$ and $w$ any vertex in $H$, we define a directed graph $G \vee_{v, w} H$ by taking the disjoint union of $G$ and $H$ and then identifying $v$ and $w$ (i.e., "gluing" $v$ to $w$ ). We call $G \vee_{v, w} H$ a sum of $G$ and $H$. (We do not define the sum $G \vee_{v, w} H$ in the case that neither of the vertices $v$ or $w$ is a source.) We will sometimes write $G \vee_{v, w} H$ simply as $G \vee H$, suppressing $v$ and $w$.

Definition 4.3 (Galled Tree). Let $\mathcal{A}$ be the smallest collection of directed acyclic graphs such that:

1. Each rooted tree is in $\mathcal{A}$.
2. Each source-sink loop is in $\mathcal{A}$.
3. if $G$ and $H$ are in $\mathcal{A}$, then so is each sum $G \vee H$.

We define a galled tree to be a graph isomorphic to one in $\mathcal{A}$. Thus, informally, a galled tree is a graph obtained by iteratively gluing rooted trees and source-sink loops along single vertices, using the sum operation specified above.

We omit the easy proof of the following:
Proposition 4.4. Any galled tree is a phylogenetic graph.
Note that the recombinants in a galled tree $G$ are in bijective correspondence with the source-sink loops in $G$.

Fig. 6 gives an example of a galled tree. It can be checked that the phylogenetic graph of Fig. 1 is not a galled tree.


Figure 6. A galled tree which can be constructed as the iterated sum of four rooted trees (solid edges) and two source-sink loops (dashed edges). The two recombinants are shown in black.

Proposition 4.5 (Equality of Temporal and Topological Novelty Profiles on Galled Trees). For any history $\mathcal{E}$ indexed by a galled tree and time function $t$,

$$
\mathcal{N}(\mathcal{E}, t)=\mathcal{T}(\mathcal{E})
$$

Proof. Suppose $\mathcal{E}$ is indexed by the galled tree $G$. We use the notation from Definition 3.8. As in the case of ordinary minimum spanning trees, a minimum spanning tree of $\bar{G}$ rel $F$ can be constructed greedily, by considering the edges of $\bar{G} \backslash F$ in order of increasing weight. In this construction, each edge of $\bar{G} \backslash F$ added to the relative minimum spanning tree can be chosen to connect a recombinant $r$ to a vertex $v$ of the source-sink loop in $G$ that has $r$ as its sink. We then have that $t(v)<t(r)$ and $d\left(\mathcal{E}_{v}, \mathcal{E}_{r}\right) \leq d\left(\mathcal{E}_{w}, \mathcal{E}_{r}\right)$ for any other vertex $w$ with $t(w)<t(r)$. Clearly, in this construction we never take $r$ to be the same recombinant more than once. The result follows.

Remark 4.6 (Galled Trees as Models for Evolution in the Low-Recombination Limit). Given a probabilistic model generating a phylogenetic graph, one may ask what the probability is of
obtaining a galled tree. This problem has previously been studied by simulation in [4], for a coalescent model of evolution. In Appendix A, we study the same problem analytically. We show that the problem reduces to the study of a finite-state Markov chain. A simple analysis of this Markov chain yields, for fixed population size $n$, a system of linear equations $L(\rho)$ depending on a recombination rate parameter $\rho$, whose solution gives the probability $P(n, \rho)$ of obtaining a galled tree. Solving these linear systems numerically for various values of $\rho$ and $n$, we observe that as $\rho$ tends to $0, P(n, \rho)$ tends to 1 .

This indicates that histories indexed by galled trees are biologically reasonable models of evolution in low-recombination settings. While, from a biological standpoint, the specific bounds on $\rho$ needed to obtain a galled tree with high probability are rather stringent in general, we do expect these bounds to hold in some settings of interest; see Section 9 for further discussion of this.

Regardless, from a mathematical perspective, the special case of galled trees seems to be a natural place to begin fleshing out theoretical foundations for the topological study of evolution.

## 5. Topological Preliminaries

In this section, we briefly review persistent homology and the related topological definitions and results we will need in the remainder of the paper. As a first application, we observe that the topological novelty profile admits a description in terms of persistent homology, and is therefore stable. We also briefly review some ideas from discrete Morse theory.

We assume that the reader is familiar with some standard concepts from elementary algebraic topology, including simplicial complexes, homology, and homotopy equivalence. Good introductions can be found in many places, e.g., [30, 39].
5.1. Persistent Homology. Our treatment of persistent homology will be terse; for a more thorough introduction to these ideas, including a discussion of some of the many applications of persistent homology to data analysis, see the surveys and textbooks [10, 11, 22, 42].

Filtrations. A filtration is a collection of topological spaces $\left\{\mathcal{F}_{r}\right\}_{r \in[0, \infty)}$ such that $\mathcal{F}_{r} \subseteq \mathcal{F}_{s}$ whenever $r \leq s$. A morphism $f: \mathcal{F} \rightarrow \mathcal{G}$ of filtrations is a collection of continuous maps $\left\{f_{r}: \mathcal{F}_{r} \rightarrow \mathcal{G}_{r}\right\}_{r \in[0, \infty)}$ such that the following diagram commutes for all $r \leq s$ :


We say $f$ is an objectwise homotopy equivalence if each $f_{r}$ is a homotopy equivalence. Intuitively, if two filtrations are connected by an objectwise homotopy equivalence, we should think of them as topologically equivalent; for further discussion of this point in the context of topological data analysis, see [7].


Figure 7. Rips complexes $\mathcal{V}(P)_{r}$ on a simple point cloud $P \subset \mathbb{R}^{2}$, for several choices of $r$.

Vietoris-Rips Filtrations. For $S$ a simplicial complex, we use square brackets to denote simplices of $S$. Thus, for example, the set of simplices of a triangle with vertex set $\{a, b, c\}$ is $\{[a],[b],[c],[a b],[b c],[a c],[a b c]\}$.

For $P$ a finite metric space and $r \in[0, \infty)$, the Vietoris-Rips complex of $P$ with scale parameter $r$, denoted $\mathcal{V}(P)_{r}$, is the simplicial complex with vertices $P$ that contains simplex $\left[p_{1}, p_{2}, \ldots, p_{n}\right]$ if and only if diameter $\left\{p_{1}, p_{2}, \ldots, p_{n}\right\} \leq 2 r$. If $r \leq s$, then $\mathcal{V}(P)_{r} \subseteq \mathcal{V}(P)_{s}$, so $\mathcal{V}(P):=\left\{\mathcal{V}(P)_{r}\right\}_{r \in[0, \infty)}$ is a filtration; see Fig. 7.

Persistence Modules. A persistence module $M$ consists of a collection of vector spaces $\left\{M_{r}\right\}_{r \in[0, \infty)}$, together with a collection of linear maps $\left\{M_{r, s}: M_{r} \rightarrow M_{s}\right\}_{r \leq s}$ such that

1. for all $r \leq s \leq t$ the following diagram commutes:

2. $M_{r, r}=\mathrm{id}_{M_{r}}$ for all $r$.

We say $M$ is pointwise finite dimensional (p.f.d.) if $\operatorname{dim} M_{r}<\infty$ for all $r$.
Similar to the definition for filtrations, a morphism $f: M \rightarrow N$ of persistence modules is a collection of linear maps $\left\{f_{r}: M_{r} \rightarrow N_{r}\right\}_{r \in[0, \infty)}$ such that for all $r \leq s$, the following diagram commutes:


We say $f$ is an isomorphism if each of the maps $f_{r}$ is an isomorphism.
Direct Sums of Persistence Modules. We assume that the reader is familiar with the definition of the direct sum of vector spaces from linear algebra. For linear maps $f: V_{1} \rightarrow W_{1}$ and $g: V_{2} \rightarrow W_{2}$, we define the direct sum

$$
f \oplus g: V_{1} \oplus V_{2} \rightarrow W_{1} \oplus W_{2}
$$

by taking $f \oplus g(v, w)=(f(v), g(w))$. We then define the sum $M \oplus N$ to be the persistent module given by

$$
(M \oplus N)_{r}=M_{r} \oplus N_{r}, \quad(M \oplus N)_{r, s}=M_{r, s} \oplus N_{r, s}
$$

We can define the direct sum of an arbitrary collection of persistence modules in the same way.

Reduced Homology. Fix a field $K$. (For example, we can take $K=\mathbb{Q}$, or $K=\mathbb{Z}_{2}$, the field with two elements.) For $i \geq 0$, let $\tilde{H}_{i}$ denote the $i^{\text {th }}$ reduced singular homology functor with coefficients in $K$. Thus, $\tilde{H}_{i}$ maps each topological space $S$ to a $K$-vector space $\tilde{H}_{i}(S)$, and maps each continuous function $f: S \rightarrow T$ to a linear map $f_{*}: \tilde{H}_{i}(S) \rightarrow \tilde{H}_{i}(T)$. Applying $\tilde{H}_{i}$ to each space and each inclusion map in a filtration $\mathcal{F}$ gives us a persistence module $\tilde{H}_{i}(\mathcal{F})$. Moreover, a morphism of filtrations $f: \mathcal{F} \rightarrow \mathcal{G}$ induces an morphism $f_{*}: \tilde{H}_{i}(\mathcal{F}) \rightarrow \tilde{H}_{i}(\mathcal{G})$.

Lemma 5.1. If a morphism of filtrations $f: \mathcal{F} \rightarrow \mathcal{G}$ is an objectwise homotopy equivalence, then for any $i \geq 0, f_{*}: \tilde{H}_{i}(\mathcal{F}) \rightarrow \tilde{H}_{i}(\mathcal{G})$ is an isomorphism.

Proof. It is a standard fact that if a continuous map $g$ is a homotopy equivalence, then $\tilde{H}_{i}(g)$ is an isomorphism. This gives the result.

Barcodes. We say $\mathcal{J} \subseteq \mathbb{R}$ is an interval if $\mathcal{J}$ is nonempty and connected. For $\mathcal{J}$ an interval, define the interval module $I^{\mathcal{J}}$ to be the persistence module such that

$$
I_{r}^{\mathcal{J}}=\left\{\begin{array}{lll}
K & \text { if } r \in \mathcal{J}, & I_{r, s}^{\mathcal{J}}= \begin{cases}\operatorname{id}_{K} & \text { if } r \leq s \in \mathcal{J} \\
0 & \text { otherwise } .\end{cases} \\
0 & \text { otherwise }
\end{array}\right.
$$

Theorem 5.2 (Structure of Persistence Modules [19]). If $M$ is a p.f.d. persistence module, then there exists a unique collection of intervals $\mathcal{B}_{M}$ such that

$$
M \cong \oplus_{\mathcal{J} \in \mathcal{B}_{M}} I^{\mathcal{J}}
$$

We call $\mathcal{B}_{M}$ the barcode of $M$. For $\mathcal{F}$ a filtration, we write $\mathcal{B}_{\tilde{H}_{i}(\mathcal{F})}$ simply as $\mathcal{B}_{i}(\mathcal{F})$. Similarly, for $P$ a finite metric space, we write $\mathcal{B}_{i}(\mathcal{V}(P))$ simply as $\mathcal{B}_{i}(P)$.

Remark 5.3. As it is defined using reduced homology, $\mathcal{B}_{0}(\mathcal{F})$ differs from the $0^{\text {th }}$ barcode constructed using unreduced homology by the removal of an infinite length interval.

Definition 5.4 (Gromov-Hausdorff Distance). Given two subspaces $P, Q$ of a metric space $Z$, we define the Hausdorff distance between $P$ and $Q$, by

$$
d_{H}(P, Q):=\max \left\{\sup _{p \in P} \inf _{q \in Q} d(p, q), \sup _{q \in Q} \inf _{p \in P} d(p, q)\right\} .
$$

For $P$ and $Q$ any compact metric spaces, define $d_{G H}(P, Q)$, the Gromov-Hausdorff distance between $P$ and $Q$, to be the infimum of $d_{H}(\gamma(P), \kappa(Q))$ over all isometric embeddings $\gamma: P \rightarrow Z, \kappa: Q \rightarrow Z$ into a third metric space $Z$.

The following stability result is well known, and plays a central role in topological data analysis. Let $d_{B}$ denote the bottleneck distance on persistence barcodes, as defined for example in [17].

Theorem 5.5 (Stability of Persistent Homology [15, 17, 18]). For any finite metric spaces $P, Q$ and $i \geq 0$,

$$
d_{B}\left(\mathcal{B}_{i}(P), \mathcal{B}_{i}(Q)\right) \leq d_{G H}(P, Q)
$$

The following variant of Theorem 5.5, which appears in slightly different language in [29], can be proven by a slight modification of the proof of Theorem 5.5.

Theorem 5.6 (Stability for a Metric Subspace [29, Proposition 5.6]). For finite metric spaces $P \subseteq Q$ and $i \geq 0$,

$$
d_{B}\left(\mathcal{B}_{i}(P), \mathcal{B}_{i}^{S}(Q)\right) \leq \frac{1}{2} d_{H}(P, Q)
$$

where $\mathcal{B}_{i}^{S}(Q)$ is the barcode obtained by shifting each interval of $\mathcal{B}_{i}(Q)$ to the right by $\frac{1}{2} d_{H}(P, Q)$.
5.2. Stability of the Topological Novelty Profile. For $\mathcal{E}$ a history indexed by $G$ and $F^{G} \subseteq G$ the forest of Definition 3.8, define a filtration $\mathcal{F}$ by $\mathcal{F}_{r}:=F^{G} \cup \mathcal{V}(\mathcal{E})_{r}$. (In the expression $\mathcal{V}(\mathcal{E}), \mathcal{E}$ is understood to denote the metric space of Definition 2.5.) It's easy to check that $\mathcal{T}(\mathcal{E})$, the topological novelty profile of $\mathcal{E}$, is exactly the list of right endpoints of intervals in $\mathcal{B}_{0}(\mathcal{F})$, possibly with some copies of 0 added in.

We can use this description to obtain a simple stability result for topological novelty profiles analogous to the one for temporal novelty profiles mentioned in Remark 3.4:

Proposition 5.7. Given histories $\mathcal{E}$ and $\mathcal{E}^{\prime}$ indexed by the same phylogenetic graph $G$ with $\left|d\left(\mathcal{E}_{v}, \mathcal{E}_{w}\right)-d\left(\mathcal{E}_{v}^{\prime}, \mathcal{E}_{w}^{\prime}\right)\right| \leq \epsilon$ for all vertices $v, w$ of $G$, we have

$$
d_{\infty}\left(\mathcal{T}(\mathcal{E}), \mathcal{T}\left(\mathcal{E}^{\prime}\right)\right) \leq \epsilon
$$

Proof. This follows immediately from a generalized version of the stability theorem for persistent homology, as described in [14], [16], and [5].
5.3. Discrete Morse Theory. The proof of our main results relies on discrete Morse theory (DMT), a well known combinatorial theory concerning topology-preserving collapses of cell complexes. We will not need the full strength of standard DMT; we review only what we need for our proof. See [26] for a detailed introduction to DMT.

Recall that for $G$ a graph with no self-edges $(v, v)$, a matching $X$ in $G$ is a subset of the edges of $G$ such that no two edges in $X$ are incident to the same vertex. For $S$ a simplicial complex, the Hasse graph $G_{S}$ of $S$ is the directed graph with vertices the simplices of $S$ and an edge from $s$ to $s^{\prime}$ if and only if $s^{\prime}$ is a codimension-1 face of $s$. A matching $X$ in $G_{S}$ is said to be acyclic if when we modify the graph $G_{S}$ by reversing the orientation of all edges in $X$, while leaving the orientation of all other edges unchanged, we obtain a directed acyclic graph.

A discrete gradient vector field $(D G V F) X$ on $S$ is an acyclic matching in $G_{S}$. A simplex $\sigma \in S$ is called critical in $X$ if $\sigma$ is not matched in $X$.

The acyclicity condition admits an alternative formulation which is often convenient. Given a matching $X$ in $G_{S}$, we define an $X$-path to be a sequence of simplices in $S$

$$
\sigma_{0}, \tau_{0}, \sigma_{1}, \tau_{1}, \ldots, \sigma_{m}, \tau_{m}, \sigma_{m+1}
$$

such that for each $j \in\{0, \ldots, m\}$, the following are true:

- $\sigma_{j}$ is a face of $\tau_{j}$ and $X$ matches $\sigma_{j}$ to $\tau_{j}$,
- $\sigma_{j+1}$ is a codimension- 1 face of $\tau_{j}$,
- $\sigma_{j} \neq \sigma_{j+1}$.

We say the $X$-path is a non-trivial if $m \geq 0$, and closed if $\sigma_{0}=\sigma_{m+1}$.

Proposition 5.8 ([26, Theorem 6.2]). A matching $X$ in $G_{S}$ is a $D G V F$ if and only if there exists no non-trivial closed $X$-path.

The following is one of the basic results of discrete Morse theory:
Proposition 5.9 ([37, Theorem 11.13]).
(i) Suppose that $X$ is a DGFV on a finite simplicial complex $S$. Then $S$ is homotopy equivalent to a $C W$-complex with exactly one cell of dimension $i$ for each critical $i$-simplex of $X$.
(ii) If the critical simplices of $X$ form a subcomplex $S^{\prime} \subseteq S$, then in fact $S$ deformation retracts onto $S^{\prime}$.

## 6. Barcodes of Histories Indexed by Galled Trees

The topological novelty profile and $0^{\text {th }}$ persistence barcode of an evolutionary history are closely related by the following result, whose easy verification we leave to the reader:

For $\mathcal{B}$ a barcode, let lengths $(\mathcal{B})$ denote the list of lengths of intervals of $\mathcal{B}$, sorted in descending order.

Proposition 6.1. Suppose we are given a history $\mathcal{E}$ and $\delta>0$ such that $d\left(\mathcal{E}_{v}, \mathcal{E}_{w}\right)<\delta$ whenever $w$ is a clone with parent $v$. Then the lists obtained from lengths $\left(\mathcal{B}_{0}(\mathcal{E})\right)$ and $\mathcal{T}(\mathcal{E})$ by removing all entries less than $\delta$ are equal.

This suggests that in some cases, $0^{\text {th }}$ barcodes may be useful in the study of recombination. However, in cases where the minimum $\delta$ satisfying the condition of Proposition 6.1 is large, or where we only have a subsample of the history, $0^{\text {th }}$ barcodes may not offer useful information. This, together with the earlier theoretical result of Chan, Rabadan and Carlsson relating recombination to higher persistence barcodes (Theorem 6.4 below), motivates us to consider the relationship between the topological novelty profile and the higher barcodes of a history.

In this section, we present our main result relating barcodes to novelty profiles in the galled tree setting (Theorem 6.25). The technical heart of our proof is a topological description of the Vietoris-Rips filtration of an almost linear metric space, which we give in Section 6.3. Our arguments rely heavily on discrete Morse theory.
6.1. Barcodes of Histories indexed by Trees. We first review the key result of Chan et al. on the barcodes of histories indexed by trees.
Definition 6.2 (Tree-Like Metric Space). We call an undirected tree with a non-negative weight function on its edges a weighted tree. A metric space $P$ is called tree-like if it is isometric to a subspace of a metric space arising from the shortest-path metric on a weighted tree.

Proposition 6.3 ([13, Supplementary Information, Theorem 2.1]). If $P$ is a tree-like metric space, then for all $r \in[0, \infty)$, each component of $\mathcal{V}(P)_{r}$ is contractible. Hence, $\mathcal{B}_{i}(P)=\emptyset$ for $i \geq 1$.

In [13], only the part of Proposition 6.3 about triviality of barcodes is stated, and not the stronger contractibility result. However, the contractibility result follows immediately from the proof given in [13], using the nerve theorem [30, Chapter 4.G] in place of a Mayer-Vietoris argument.

The following result, due to Chan et al., makes precise the idea that for $i \geq 1$, a non-empty barcode $\mathcal{B}_{i}(S)$ serves as a certificate that recombination is present in the history from which $S$ was sampled.
Theorem 6.4 ([13]). If $G$ is a tree, $\mathcal{E}$ is a history indexed by $G$, and $S \subseteq \mathcal{E}$, then $\mathcal{B}_{i}(S)=\emptyset$ for $i \geq 1$.
Proof. If $S$ is a subset of a history indexed by a tree, then met $S$ is easily seen to be tree-like. Hence, the result follows from Proposition 6.3.

Remark 6.5. In the absence of recombination, homoplasies (recurrent mutations that violate the infinite sites assumption) can lead to a metric space that is not tree-like. However, as indicated in Remark 2.7, a small number of homoplasies causes a correspondingly small deviation from tree-likeness (with respect to Gromov-Hausdorff distance). A single recombination event, on the other hand, can yield a metric space that is arbitrarily far from a tree-like one.
6.2. Metric Decomposition of an Evolutionary History. Define a based metric space simply to be a metric space $P$, together with a choice of basepoint $p \in P$.
Definition 6.6 (Sum of Based Metric Spaces). For based metric spaces $P$ and $Q$ with basepoints $p \in P, q \in Q$, we regard the wedge sum $P \vee Q$ as a metric space, with the metric given by

$$
d_{P \vee Q}(x, y)= \begin{cases}d_{P}(x, y) & \text { if } x, y \in P \\ d_{Q}(x, y) & \text { if } x, y \in Q \\ d_{P}(x, p)+d_{Q}(q, y) & \text { if } x \in P, y \in Q\end{cases}
$$

For based metric spaces $P$ and $Q$, let $\mathcal{V}(P) \vee \mathcal{V}(Q)$ denote the wedge sum filtration, given by

$$
(\mathcal{V}(P) \vee \mathcal{V}(Q))_{r}:=\mathcal{V}(P)_{r} \vee \mathcal{V}(Q)_{r}
$$

Proposition 6.7. For finite based metric spaces $P$ and $Q$, the inclusion

$$
\mathcal{V}(P) \vee \mathcal{V}(Q) \hookrightarrow \mathcal{V}(P \vee Q)
$$

is an objectwise homotopy equivalence. In particular, for any $i \geq 0$,

$$
\mathcal{B}_{i}(P \vee Q)=\mathcal{B}_{i}(P) \cup \mathcal{B}_{i}(Q) .
$$

Proof. We give a proof using discrete Morse theory. For $r \in[0, \infty)$, if $\sigma$ is a simplex in $\mathcal{V}(P \vee Q)_{r}$ containing vertices in both $P$ and $Q$ but not the common vertex $p=q$, then the simplex $\{p=q\} \cup \sigma$ is clearly also in $\mathcal{V}(P \vee Q)_{r}$. We define a DGVF on $\mathcal{V}(P \vee Q)_{r}$ by matching each such simplex $\sigma$ to $\{p=q\} \cup \sigma$. It is clear that this matching is acyclic, hence indeed gives a well-defined DGVF whose set of critical simplices is $(\mathcal{V}(P) \vee \mathcal{V}(Q))_{r}$. Thus, by Proposition 5.9 (ii), the inclusion

$$
(\mathcal{V}(P) \vee \mathcal{V}(Q))_{r} \rightarrow \mathcal{V}(P \vee Q)_{r}
$$

is a homotopy equivalence.
To check that

$$
\mathcal{B}_{i}(P \vee Q)=\mathcal{B}_{i}(P) \cup \mathcal{B}_{i}(Q),
$$

note that by Lemma 5.1,

$$
\mathcal{B}_{i}(P \vee Q)=\mathcal{B}_{i}(\mathcal{V}(P) \vee \mathcal{V}(Q)),
$$

so it suffices to check that

$$
\mathcal{B}_{i}(\mathcal{V}(P) \vee \mathcal{V}(Q))=\mathcal{B}_{i}(P) \cup \mathcal{B}_{i}(Q)
$$

A standard result on the homology of wedge sums of topological spaces [30, Corollary 2.25] furnishes isomorphisms of vector spaces

$$
\tilde{H}_{i}(\mathcal{V}(P) \vee \mathcal{V}(Q))_{r} \rightarrow \tilde{H}_{i}(\mathcal{V}(P))_{r} \oplus \tilde{H}_{i}(\mathcal{V}(Q))_{r}
$$

for each $r \in[0, \infty)$, and these isomorphisms are natural, i.e., they assemble into an isomorphism of persistence modules

$$
\tilde{H}_{i}(\mathcal{V}(P) \vee \mathcal{V}(Q)) \rightarrow \tilde{H}_{i}(\mathcal{V}(P)) \oplus \tilde{H}_{i}(\mathcal{V}(Q)) .
$$

This implies that

$$
\mathcal{B}_{i}(\mathcal{V}(P) \vee \mathcal{V}(Q))=\mathcal{B}_{i}(P) \cup \mathcal{B}_{i}(Q)
$$

Remark 6.8. Proposition 6.7 has a category-theoretic interpretation: It says that reduced persistent homology commutes with coproducts in the categories of based metric spaces and persistence modules, where morphisms of metric spaces are 1-Lipschitz maps sending basepoint to basepoint.

Remark 6.9. Proposition 6.7 has also been discovered independently by the authors of [2]. Their work also establishes the result for infinite based metric spaces and for Čech filtrations.

We leave the easy verification of the following to the reader:
Proposition 6.10. Suppose $G$ is a phylogenetic graph with $G=G^{1} \vee G^{2}$ for subgraphs $G^{1}, G^{2} \subseteq G, \mathcal{E}$ is a history indexed by $G$, and $\mathcal{E}^{1}$ and $\mathcal{E}^{2}$ are the respective restrictions of $\mathcal{E}$ to $G^{1}$ and $G^{2}$. Then

$$
\operatorname{met} \mathcal{E} \cong \operatorname{met} \mathcal{E}^{1} \vee \operatorname{met} \mathcal{E}^{2}
$$

Theorem 6.11. Suppose a galled tree $G$ is an iterated sum of source-sink loops $G^{1}, \ldots, G^{k} \subset$ $G$ and rooted trees $G^{k+1}, \ldots, G^{l} \subset G$, and that $\mathcal{E}$ is a history indexed by $G$. Let $\mathcal{E}^{j}$ denote the restriction of $\mathcal{E}$ to $G^{j}$.
(i) There is an objectwise homotopy equivalence from an iterated wedge sum of the filtrations $\mathcal{V}\left(\mathcal{E}^{j}\right)$ to $\mathcal{V}(\mathcal{E})$.
(ii) For $i \geq 1$,

$$
\mathcal{B}_{i}(\mathcal{E})=\bigcup_{j=1}^{k} \mathcal{B}_{i}\left(\mathcal{E}^{j}\right)
$$

Proof. (i) follows from Propositions 6.7 and 6.10. (ii) follows from (i) and Proposition 6.3.
6.3. Vietoris-Rips Filtrations of Almost Linear Metric Spaces. As mentioned in the introduction, we say a non-empty finite metric space $P$ is almost linear if there is a point $p \in P$ such that $P \backslash\{p\}$ is isometric to a finite subset of $\mathbb{R}$. We call any such point $p a$ distinguished point. See Fig. 8.

Proposition 6.12. If $\mathcal{E}$ is a history indexed by a source-sink loop, then met $\mathcal{E}$ is almost linear.

Proof. Let $p$ be the unique recombinant. $\operatorname{met} \mathcal{E} \backslash\left\{\mathcal{E}_{p}\right\}$ is isometric to a subset of $\mathbb{R}$.

## 000000

Figure 8. An almost linear metric space embedded in $\mathbb{R}^{2}$. The unique distinguished point is shown in solid black. Note that not all almost linear metric spaces can be embedded in $\mathbb{R}^{2}$.

In view of Theorem 6.11 and Proposition 6.12, to understand the topology of Vietoris-Rips filtrations of histories indexed by galled trees, it suffices to understand the topology of Vietoris-Rips filtrations of almost linear metric spaces. We now describe the latter:

Theorem 6.13 (Topology of the Vietoris-Rips Filtration of an Almost Linear Metric Space). Let $P$ be an almost linear metric space with distinguished point $p$.
(i) For each $r \in[0, \infty)$, the connected component $C_{r}$ of $\mathcal{V}(P)_{r}$ containing $p$ is either contractible or homotopy equivalent to a circle, and each other component of $\mathcal{V}(P)_{r}$ is contractible. In particular, $\mathcal{B}_{i}(P)=0$ for $i \geq 2$.
(ii) If $C_{r}$ and $C_{r^{\prime}}$ are both homotopy equivalent to circles and $r \leq r^{\prime}$, then the inclusion $C_{r} \hookrightarrow C_{r^{\prime}}$ is a homotopy equivalence. Thus, $\mathcal{B}_{1}(P)$ has at most one interval.
(iii) The unique interval of $\mathcal{B}_{1}(P)$, when it exists, has length at most $d(p, P \backslash\{p\})$ and is contained in the interval

$$
[d(p, P \backslash\{p\}) \text {, diameter }(P \backslash\{p\}) / 2))
$$

Remark 6.14. Together, Theorem 6.11 (i), Proposition 6.12, and Theorem 6.13 (i) tell us that for $\mathcal{E}$ a history indexed by a galled tree $G$ and $r \in[0, \infty)$, each component of $\mathcal{V}(\mathcal{E})_{r}$ is homotopy equivalent to a bouquet of circles.

Remark 6.15. In analogy with the definition of an almost linear metric space, we can define an almost tree-like metric space to be one obtained from a tree-like metric space by adding a single point. In an earlier version of this paper, we conjectured that Theorem 6.13 also holds for almost tree-like metric spaces, but Matthew Zaremsky showed us the following simple counterexample: Let $d$ be the metric on $\{a, b, c, z\}$ given by

$$
d(a, z)=d(b, z)=d(c, z)=1, \quad d(a, b)=d(a, c)=d(b, c)=2 .
$$

This metric is tree-like; we can take the tree to be the star centered at $z$. We extend $d$ to a metric on $\{a, b, c, z, p\}$ by taking $d(p, z)=2$ and $d(p, a)=d(p, b)=d(p, c)=1$. The resulting metric space $P$ is almost tree-like, but $\mathcal{B}_{1}(P)$ consists of two intervals. Thus, Theorem 6.13 (ii) does not extend to almost tree-like metric spaces.

We build up to the proof of Theorem 6.13 with several definitions and lemmas. In what follows, let $P$ be an almost linear metric space with distinguished point $p$, and let $r \in[0, \infty)$ be such that $\mathcal{V}(P)_{r}$ is connected. By choosing an isometric embedding $P \backslash\{p\} \hookrightarrow \mathbb{R}$, we may regard $P \backslash\{p\}$ as a subset of $\mathbb{R}$.
Definition 6.16. Let $P_{\text {left }} \subset P \backslash\{p\}$ denote the set of points $y$ such that

1. $[p, y] \notin \mathcal{V}(P)_{r}$,
2. there is no $w \in P \backslash\{p\}$ satisfying each of the following conditions:

- $w<y$,
- $w$ and $y$ lie in the same connected component of $\mathcal{V}(P \backslash\{p\})_{r}$,
- $[p, w] \in \mathcal{V}(P)_{r}$.

See Fig. 9 for an illustration of $P_{\text {left }}$.


Figure 9. Illustration of the Vietoris-Rips complex $\mathcal{V}(P)_{r}$ for an almost linear metric space $P$, and some choice of scale parameter $r$. Here, the metric on $P$ is not assumed to be the one given by the shown embedding of the points in the plane. The distinguished point is solid black, points of $P_{\text {left }}$ are white, and the remaining points are gray.

Lemma 6.17. $\mathcal{V}(P)_{r}$ deformation retracts onto $\mathcal{V}\left(P \backslash P_{\text {left }}\right)_{r}$.
Proof. We give a simple discrete Morse theory argument. Define a DGVF $W$ on $\mathcal{V}(P)_{r}$ as follows: For $j \geq 2$ and

$$
\sigma:=\left[a_{1}<a_{2}<\cdots<a_{j}\right]
$$

a simplex in $\mathcal{V}(P)_{r}$ such that $a_{1} \in P_{\text {left }}$ and $a_{2}$ is the point in $P$ immediately to the right of $a_{1}, W$ matches $\sigma$ to its face $\left[a_{1}, a_{3}, \ldots, a_{j}\right]$. To see that $W$ is acyclic, note that for any $W$-path

$$
\sigma_{0}, \tau_{0}, \ldots, \sigma_{m}, \tau_{m}, \sigma_{m+1}
$$

the $\tau_{j}$ are strictly increasing with respect to the lexicographical order induced by the vertex ordering. If $m \geq 0$ and $\sigma_{0}=\sigma_{m+1}$, then

$$
\sigma_{0}, \tau_{0}, \ldots, \sigma_{m}, \tau_{m}, \sigma_{0}, \tau_{0}, \sigma_{1}
$$

is a $W$-path with $\tau_{m}<\tau_{0}$, so there cannot exist a non-trivial closed $W$-path. Therefore $W$ is acyclic.

Furthermore, $W$ matches every simplex containing a point in $P_{\text {left }}$, so the critical simplices of $W$ form the subcomplex $\mathcal{V}\left(P \backslash P_{\text {left }}\right)_{r}$. Hence, $\mathcal{V}(P)_{r}$ deformation retracts onto $\mathcal{V}\left(P \backslash P_{\text {left }}\right)_{r}$ by Proposition 5.9 (ii).

Let us now assume that $P_{\text {left }}=\emptyset$. We next define a discrete gradient vector field $Y$ on $\mathcal{V}:=\mathcal{V}(P)_{r}$. We do so in two steps, first giving a simple definition of a DGVF $X$ on $\mathcal{V}:=\mathcal{V}(P)_{r}$, and then extending this by matching more simplices. To start, we order the vertices in $P$ by taking $\{p\}$ to be the minimum, and ordering $P-\{p\}$ from left to right, via the chosen embedding of $P-\{p\}$ into $\mathbb{R}$. Henceforth, it will be our convention that the vertices of a simplex in $\mathcal{V}$ are always written in increasing order.

The definition of $X$ is an instance of a general construction due to Matt Kahle [35, Section 5], which in fact gives a DGVF on any simplicial complex with ordered vertex set:

Definition 6.18 (The Discrete Gradient Vector Field $X$ ). If a simplex $\sigma=\left[a_{1}, a_{2}, \ldots, a_{j}\right]$ of on $\mathcal{V}$ has a coface $a_{0} \cup \sigma:=\left[a_{0}, a_{1}, a_{2}, \ldots, a_{j}\right]$ with $a_{0}<a_{1}$, then $X$ matches $\sigma$ to $a_{0} \cup \sigma$ with $a_{0}$ as small as possible. $X$ matches no other simplices. It is easy to check that this in fact gives a well-defined DGVF.

See Fig. 10 for an illustration of the DGVF $X$.


Figure 10. Illustration of the discrete gradient vector field $X$ on a Vietoris-Rips complex of an almost-linear metric space, with the bottom vertex ordered first, and the remaining vertices ordered left-to-right. Matched simplices are gray, and matched pairs are denoted with an arrow pointing away from the simplex of lower dimension. Critical simplices are black. Thus, $X$ has a one critical 0 -simplex, two critical 1 -simplices, and one critical 2 -simplex.

Clearly, $[p]$ is critical in $X$, and since we assume that $P_{\text {left }}=\emptyset$, no other vertex is critical. The following describes the remaining critical simplices in $X$ :

Lemma 6.19. For $j \geq 2$, a simplex $\left[a_{1}, \ldots, a_{j}\right]$ is critical in $X$ if and only if the following three conditions are satisfied:

1. $a_{1} \neq p$,
2. $\left[q, a_{1}, \ldots, a_{j}\right] \notin \mathcal{V}$ for any $q<a_{1}$,
3. $\left[p, a_{2}, a_{3}, \ldots a_{j}\right] \in \mathcal{V}$.

Proof. If all three conditions hold, then by condition $2, \sigma:=\left[a_{1}, \ldots, a_{j}\right]$ cannot be the simplex of lower dimension in a pair matched by $X$, and by condition $3, X$ matches $\left[a_{2}, a_{3}, \ldots a_{j}\right]$ to [ $p, a_{2}, a_{3}, \ldots a_{j}$ ], so by condition $1, \sigma$ cannot be the simplex of higher dimension in a pair matched by $X$. Thus $\sigma$ is critical in $X$.

Conversely, if $\sigma$ is critical in $X$, then condition 1 holds, for else $\sigma$ would match to $\left[a_{2}, a_{3}, \ldots a_{j}\right]$. Condition 2 holds, for else $\sigma$ would match to a simplex of higher dimension. Finally, condition 3 holds, for else $\left[a_{2}, a_{3}, \ldots, a_{j}\right]$ would match to some simplex $\left[q, a_{2}, \ldots, a_{j}\right] \in$ $\mathcal{V}$ with $p<q<a_{1}$, implying that $\left[q, a_{1}, \ldots, a_{j}\right] \in \mathcal{V}$, and hence contradicting the criticality of $\sigma$.

Remark 6.20. Note that Lemma 6.19 implies in particular that if $\left[a_{1}, \ldots, a_{j}\right]$ is critical, then $a_{1}$ is not incident to $p$, since otherwise, in view of condition 3 , condition 2 would be violated.

Lemma 6.19 suggests a way to extend $X$ to a DGVF $Y$ with the desired properties:
Definition 6.21 (The Discrete Gradient Vector Field $Y$ ). For $\left[a_{1}, a_{2}, a_{3}, \ldots, a_{j}\right]$ a critical simplex for $X$ with $j \geq 3$, suppose there exists no vertex $b$ such that $[p, b] \in \mathcal{V}$ and $a_{1}<b<a_{2}$. It follows easily from Lemma 6.19 that $\left[a_{1}, a_{3}, \ldots, a_{j}\right]$ is also critical in $X$. We
match $\left[a_{1}, a_{2}, a_{3}, \ldots, a_{j}\right]$ to $\left[a_{1}, a_{3}, \ldots, a_{j}\right]$ in $Y$. We take all matched pairs in $Y \backslash X$ to be of this form.

Lemma 6.22. The matching $Y$ is acyclic, hence a $D G V F$.
Proof. We claim that in any $Y$-path

$$
\sigma_{0}, \tau_{0}, \ldots, \sigma_{m}, \tau_{m}, \sigma_{m+1}
$$

no two distinct $\tau_{j}$ are equal. From this, it follows that there does not exist a non-trivial closed $Y$-path, so $Y$ is indeed acyclic. To verify the claim, we make three simple observations: Letting $\tau_{j}^{1}$ denote the minimum vertex in $\tau_{j}$, we have that for any $j \in\{0, \ldots, m-1\}$,

1. If $\tau_{j+1}$ is matched by $X$, then $\tau_{j}^{1}>\tau_{j+1}^{1}$.
2. If $\tau_{j+1}$ is matched by $Y \backslash X$, then $\tau_{j}$ is matched by $X$ and $\tau_{j+1}^{1}=\tau_{j}^{1}$.
3. $\tau_{j} \neq \tau_{j+1}$.

By observations 1 and 2, we have that $\tau_{j}^{1}>\tau_{k}^{1}$ for all $k \in\{j+2, j+3, \ldots, m\}$. The claim follows from this and observation 3 .

Fig. 11 illustrates the extension of the DGVF $X$ of example Fig. 10 to the DGVF $Y$.


Figure 11. The extension of the DGVF $X$ of example Fig. 10 to the DGVF $Y . Y$ contains one pair of matched simplices not in $X$ : The curved 1-simplex in the top of the figure now is matched with its coface. Thus, $Y$ has two critical simplices: A critical 0 -simplex and a critical 1 -simplex.

Lemma 6.23. The critical simplices of $Y$ are $[p]$ and the 1 -simplices $\left[a_{1}, a_{2}\right]$ such that

1. $\left[a_{1}, a_{2}\right]$ satisfies the conditions of Lemma 6.19 and
2. $[p, b] \notin \mathcal{V}$ for all $a_{1}<b<a_{2}$.

In particular, $Y$ has a single critical 0-simplex, and no critical simplices of dimension greater than one.

Proof. Since $Y$ is an extension of $X$, any critical simplex of $Y$ is a critical simplex of $X$. It is easy to see that $Y$ matches every critical simplex of $X$, except $[p]$ and those 1 -simplices satisfying condition 2. A 1-simplex is critical in $X$ if and only if it satisfies the conditions of Lemma 6.19, so the result follows.

Lemma 6.24. For $P$ an almost linear metric space and $r \in[0, \infty)$, each component of $\mathcal{V}=\mathcal{V}(P)_{r}$ is contractible or deformation retracts onto a wedge sum of finitely many circles.

Proof. Any component of $\mathcal{V}$ not containing $p$ is tree-like, and so is contractible by Proposition 6.3. Thus, we may assume loss of generality that $\mathcal{V}$ is connected. Moreover, by

Lemma 6.17, we may assume without loss of generality that $P_{\text {left }}=\emptyset$. The DGVF $Y$ on $\mathcal{V}$ is defined under these assumptions. The result now follows from Lemma 6.23 and Proposition 5.9 (i).

Proof of Theorem 6.13 (i). As in the proof of Lemma 6.24, we may assume without loss of generality that $\mathcal{V}:=\mathcal{V}(P)_{r}$ is connected, and that $P_{\text {left }}=\emptyset$. The fundamental group of a wedge sum of circles is free [30, Example 1.21], so by Lemma 6.24, $\pi_{1}(\mathcal{V}, p)$ is free. To establish Theorem $6.13(\mathrm{i})$, it suffices to show that $\pi_{1}(\mathcal{V}, p)$ is trivial or cyclic.

To show this, we first note that the DGVF $Y$ provides us with a basis for $\pi_{1}(\mathcal{V}, p)$, as follows: Let $\Gamma$ denote the set of critical 1-simplies of $Y$, as described by Lemma 6.23. For $\sigma=[b, c] \in \Gamma$ with $b<c$, let $a \in P-\{p\}$ denote the maximum vertex such that $a<b$ and $[p, a] \in \mathcal{V}$. Such $a$ always exists by our assumption that $P_{\text {left }}=\emptyset$. Let us regard $S^{1}$ as a based topological space, with the basepoint denoted as 1 , and let $\gamma_{\sigma}: S^{1} \rightarrow[p, a] \cup[a, c] \cup[p, c]$ be a homeomorphism sending 1 to $p$.

We now observe that $G:=\left\{\gamma_{\sigma} \mid \sigma \in \Gamma\right\}$ is a basis for $\pi_{1}(\mathcal{V}, p)$. For $\sigma \in \Gamma$, let $S_{\sigma}^{1}$ denote a copy of $S^{1}$. The proof of Proposition 5.9 (i) presented in [37] gives a (not necessarily unique) homotopy equivalence $h: \mathcal{V} \rightarrow \vee_{\sigma \in \Gamma} S_{\sigma}^{1}$ mapping the interior of $\sigma$ homeomorphically to $S_{\sigma}^{1} \backslash\{1\}$, so that $h \circ \gamma_{\sigma}$ is homotopic either to the inclusion $\iota_{\sigma}: S_{\sigma}^{1} \hookrightarrow \vee_{\sigma \in \Gamma} S_{\sigma}^{1}$, or to its inverse in $\pi_{1}\left(\vee_{\sigma \in \Gamma} S_{\sigma}^{1}, 1\right)$. Since $h$ is a homotopy equivalence and $\left\{\iota_{\sigma} \mid \sigma \in \Gamma\right\}$ is a basis for $\pi_{1}\left(\vee_{\sigma \in \Gamma} S_{\sigma}^{1}, 1\right)$, we see that $G$ is a basis for $\pi_{1}(\mathcal{V}, p)$, as desired.

To finish the proof of Theorem 6.13 (i), it remains to show that $|G| \leq 1$. To do so, we apply the triangle inequality. Our argument is illustrated in Fig. 12. For $[b, c]=\sigma \in \Gamma$ with $b<c$, let $a<b$ be as above, and for $\left[b^{\prime}, c^{\prime}\right]=\sigma^{\prime} \in \Gamma$ with $b^{\prime}<c^{\prime}$, define $a^{\prime}<b^{\prime}$ in the same way. To arrive at a contradiction, suppose $\sigma \neq \sigma^{\prime}$. Then either $c \leq a^{\prime}$ or $c^{\prime} \leq a$. Switching the labels of $\sigma$ and $\sigma^{\prime}$ if necessary, we may assume without loss of generality that $c \leq a^{\prime}$. We have $[p, a],\left[p, c^{\prime}\right] \in \mathcal{V}$, so $d(a, p) \leq 2 r$ and $d\left(p, c^{\prime}\right) \leq 2 r$. By the triangle inequality, $d\left(a, c^{\prime}\right) \leq 4 r$. Thus, since $P \backslash\{p\}$ is isometric to a subset of $\mathbb{R}$, we have

$$
d(a, c)+d\left(a^{\prime}, c^{\prime}\right) \leq d\left(a, c^{\prime}\right) \leq 4 r
$$

Therefore either $d(a, c) \leq 2 r$ or $d\left(a^{\prime}, c^{\prime}\right) \leq 2 r$, so either $[a, c] \in \mathcal{V}$ or $\left[a^{\prime}, c^{\prime}\right] \in \mathcal{V}$. But then either $\gamma_{\sigma}$ or $\gamma_{\sigma^{\prime}}$ is nullhomotopic in $\mathcal{V}$, contradicting that $G$ is a basis for $\pi_{1}(\mathcal{V}, p)$.


Figure 12. Illustration of the argument by contradiction that $|G| \leq 1$ in the proof of Theorem 6.13 (i). Critical simplices are black and matched simplices are gray. By the triangle inequality, $d\left(a, c^{\prime}\right) \leq 4 r$, so since $\left\{a<b<c \leq a^{\prime}<b^{\prime}<c^{\prime}\right\}$ is isometric to a subset of $\mathbb{R}$, either $[a, c] \in \mathcal{V}$ or $\left[a^{\prime}, c^{\prime}\right] \in \mathcal{V}$.

Proof of Theorem 6.13 (ii). As in the statement of the theorem, let $C_{r}$ denote the component of $\mathcal{V}(P)_{r}$ containing $\{p\}$. We need to show that for $r \leq r^{\prime} \in[0, \infty)$, if $C_{r} \simeq S^{1} \simeq C_{r^{\prime}}$, then the inclusion $C_{r} \hookrightarrow C_{r^{\prime}}$ is a homotopy equivalence. Let $\gamma_{\sigma}: S^{1} \rightarrow C_{r}$ and $\gamma_{\sigma^{\prime}}: S^{1} \rightarrow C_{r^{\prime}}$ be the generators for $\pi_{1}\left(C_{r}, p\right)$ and $\pi_{1}\left(C_{r^{\prime}}, p\right)$ specified in the proof of Theorem $6.13(\mathrm{i})$ above.

Given the way $\gamma_{\sigma}$ and $\gamma_{\sigma^{\prime}}$ are defined, exactly one of the following must be true:

1. $c \leq a^{\prime}$,
2. $c^{\prime} \leq a$,
3. $a \leq a^{\prime}<c^{\prime} \leq c$.

We show that we cannot have $c \leq a^{\prime}$ using essentially the same triangle inequality argument we used in the proof of Theorem 6.13 (i): Suppose otherwise. Then $d(a, p)<2 r$ and $d\left(c^{\prime}, p\right)<2 r^{\prime}$. By the triangle inequality, $d\left(a, c^{\prime}\right) \leq 2\left(r+r^{\prime}\right)$, so we have

$$
d(a, c)+d\left(a^{\prime}, c^{\prime}\right) \leq d\left(a, c^{\prime}\right) \leq 2\left(r+r^{\prime}\right)
$$

Therefore either $d(a, c) \leq 2 r$ or $d\left(a^{\prime}, c^{\prime}\right) \leq 2 r^{\prime}$, leading to a contradiction as above.
The same argument shows that we cannot have $c^{\prime} \leq a$. Therefore, we must have $a \leq a^{\prime}<$ $c^{\prime} \leq c$.

We will show that if $a \neq a^{\prime}$, then $\left[a, a^{\prime}\right] \in C_{r^{\prime}}$ : We have $d(a, p) \leq 2 r$ and $d\left(c^{\prime}, p\right) \leq 2 r^{\prime}$, so by the triangle inequality, $d\left(a, c^{\prime}\right) \leq 2\left(r+r^{\prime}\right)$. Therefore either $d\left(a, a^{\prime}\right) \leq 2 r^{\prime}$ or $d\left(a^{\prime}, c^{\prime}\right) \leq 2 r$. But since $\gamma_{\sigma^{\prime}}$ is not nullhomotopic by assumption, we must have $d\left(a^{\prime}, c^{\prime}\right)>2 r^{\prime} \geq 2 r$, so $d\left(a, a^{\prime}\right) \leq 2 r^{\prime}$. Thus $\left[a, a^{\prime}\right] \in C_{r^{\prime}}$, as desired. It follows that $\left[p, a, a^{\prime}\right] \in C_{r^{\prime}}$.

The symmetric argument shows that if $c^{\prime} \neq c$, then $\left[p, c^{\prime}, c\right] \in C_{r^{\prime}}$. Letting

$$
j: C_{r} \hookrightarrow C_{r^{\prime}}
$$

denote the inclusion, we thus have that $j \circ \gamma_{\sigma} \sim \gamma_{\sigma^{\prime}}$. Since $\gamma_{\sigma}$ and $\gamma_{\sigma^{\prime}}$ are both homotopy equivalences, $j$ must be a homotopy equivalence as well.

Proof of Theorem 6.13 (iii). Given the form of the set $G$ of generators for $\pi_{1}\left(\mathcal{V}(P)_{r}, p\right)$ given in the proof of Theorem 6.13 (i), it is clear that if

$$
r \notin[d(p, P \backslash\{p\}), \text { diameter }(P \backslash\{p\}) / 2),
$$

then $\pi_{1}\left(\mathcal{V}(P)_{r}, p\right)$ is trivial. By $(i)$ then, each component of $\mathcal{V}(P)_{r}$ is contractible. Hence, the unique interval of $\mathcal{B}_{1}(P)$, if it exists, is contained in

$$
[d(p, P \backslash\{p\}), \operatorname{diameter}(P \backslash\{p\}) / 2) .
$$

To finish the proof of (iii), we need to show that the unique bar of $\mathcal{B}_{1}(P)$ is of length at most $d(p, P \backslash\{p\})$. This follows from the stability of persistent homology. To see this, note that since $P \backslash\{p\}$ is isometric to a subset of $\mathbb{R}$, it is tree-like, so Proposition 6.3 gives that $\mathcal{B}_{1}(P \backslash\{p\})=\emptyset$. Therefore, by Theorem 5.6,

$$
2 d_{B}\left(\mathcal{B}_{1}(P), \emptyset\right)=2 d_{B}\left(\mathcal{B}_{1}(P), \mathcal{B}_{1}(P \backslash\{p\})\right) \leq d_{H}(P, P \backslash\{p\})=d(p, P \backslash\{p\}),
$$

where the last equality follows from the definition of $d_{H}$. The bottleneck distance of any barcode $\mathcal{B}$ to the empty barcode is half the length of the longest interval of $\mathcal{B}$, so the result follows.
6.4. Inference about Recombination from Barcodes. As an immediate corollary of the results of Sections 6.2 and 6.3, we now obtain our main result relating barcodes to recombination in the galled tree setting.

Recall from Section 3 that $\mathcal{T}(\mathcal{E})$ denotes the topological novelty profile of a history $\mathcal{E}$, and that the temporal novelty of a recombinant $r$ (with respect to some choice of time function) is denoted as $\mathcal{N}(r)$. Recall also from Proposition 4.5 that when $\mathcal{E}$ is indexed by a galled tree, $\mathcal{T}(\mathcal{E})$ is equal to the temporal novelty profile of $\mathcal{E}$, with respect to any time function.

For $G$ a phylogenetic graph, let $\mathcal{R}^{G}$ denote the set of recombinants of $G$. As in the beginning of Section 6 , for $\mathcal{B}$ a barcode, let lengths $(\mathcal{B})$ denote the list of lengths of intervals of $\mathcal{B}$, sorted in descending order.

Theorem 6.25. Let $\mathcal{E}$ be a history indexed by a galled tree $G$.
(i) Theorem 6.11 (ii) and Theorem 6.13 (ii) yield a canonical injection

$$
\phi: \mathcal{B}_{1}(\mathcal{E}) \hookrightarrow \mathcal{R}^{G},
$$

such that length $(I) \leq \mathcal{N}(\phi(I))$ for all $I \in \mathcal{B}_{1}(\mathcal{E})$. In particular,

$$
\text { lengths }\left(\mathcal{B}_{1}(\mathcal{E})\right) \leq \mathcal{T}(\mathcal{E})
$$

(ii) $\mathcal{B}_{i}(\mathcal{E})=\emptyset$ for $i \geq 2$.

Proof. For $G$ a galled tree, each $r \in \mathcal{R}^{G}$ corresponds to an entry of $\mathcal{T}(\mathcal{E})$; in fact, this entry is easily seen to be $d\left(\mathcal{E}^{L \backslash r}, \mathcal{E}_{r}\right)$, where $L$ denotes the source-sink loop corresponding to $R$, and $\mathcal{E}^{L \backslash r}$ denotes the restriction of $\mathcal{E}$ to vertices of $L$ other than $r$. (i) now follows from Theorem 6.13 (iii).
(ii) is immediate from Theorem 6.11 (ii), Proposition 6.12, and Theorem 6.13 (i).

Example 6.26. Given the analogy between Theorem 6.4 (for trees) and Theorem 6.25 (for galled trees), and the fact that Theorem 6.4 holds for arbitrary subsamples of a history, it is natural to ask whether Theorem 6.25 also holds for arbitrary subsamples. The example shown in Fig. 13 demonstrates that Theorem 6.25 (i) does not hold for arbitrary subsamples; the example, discovered by computer, is a subset $S$ of a history $\mathcal{E}$ indexed by a galled tree with a single recombinant, for which $\mathcal{B}_{1}(S)=\{[5,6),[5,6)\}$ and $\mathcal{B}_{1}(\mathcal{E})=\emptyset$. We conjecture that Theorem 6.25 (ii) also does not hold for arbitrary subsets.

Nevertheless, it may be the case that for reasonable random models of histories indexed by galled trees, violations of Theorem 6.25 are relatively rare. We provide some preliminary numerical evidence for this in Appendix B, focusing on how often the number of intervals in $B_{1}(S)$ of a sample $S$ exceeds the number of recombinants in the underlying history.


Figure 13. A subset $S$ of a history $\mathcal{E}$ indexed by a galled tree with one recombinant, for which $\left|\mathcal{B}_{1}(S)\right|=2$ and $B_{1}(\mathcal{E})=\emptyset$. Nodes corresponding to elements of $S$ are shown in black; the remaining nodes are shown in white.

## 7. Relaxing the Complete Sampling and Galled Tree Assumptions

Theorem 6.25, the main result of the previous section, holds under the assumption that our evolutionary history is indexed by a galled tree, and that all organisms in the history have been sampled. In this section, we apply the stability of persistent homology to extend the theorem to the case of an arbitrary (noisy) subsample of a history indexed by an arbitrary phylogenetic graph.
7.1. Relaxing the Complete Sampling Assumption. First, we extend Theorem 6.25 to the case of a noisy subsample. Given a list of non-negative numbers $L$, let $\operatorname{Trim}(L, \delta)$ be the list obtained by removing each of the numbers less than or equal to $\delta$ and subtracting $\delta$ from each of the remaining numbers.

Corollary 7.1. Let $\mathcal{E}$ be a history indexed by a galled tree and let $S$ be a finite metric space with $d_{G H}(\mathcal{E}, S)=\delta$. Then
(i) $\operatorname{Trim}\left(\right.$ lengths $\left.\left(\mathcal{B}_{1}(S)\right), 2 \delta\right) \leq \mathcal{T}(\mathcal{E})$.
(ii) For $i \geq 2$, each interval of $\mathcal{B}_{i}(S)$ has length at most $2 \delta$.

Proof. This follows immediately from Theorems 5.5 and 6.25.
7.2. Relaxing the Galled Tree Assumption. As an application of Corollary 7.1 (i), we next also relax the assumption that $\mathcal{E}$ is indexed by a galled tree, yielding a further extension of Theorem 6.25 which applies to any phylogenetic graph.

For $G$ any phylogenetic graph and $\mathcal{E}$ a history indexed by $G$, let

$$
\bigcup \mathcal{E}:=\bigcup_{v \in V} \mathcal{E}_{v} .
$$

Thus, $\bigcup \mathcal{E}$ is set of all mutations appearing in the history $\mathcal{E}$.

For $M \subseteq \bigcup \mathcal{E}$ any subset and $v \in V$, let $\mathcal{E}_{v}^{M}=\mathcal{E}_{v} \backslash M$. Let $G^{M}$ denote a subgraph of $G$ obtained by removing edges as follows: Suppose $w$ is a recombinant of $G$ with parents $u, v$. If $\mathcal{E}_{w}^{M}=\mathcal{E}_{u}^{M} \neq \mathcal{E}_{v}^{M}$, we remove the edge $(v, w)$ from $G$. If $\mathcal{E}_{w}^{M}=\mathcal{E}_{u}^{M}=\mathcal{E}_{v}^{M}$ we remove exactly one of the edges $(u, w)$ and $(v, w)$, choosing arbitrarily. It is easy to check that the sets $\mathcal{E}_{v}^{M}$ then give a well-defined evolutionary history $\mathcal{E}^{M}$ indexed by $G^{M}$.

Definition 7.2. We let
$\operatorname{Gall}(\mathcal{E}):=\min \left\{|M| \mid M \subseteq \bigcup \mathcal{E}\right.$ such that $G^{M}$ can be be chosen to be a galled tree $\}$.
Informally, $\operatorname{Gall}(\mathcal{E})$ is the number of mutations in $\mathcal{E}$ which must be ignored to obtain a history indexed by a galled tree by pruning edges in $G$.

The following is our most general result relating barcodes and recombination:
Corollary 7.3. Let $\mathcal{E}$ be a history indexed by an arbitrary phylogenetic graph $G$, and let $S$ be a finite metric space with $d_{G H}(\mathcal{E}, S)=\delta$. Then
(i) $\operatorname{Trim}\left(\operatorname{lengths}\left(\mathcal{B}_{1}(S)\right), 3 \operatorname{Gall}(\mathcal{E})+2 \delta\right) \leq \mathcal{T}(\mathcal{E})$.
(ii) For $i \geq 2$, each interval of $\mathcal{B}_{i}(S)$ has length at most $2(\operatorname{Gall}(\mathcal{E})+\delta)$.

Proof. Choose $M \subseteq \bigcup \mathcal{E}$ and a galled tree $G^{M}$ as above, such that $|M|=\operatorname{Gall}(\mathcal{E})$. Note that $d_{G H}\left(\mathcal{E}^{M}, \mathcal{E}\right) \leq|M|=\operatorname{Gall}(\mathcal{E})$, so by the triangle inequality, $d_{G H}\left(\mathcal{E}^{M}, S\right) \leq \operatorname{Gall}(\mathcal{E})+\delta$. (ii) then follows from Corollary 7.1 (ii). By Corollary 7.1 (i),

$$
\begin{equation*}
\operatorname{Trim}\left(\operatorname{lengths}\left(\mathcal{B}_{1}(S)\right), 2(\operatorname{Gall}(\mathcal{E})+\delta)\right) \leq \mathcal{T}\left(\mathcal{E}^{M}\right) \tag{1}
\end{equation*}
$$

Letting $\overline{\mathcal{T}}\left(\mathcal{E}^{M}\right)$ be the vector of length $|\mathcal{T}(\mathcal{E})|$ obtained by adding some 0 's to the end of $\mathcal{T}\left(\mathcal{E}^{M}\right)$, we have by Proposition 5.7 that $d_{\infty}\left(\mathcal{T}\left(\mathcal{E}^{M}\right), \mathcal{T}(\mathcal{E})\right) \leq|M|$. Together with (1), this implies that

$$
\operatorname{Trim}\left(\operatorname{lengths}\left(\mathcal{B}_{1}(S)\right), 3 \operatorname{Gall}(\mathcal{E})+2 \delta\right) \leq \mathcal{T}(\mathcal{E})
$$

which gives (i).
Remark 7.4. Clearly, for Corollary 7.3 to yield a strong bound, $\operatorname{Gall}(\mathcal{E})$ must be small. One might expect $\operatorname{Gall}(\mathcal{E})$ to be small but non-zero when recombination events typically affect short genome tracts (e.g., when they are gene conversion events).

## 8. Random Histories Indexed by Galled Trees

The results we have presented so far have been deterministic. In Section 8.1 below, we observe that in a wide class of probabilistic models of genetic sequence evolution on galled trees, the intervals of the first persistence barcode are independent random variables. Thus, to understand the statistical properties of these barcodes, it suffices to understand the special case that the galled tree is a source-sink loop. In Section 8.2, we study this special case numerically, for one choice of probabilistic model.
8.1. Independence of Intervals in the First Barcode. In this section, we assume the reader is familiar with basic elements of the measure-theoretic formulation of probability theory [21] and with the definition of conditional independence given a random variable [36, Chapter 5].

Notation. Suppose $X, Y$, and $Z$ are random variables on the same probability space. If $X$ is independent of $Y$, we write $X \Perp Y$. If $X$ is independent of $Y$ given $Z$, we write $X \Perp Y \mid Z$.

For $P$ a poset and $p \in P$, let

$$
\operatorname{nd}(p):=\{q \in P \mid p \not \leq q\} .
$$

(Here, nd stands for non-descendants.) In what follows, $P$ will often be the vertex set of a directed acylic graph, with the partial order induced by the graph.

If $\mathcal{E}$ is an evolutionary history indexed by $G, V$ is the vertex set of $G$, and $S \subseteq V$, we write $\mathcal{E}_{S}:=\left\{\mathcal{E}_{v} \mid v \in S\right\}$. Similarly, for $v, w \in V$, let $\mathcal{E}_{v \backslash w}:=\mathcal{E}_{v} \backslash \mathcal{E}_{w}$, and $\mathcal{E}_{v \cap w}:=\mathcal{E}_{v} \cap \mathcal{E}_{w}$. We will also use these notation conventions in combination with one another, so that e.g., $\mathcal{E}_{r \backslash(p \cap q)}$ is understood to denote $\mathcal{E}_{r} \backslash\left(\mathcal{E}_{p} \cap \mathcal{E}_{q}\right)$.

Definition 8.1 (Random History). For $G$ a fixed phylogenetic graph with vertices $V$, a random (evolutionary) history $\mathcal{E}$ indexed by $G$ consists of the following data:

- A probability space $\Omega$.
- A countable set $X_{v}$ for each $v \in V$, such that each element of $X_{v}$ is itself a set. We equip $X_{v}$ with the discrete $\sigma$-algebra.
- For each $v \in V$, a random variable $\mathcal{E}_{v}: \Omega \rightarrow X_{v}$ such that for each $\omega \in \Omega,\left\{\mathcal{E}_{v}(\omega)\right\}_{v \in V}$ is an evolutionary history.

Definition 8.2 (Locally Markov History). Suppose that $\mathcal{E}$ is a random history indexed by a phylogenetic graph $G$ with vertices $V . \mathcal{E}$ is said to be locally Markov if for each $v \in V$,

$$
\begin{equation*}
\mathcal{E}_{v} \Perp \mathcal{E}_{\mathrm{nd}(v)} \mid\left\{\mathcal{E}_{p} \mid p \text { a parent of } v\right\} . \tag{2}
\end{equation*}
$$

A locally Markov history is a special case of a Bayesian network, a widely used probabilistic model [38].

The assumption that a random history is locally Markov is quite natural; informally, this says that the genome of each organism depends only on the genomes of its parents. However, the next example shows that for $\mathcal{E}$ a locally Markov history, it is not necessarily the case that the intervals in $\mathcal{B}_{1}(\mathcal{E})$ are independent.

Example 8.3. In the locally Markov history $\mathcal{E}$ of Fig. 14, the mutations from the top sourcesink loop are passed down to the bottom source-sink loop, where they serve as "instructions" for how clonal mutations occur in the bottom loop. Thus, the intervals in $\mathcal{B}_{1}(\mathcal{E})$ associated to the two recombinants are not independent.

For each vertex $v \neq w, \mathcal{E}_{v}$ is completely determined by its parents. The top recombinant corresponds to an interval $\left[\frac{1}{2}, 1\right.$ ) in the first barcode with probability $\frac{1}{2}$, and to an empty interval with probability $\frac{1}{2}$. If the top recombinant corresponds to $\left[\frac{1}{2}, 1\right)$, then the bottom recombinant corresponds to $[1,2)$; otherwise, the bottom recombinant corresponds to $\left[\frac{1}{2}, 1\right)$.

Motivated by the above, we introduce the following subclass of locally Markov histories:
Definition 8.4 (Phylogenetically Markov History). Suppose that $\mathcal{E}$ is a random history indexed by a fixed phylogenetic graph $G$ with vertices $V$. We say $\mathcal{E}$ is phylogenetically Markov if

1. $\mathcal{E}$ is locally Markov.
2. $\mathcal{E}_{v \backslash p} \Perp \mathcal{E}_{p}$ for all clones $v \in V$ with parent $p$.


Figure 14. A locally Markov history $\mathcal{E}$ for which the intervals in the $1^{\text {st }}$ persistence barcode corresponding to the two recombinants are not independent.
3. For all recombinants $r$ with parents $p$ and $q$,

$$
\mathcal{E}_{r \backslash(p \cap q)} \Perp \mathcal{E}_{\mathrm{nd}(r)} \mid \mathcal{E}_{\{p \backslash(p \cap q), q \backslash(p \cap q)\}} .
$$

To parse condition 3, recall that a recombinant $r$ 's genome necessarily inherits what is common to both parents $p$ and $q$; condition 3 states that the rest of $r$ 's genome is independent of the genomes of all non-descendants of $r$, given the rest of the genomes of each parent.
Remark 8.5. Definition 8.4 is slightly redundant, in the sense that condition 3 implies the local Markov property for each recombinant; this follows from Lemma 8.8 below. One might hope that one could obtain an equivalent definition by replacing condition 3 in Definition 8.4 with the simpler condition that $\mathcal{E}_{r \backslash(p \cap q)} \Perp \mathcal{E}_{p \cap q}$, but in fact this is strictly weaker. It can be shown that our independence result (Theorem 8.7 below) does not hold for this weaker condition.

Example 8.6. Assume that $G$ is endowed with a time function $t: V \rightarrow \mathbb{R}$, as defined in Section 3. We specify (up to choice of labels for mutations) a phylogenetically Markov history $\mathcal{E}$, the Poisson history indexed by $G$ :

- $\mathcal{E}_{r}=\emptyset$, for $r$ the root of $G$.
- If $w$ is a clone with ancestor $v,\left|\mathcal{E}_{w} \backslash \mathcal{E}_{v}\right|$ is Poisson distributed with parameter $t(w)-t(v)$.
- If $w$ is a recombinant with parents $u$ and $v$, then for each $m \in \mathcal{E}_{u} \backslash \mathcal{E}_{v}, P\left(m \in \mathcal{E}_{w}\right)=p_{w}$, and for each $m \in \mathcal{E}_{v} \backslash \mathcal{E}_{u}, P\left(m \in \mathcal{E}_{w}\right)=1-p_{w}$. Here, we may either take $p_{w}=1 / 2$ for all $w$, or take the $p_{w}$ to be i.i.d. random variables with the uniform distribution on $[0,1]$.
Let

$$
\mathcal{I}:=\{[a, b) \mid a<b \in\{0,1,2, \ldots\}\} \cup\{\emptyset\} .
$$

Thus, $\mathcal{I}$ is a collection of intervals with integer endpoints, together with the empty interval. For $\mathcal{E}$ a history indexed by a galled tree $G$ and $r$ a recombinant in $G$, let $\mathcal{I}^{\mathcal{E}}(r) \in \mathcal{I}$ denote the unique interval in $\mathcal{B}_{1}(\mathcal{E})$ corresponding to $r$, if such an interval exists (see Theorem 6.25), and let $\mathcal{I}^{\mathcal{E}}(r)=\emptyset$ otherwise. As in Section 6.4, we let $\mathcal{R}^{G}$ denote the set of recombinants of $G$.

Here is the main result of this section:
Theorem 8.7. For $\mathcal{E}$ a phylogenetically Markov history indexed by a galled tree $G$, the random variables $\left\{\mathcal{I}^{\mathcal{E}}(r)\right\}_{r \in \mathcal{R}^{G}}$ are independent.

The proof of the theorem will use several standard facts about conditional independence, which we record in the following lemma.

Lemma 8.8. Assume $h$ is a measurable function whose domain is the codomain of the random variable $X$.
(i) If $X \Perp Y \mid Z$, then $Y \Perp X \mid Z$.
(ii) If $X \Perp Y \mid Z$, then $h(X) \Perp Y \mid Z$.
(iii) If $X \Perp Y \mid Z$, then $X \Perp Y \mid(Z, h(X))$.
(iv) If $X \Perp Y \mid Z$ and $W \Perp Y \mid(X, Z)$, then $(W, X) \Perp Y \mid Z$.
(v) If $X \Perp Y \mid Z$, then $(Z, X) \Perp Y \mid Z$.

Note that by taking $Z$ to be the identity random variable, we also obtain unconditional versions of (i)-(iv) above.

Proof. Properties (i)-(iv) appear in many places; see e.g. [38, Chapter 3]. We prove (v). Taking $W=Z$ in (iv), it suffices to show that $Z \Perp Y \mid(X, Z)$. By (ii), for this it is enough to show that $(X, Z) \Perp Y \mid(X, Z)$. But it is easy to check that in general, $A \Perp B \mid A$.

Lemma 8.9. If $\left\{X_{a}\right\}_{a \in P}$ is a collection of random variables indexed by a finite poset $P$ and $X_{a} \Perp X_{\operatorname{nd}(a)}$ for each $a \in P$, then the $\left\{X_{a}\right\}_{a \in P}$ are independent.
Proof. Choose a total order compatible with the partial order on $P$, and relabel the random variables with respect to this order as $X_{1}, \ldots, X_{|P|}$. We show by induction that $X_{1}, \ldots, X_{m}$ are independent for each $m \in\{1, \ldots,|P|\}$. The base case is trivial. Now suppose $X_{1}, \ldots, X_{m-1}$ are independent. The elements of $P$ corresponding to the indices $1, \ldots, m-1$ are in $\operatorname{nd}(m)$, so $X_{m}$ is independent of $X_{1}, \ldots, X_{m-1}$. By this and the induction hypothesis, $X_{1}, \ldots, X_{m}$ are independent.

Proof of Theorem 8.7. Order the vertices of $G$ arbitrarily. For $r \in \mathcal{R}^{G}$, let $D^{r}$ denote the distance matrix obtained by restricting $\mathcal{E}$ to the source-sink loop $L^{r}$ of $G$ with $\operatorname{sink} r$. The images of independent random variables under measurable functions remain independent, so in view of the results of Section 6 , it suffices to show that the $\left\{D^{r}\right\}_{r \in \mathcal{R}^{G}}$ are independent. Now for each recombinant $r$, let $V^{r}$ denote the vertices of $L^{r}$, and let $q_{r}$ denote the unique source of $L^{r}$. Since $\mathcal{E}_{q_{r}} \subseteq \mathcal{E}_{v}$ for all $v \in V^{r}$, clearly $D^{r}$ is determined by $A^{r}:=\left\{\mathcal{E}_{v \backslash q_{r}}\right\}_{v \in V^{r} \backslash\left\{q_{r}\right\}}$. Therefore, it in fact suffices to show that the random sets $\left\{A^{r}\right\}_{r \in \mathcal{R}^{G}}$ are independent.

We define a partial order on $\mathcal{R}^{G}$ by writing $r \leq r^{\prime}$ if for some $v \in V^{r} \backslash\left\{q_{r}\right\}$, there is a directed path from $v$ to $q_{r}^{\prime}$ in $G$; it is easy to check that this is in fact a partial order. This partial order induces a partial order on $\left\{A^{r}\right\}_{r \in \mathcal{R}^{G}}$. We establish the independence of the
$\left\{A^{r}\right\}_{r \in \mathcal{R}^{G}}$ by applying Lemma 8.9, using this partial order. Let

$$
\operatorname{nd}\left(L^{r}\right):=\bigcap_{v \in V^{r} \backslash\left\{q_{r}\right\}} \operatorname{nd}(v)
$$

If $r \not \leq r^{\prime}$ then $V^{r^{\prime}} \subseteq \operatorname{nd}\left(L^{r}\right)$, so $\left\{A^{r^{\prime}} \mid A^{r} \not \leq A^{r^{\prime}}\right\}$ is completely determined by $\mathcal{E}_{\text {nd }\left(L^{r}\right)}$. Thus, it suffices to show that $A^{r} \Perp \mathcal{E}_{\mathrm{nd}\left(L^{r}\right)}$ for each $r \in \mathcal{R}^{G}$.

Let us fix $r \in \mathcal{R}^{G}$ and write $q=q^{r}$. Choose a total order on $V^{r} \backslash\{r, q\}$ compatible with the partial order on $V$, and write the elements in increasing order as $\left\{c_{1}, \ldots, c_{m}\right\}$. For $j \in\{1, \ldots, m\}$, let $B_{j}:=\left\{\mathcal{E}_{c_{i} \backslash q}\right\}_{1 \leq i \leq j} \subset A^{r}$. We show by induction on $j$ that $B_{j} \Perp \mathcal{E}_{\text {nd }\left(L^{r}\right)}$ for each $j$.

First, consider the base case $j=1$. In the remainder of the proof, the five statements of Lemma 8.8 will be denoted simply as (i)-(v). By the definition of a phylogenetically Markov history, we have $\mathcal{E}_{c_{1}} \Perp \mathcal{E}_{\text {nd }\left(c_{1}\right)} \mid \mathcal{E}_{q}$, so by (ii), we have $\mathcal{E}_{c_{1}} \Perp \mathcal{E}_{\text {nd }\left(L^{r}\right)} \mid \mathcal{E}_{q}$. By (v) then, $\mathcal{E}_{\left\{c_{1}, q\right\}} \Perp \mathcal{E}_{\text {nd }\left(L^{r}\right)} \mid \mathcal{E}_{q}$, so by (ii), $\mathcal{E}_{c_{1} \backslash q} \Perp \mathcal{E}_{\text {nd }\left(L^{r}\right)} \mid \mathcal{E}_{q}$. The definition of a phylogenetically Markov history also gives that $\mathcal{E}_{c_{1} \backslash q} \Perp \mathcal{E}_{q}$. Applying (iv) and (ii), we find that $\mathcal{E}_{c_{1} \backslash q} \Perp \mathcal{E}_{\text {nd }\left(L^{r}\right)}$. This shows that $B_{1} \Perp \mathcal{E}_{\text {nd }\left(L^{r}\right)}$.

The induction step is similar to the above. Let $p$ denote the parent of $c_{j}$. $\mathcal{E}_{c_{j}} \Perp \mathcal{E}_{\mathrm{nd}\left(c_{j}\right)} \mid \mathcal{E}_{p}$, so $\mathcal{E}_{c_{j} \backslash p} \Perp \mathcal{E}_{\text {nd }\left(c_{j}\right)} \mid \mathcal{E}_{p}$. Moreover, $\mathcal{E}_{c_{j} \backslash p} \Perp \mathcal{E}_{p}$, so $\mathcal{E}_{c_{j} \backslash p} \Perp \mathcal{E}_{\text {nd }\left(c_{j}\right)}$. Then by (ii) and (iii), $\mathcal{E}_{c_{j} \backslash p} \Perp \mathcal{E}_{\text {nd }\left(L^{r}\right)} \mid B_{j-1}$. By (iv) and the induction hypothesis, we thus have that $B_{j} \Perp \mathcal{E}_{\operatorname{nd}\left(L^{r}\right)}$, as desired.

Finally, we show that $A^{r} \Perp \mathcal{E}_{\operatorname{nd}\left(L^{r}\right)}$. Let $p_{1}$ and $p_{2}$ denote the parents of $r$. By the third condition in the definition of a phylogenetically Markov history, we have $\mathcal{E}_{r \backslash q} \Perp$ $\mathcal{E}_{\mathrm{nd}(r)} \mid \mathcal{E}_{\left\{p_{1} \backslash q, p_{2} \backslash q\right\}}$. By (iii), $\mathcal{E}_{r \backslash q} \Perp \mathcal{E}_{\mathrm{nd}(r)} \mid\left(\mathcal{E}_{\left\{p_{1} \backslash q, p_{2} \backslash q\right\}}, B_{m}\right) . \sigma\left(\mathcal{E}_{\left\{p_{1} \backslash q, p_{2} \backslash q\right\}}, B_{m}\right)=\sigma\left(B_{m}\right)$ since $\mathcal{E}_{\left\{p_{1} \backslash q, p_{2} \backslash q\right\}}=h\left(B_{m}\right)$ for some measurable function $h$, so $\mathcal{E}_{r \backslash q} \Perp \mathcal{E}_{\mathrm{nd}(r)} \mid B_{m}$ by the definition of conditional independence. By (ii), $\mathcal{E}_{r \backslash q} \Perp \mathcal{E}_{\mathrm{nd}\left(L^{r}\right)} \mid B_{m}$. We have also shown that $B_{m} \Perp \mathcal{E}_{\mathrm{nd}\left(L^{r}\right)}$, so by (iv), $A^{r} \Perp \mathcal{E}_{\mathrm{nd}\left(L^{r}\right)}$.
8.2. The Barcode of a Random History on a Source-Sink Loop: Numerical Results. Theorem 8.7 tells us that for a phylogenetically Markov history indexed by a galled tree, to understand the distribution of the $1^{\text {st }}$ barcode, it suffices to understand this for each source-sink loop in the galled tree. Working with a simple random model of a history $\mathcal{E}$ indexed by a source-sink loop, we now study the distribution of $\mathcal{B}_{1}(\mathcal{E})$ numerically. Recall that by Theorem 6.13 (ii), $\mathcal{B}_{1}(\mathcal{E})$ has at most one interval. We consider here the probability that $\mathcal{B}_{1}(\mathcal{E})$ is nontrivial, as well as the average length of the interval.

In our simulations, we find that in the limit of high novelty, persistent homology captures between $14 \%$ and $35 \%$ of recombination events, the exact value depending on mutational parameters. In typical simulations where a recombination event is detected, the bar length is well below the theoretical maximum provided by Theorem 6.25 (i), scaling roughly as the square root of novelty.

Details of the Computations. We now specify the random model of a history indexed by a source-sink loop that we use in our simulations. The model depends on parameters $m$ and $k$. Each random history generated by this model consists of: a left parent with no mutations; a right parent with mutations $\{1, \ldots, m\}$; a recombinant with some subset of these mutations;
and $k$ "intermediate sequences," each randomly sampled (with replacement) from the set

$$
\{\{1\},\{1,2\},\{1,2,3\}, \ldots,\{1, \ldots, m-1\}\} .
$$

In our simulations, we consider values of $m$ between 2 and 200, and values of $k$ between 1 and 50. In addition, we consider a "maximal sampling" scenario, in which all possible intermediate sequences were included in the sample; for the purpose of visualization, this scenario is assigned parameter value $k=51$. To construct the recombinant, we select each of the $m$ mutations with probability $\alpha$. In one set of our simulations, we set $\alpha=0.5$ (simulating a recombination breakpoint at the midpoint of the genome); in a second set of simulations, we choose $\alpha$ randomly from the uniform distribution on $[0,1]$. For each sampled history, we compute both the novelty of the recombinant and the persistent homology of the sample.

Results. The results of our simulations are given in figure Fig. 15. To obtain each subfigure, we aggregated the data for the various values of the parameter $m$. We see that the rate of detection of a recombinant increases with novelty, up to about $37 \%$ (midpoint recombination breakpoint, Fig. 15 (a)) or $20 \%$ (uniform recombination breakpoint, Fig. 15 (b)). For high novelty, increasing $k$ improves detection only up to about $k=7$, after which detection falls to about $28 \%$ (midpoint recombination breakpoint) or $16 \%$ (uniform recombination breakpoint) for high $k$.

Bar length typically falls well below the upper bound given by the novelty of the recombinant. In particular, for simulations where the recombination event was detected, bar length scales roughly as the square root of novelty (Fig. 15 (c,d)). For cases with high novelty, median bar length ranges from about $25 \%$ of the square root of novelty (if many intermediate sequences are sampled, upper right corner of Fig. $15(\mathrm{e}, \mathrm{f})$ ) to $35 \%$ of the square root of novelty (if few intermediate sequences are sampled, upper left corner of Fig. 15 (e,f)).


Figure 15. Sensitivity of persistent homology in simulations of single recombination events. Panels (a,b): Fraction of simulations in which the recombination event was detected. Recombination breakpoint is either at the midpoint of (panel a) or uniformly distributed along (panel b) the genome. Panels (c,d): Each translucent point marks the result of a single simulation in which the recombination event was detected, the red line tracks the average bar length among all simulations with equal novelty, and the black line shows the least-squares fit of parameters $a$ and $b$ among functions $y=a \sqrt{x}+b$. Panel (c): Recombination breakpoint at midpoint of the genome. The fit of all 684,026 cases is $y=0.30 \times \sqrt{x}+0.40$. Panel (d): Recombination breakpoint uniformly distributed along the genome. The fit of all 134,830 cases is $y=0.26 \times \sqrt{x}+0.51$. Panels (e,f): Median of ratio of bar length to square root of novelty, conditional on the recombination event being detected (i.e., bar length $\geq 1$ ). Recombination breakpoint is either at the midpoint of (panel e) or uniformly distributed along (panel f) the genome.

## 9. Discussion

In this paper, we have introduced novelty profiles, simple statistics of an evolutionary history which not only count the number of recombination events in the history, but also quantify the contribution recombination makes to genetic diversity. We have studied the problem of inferring information about a novelty profile from the persistent homology of sampled data, and have shown that under certain conditions, persistence barcodes of genomic data can be interpreted as lower bounds on novelty profiles. Our results provide mathematical foundations for several earlier works which have used persistent homology to study recombination.
9.1. Potential Applications of the Novelty Profile. Understanding the precise mechanisms by which recombination contributes to evolution is a long-standing problem in evolutionary biology. Progress on this problem depends critically on the availability of suitable quantitative descriptors of recombination. As a measure of the contribution of recombination to genetic diversity, the novelty profile captures information about recombination not captured by standard measures such as the recombination rate. This information may be helpful for understanding how recombination drives evolution.

We describe one class of potential applications in this direction. It is well known that reticulate evolution plays an important role both in the spread of infection and in the development of drug resistance. Examples include the emergence of a norovirus pandemic [23], outbreaks of influenza (e.g., the Swine flu pandemic of 2009) [34, 47, 51], the emergence of resistance to anti-viral medication in HIV [41], and the spread of antibiotic resistance in bacteria (e.g., E. coli) [20, 46].

The novelty profile could be useful for developing a fuller quantitative understanding of the role of recombination in such epidemiological events. We hypothesize that, compared to a count of recombination events alone, the novelty profile of a pathogen better predicts both future outbreaks of infection and the proliferation of drug resistance. It could be interesting to test this hypothesis in simulation. To test the hypothesis on real biological data, one needs well-behaved estimators of (statistics of) the novelty profile; our main bounds represent progress in this direction, and in Section 9.4 below, we discuss alternative estimation approaches which may be more practical. If the novelty profile is indeed predictive of outbreaks or of the proliferation of drug resistance, statistics derived from estimates of novelty profiles could potentially inform public health responses to infectious disease.
9.2. On the Assumptions Underlying our Main Results. Our main results relating barcodes to novelty profiles depend on strong assumptions about the evolving population and genomic sampling. In their simplest form, our results assume that the evolutionary history $\mathcal{E}$ is indexed by a galled tree and that all genomes in the history are included in our sample. Using the stability of persistent homology, we have extended these results to hold for an arbitrary sample $S$ of an arbitrary evolutionary history $\mathcal{E}$. The strength of the bounds provided by these extended results, relative to the ideal case of a galled tree with every genome sampled, is controlled by $d_{G H}(S, \mathcal{E})$, the Gromov-Hausdorff distance between $S$ and $\mathcal{E}$, and $\operatorname{Gall}(\mathcal{E})$, the number of mutations in $\mathcal{E}$ which must be ignored to obtain a history indexed by a galled tree. In cases where $d_{G H}(S, \mathcal{E})$ and $\operatorname{Gall}(\mathcal{E})$ can be assumed to be small relative to the lengths of intervals in the barcodes $\mathcal{B}_{i}(S)$, our results provide an informative
lower bound on the novelty profile, though the numerical results of Section 8.2 suggest that this bound is typically far from tight.

These results raise three key questions about applications of our work: First, under what circumstances can real-world genomic samples be expected to exhibit small enough values of $d_{G H}(S, \mathcal{E})$ and $\operatorname{Gall}(\mathcal{E})$ for our bounds on the novelty profile to be useful? Second, can the theory in this paper be extended to yield a useful topological bound on the novelty profile even when $d_{G H}(S, \mathcal{E})$ are $\operatorname{Gall}(\mathcal{E})$ not necessarily small? And third, can more sensitive bounds on the novelty profile be obtained? We discuss each of these questions below.

The Small $\operatorname{Gall}(\mathcal{E})$ Condition. While restrictive, the assumption that $\operatorname{Gall}(\mathcal{E})$ is small is biologically plausible in some settings. As shown in Appendix A in the context of the coalescent model, the phylogenetic graph $G$ indexing a history $\mathcal{E}$ will be a galled tree with high probability if and only if a relatively strong condition on the rareness of recombination is satisfied. It is important to note that this condition depends not only on the species studied, but also the sample size and how the genomic data is analyzed. For example, in the study of human recombination, a key methodological choice is the size of the genomic window used for the analysis; while a larger window offers more accurate estimates of recombination rate, a smaller window better localizes recombination breakpoints [9]. Since recombination is also rarer in a smaller window, a sample's ancestry is more likely to be represented by a galled tree when using a smaller window. In a sample of 125 humans, for instance, the ancestry in an average 275 bp window is predicted to be a galled tree with $90 \%$ probability (Appendix A). As recombination rates vary dramatically across segments of the human genome, the actual probability likewise varies.

The Small Gromov-Hausdorff Distance Condition. For a typical genomic sample $S, d_{G H}(S, \mathcal{E})$ can be large. Indeed, regardless of whether individuals are sampled simultaneously or longitudinally, the most recent common ancestor of two individuals in $S$ may be genetically distant from all individuals in $S$.

There are applications, however, where sampling is so dense, and so frequent, that we do expect common ancestors of sampled individuals to be genetically close to individuals in the sample. One such application is dense epidemiological sampling of HIV, for which entire countries have established long-term viral genomic surveillance [6,50]. In such cases, standard phylogentic methods, which assume that ancestors are absent from the sample, may produce misleading results [49]. As genetic sequencing continues to decline in cost, it is likely for dense longitudinal genomic samples to become more common. For such data sets, particularly for pathogens where evolutionary time scales are short compared to sampling duration, the assumption that $d_{G H}(S, \mathcal{E})$ is small may be more reasonable.

Even for samples $S$ for which $d_{G H}(S, \mathcal{E})$ is large, our simulation results using the coalescent model in Appendix B suggest that violations of the exact bound on the number of recombinations given by Theorem 6.25 (i) are relatively rare, in part because of the limited sensitivity of persistence barcodes in detecting recombination (Section 8.2). This, together with the empirical results from extensive simulations described in previous literature [8, 9, 13, 24], give us hope that our main theoretical results may be extended to probabilistic ones that yield useful bounds even when $d_{G H}(S, \mathcal{E})$ is large.

### 9.3. Directions for Further Theoretical Work.

Extending our Results to More Complex Phylogenetic Graphs. It may be possible to extend our bounds to histories indexed by iterated sums of phylogenetic graphs with at most $k$ recombination events, at least for small $k$. (The galled tree setting is the case $k=1$.) Such an extension would yield lower bounds on the novelty profile for a larger class of evolutionary histories. The next logical step would be to study the case $k=2$. In analogy with our main results, several questions arise about a history $\mathcal{E}$ indexed by a phylogenetic graph with two recombinants:

- Does $B_{1}(\mathcal{E})$ have at most two bars?
- For which degrees $i$ is $B_{i}(\mathcal{E})$ necessarily trivial?
- What can be said about the lengths of the intervals of $\mathcal{B}_{i}(\mathcal{E})$ ?

Tightening our Lower Bounds on the Novelty Profile. As shown numerically in Section 8.2, one limitation of barcodes as lower bounds on novelty profiles is their relatively low sensitivity to individual recombination events. A natural goal is to devise a more sensitive variant of our bounds, with similar theoretical guarantees. Is it possible to develop a consistent barcode estimator for the novelty profile, for a reasonable class of probabilistic models?

As a step in this direction, it would be interesting to apply our independence result, Theorem 8.7, to obtain analytic results about the probability distribution on $\mathcal{B}_{1}(\mathcal{E})$, for $\mathcal{E}$ a phylogenetically Markov random evolutionary history (e.g., Poisson) indexed by a galled tree. Ideally, such results would explain the relationships between novelty profiles and barcodes observed empirically in Section 8.2.
9.4. Practical Strategies for Estimating (Statistics of) Novelty Profiles. While this paper has focused primarily on bounding the novelty profile in terms of standard persistence barcodes, other approaches to estimation of the novelty profile may be more practical. There is a large literature on direct estimation of evolutionary histories (which, as noted earlier, are usually called ancestral recombination graphs); see for example [27] and the references therein. Though direct inference of histories is computationally difficult on larger data sets, recent approaches such as ARGWeaver [45] are powerful enough to yield biological insights from some real data sets consisting of dozens or hundreds of genomes. For such data sets, it may be feasible to use direct inference of a history to estimate the novelty profile. Indeed, once one has the history, computing the novelty profile is straightforward.

For larger data sets, where estimation of a full evolutionary history is not feasible, a machine learning approach may be effective. For context, a recent paper of Humphreys, McGuirl, Miyagi, and Blumberg [32] trains a regression model to predict recombination rates from Vietoris-Rips barcodes of genomic data. The training data is obtained from simulations. The authors demonstrate that this approach performs well, offering a good tradeoff between accuracy and scalability compared to a popular alternative approach called LDhelmet [12]. To estimate statistics of the novelty profile, it may be worthwhile to develop an approach analogous to that of [32].

## Appendix A. Probability that the Coalescent with Recombination Generates a Galled Tree

A.1. Overview of the Coalescent with Recombination. The coalescent with recombination is a commonly used model of the evolutionary process generating a population genetic sample. For a detailed introduction to the coalescent with recombination, see [52]. Here, we give only a brief, informal description.

Instead of tracking an entire population, which may include millions or billions of reproducing organisms, a coalescent model tracks only the sampled individuals and their direct ancestors, up to their most recent common ancestor. We can think of the coalescent with recombination as a dynamical model that generates a phylogenetic graph, together with a time function on it, by proceeding backward in time. We start with an initial set of $n$ vertices at some fixed final time, corresponding to $n$ distinct lineages. As we proceed backwards in time, we can merge two lineages by adding a vertex of in-degree one and out-degree two, representing a common ancestor. We can also split a lineage into two distinct ones by adding a vertex of out-degree one and in-degree two, representing a recombinant. We require that the phylogenetic graph we create is rooted, so any split must eventually resolve itself by a merge further back in time. Once all lineages merge into the common ancestor, the graph-generating process stops; one can then generate the mutations at each vertex, using, e.g., a Poisson-type model as in Example 8.6. However, in this section, we will be concerned only with the underlying phylogenetic graph.

Two parameters are needed to specify the coalescent with recombination's graph generation process: the number of leaves $n$ and a recombination rate parameter $\rho$. The rate parameter equals twice the expected number of recombination events occurring in the entire population, per generation.

## A.2. Probability of Generating a Galled Tree as the Solution of a Linear System.

 Let $P(n, \rho)$ denote the probability that the coalescent with recombination generates a galled tree, given the parameters $n$ and $\rho$. For fixed $n$, we derive a system of $\mathcal{O}\left(n^{2}\right)$ linear equations, depending on $\rho$, whose solution gives an analytic expression for $P(n, \rho)$ as a function of $\rho$. As $n$ grows large, this expression becomes very complicated. But each linear system is sparse, so it is easy to solve for $P(n, \rho)$ numerically for fixed values of $n$ and $\rho$, provided $n$ is not too large; see Fig. 17.In the coalescent model, there are two types of disallowed interactions whose occurrence prevents the resulting graph from being a galled tree. First, after a split occurs, one of the two resulting branches may again split, prior to the resolution of the first split (Fig. 16a). Second, two unrelated splits may occur (resulting in four parental branches), after which a branch from one split joins with a branch from the other split (Fig. 16b).

To compute the probability $P(n, \rho)$ that the coalescent generates a galled tree, we track the number of lineages $(k)$ and unresolved splits $(s)$ at each step of the process. The evolution of $k$ and $s$ is described by a discrete-time Markov chain whose state space is the finite set $T \cup\{X\}$, where

$$
T:=\left\{(k, s) \mid 1 \leq k \leq 2 n, \max (0, k-n) \leq s \leq \frac{k}{2}\right\}
$$

and $X$ is an absorbing "failure state" that we enter into when a disallowed interaction occurs. $P(n, \rho)$ is the probability that we eventually reach state $(1,0)$ in this Markov Chain, starting from $(n, 0)$.

To complete the description of the Markov chain, we first specify the transitions that can occur, and then specify the probabilities of each these. No self-transitions occur, with the exception that $X$ and $(1,0)$ are both absorbing states. No two split or join events occur simultaneously in the coalescent; each transition thus corresponds to a single split or merge.

There are two types of splits: We may have a disallowed split, as described above and illustrated in (Fig. 16a); or the split may be allowed, in which case the state transition is $(k, s) \mapsto(k+1, s+1)$ (Fig. 16c). There are three types of joins: We may have a disallowed join, as described above and illustrated in (Fig. 16b); a split may resolve itself (transition $(k, s) \mapsto(k-1, s-1))$ (Fig. 16d), or two branches may join without altering any split (transition $(k, s) \mapsto(k-1, s)$ ) (Fig. 16e).

The transition probabilities are obtained as ratios of rates: In the coalescent, for $k \geq 2$ the rate $r_{\mathrm{sp}}$ at which a split occurs is defined to be $\rho k / 2$, and the rate $r_{\mathrm{jn}}$ at which two branches join together is defined to be $k(k-1) / 2$ [52]. The total rate of a split or merge is then $r_{\text {tot }}:=r_{\mathrm{sp}}+r_{\mathrm{jn}}$, i.e., $r_{\mathrm{tot}}=k(k+\rho-1) / 2$.

The rate $r_{\text {asp }}$ of an allowed split is the product of $r_{\text {sp }}$ with the fraction of branches for which splits are allowed, i.e., $r_{\text {asp }}=\rho(k / 2-s)$. For $k \geq 2$, let $p_{S}(k, s)$ denote the probability of an allowed split (i.e., a transition $(k, s) \mapsto(k, s+1))$. Then

$$
p_{S}(k, s)=r_{\text {asp }} / r_{\text {tot }}=\frac{\rho(k-2 s)}{k(k+\rho-1)} .
$$

The other transition probabilities are obtained analogously. We denote them as follows: $p_{R}(k, s)$ is the probability that a split is resolved; $p_{J}(k, s)$ is the probability of an allowed join that does not resolve a split; and $p_{X}(k, s)$ is the probability of a disallowed interaction of either type. The formulas for these are given in Fig. 16.

For $(k, s) \in T$, let $f(k, s)$ be the probability that we eventually generate a galled tree, given that the current state is $(k, s)$, and for $(k, s) \in \mathbb{Z}^{2} \backslash T$, let $f(k, s)=0$. Thus, $P(n, \rho)=f(n, 0)$. The $f(k, s)$ satisfy the linear system

$$
\begin{align*}
f(k, s) & =p_{S}(k, s) f(k+1, s+1) \\
& +p_{R}(k, s) f(k-1, s-1) \\
& +p_{J}(k, s) f(k-1, s) \quad \text { for }(k, s) \in T \backslash\{(1,0)\},  \tag{3}\\
f(1,0) & =1
\end{align*}
$$

Fig. 17 gives values of $P(n, \rho)$ for several choices of $\rho$ and $n$, obtained by solving this linear system numerically. We observe that for fixed $n, P(n, \rho)$ tends to 1 as $\rho$ tends to 0 . Varying both parameters, $1-P(n, \rho)$ appears to decrease on the order $\rho^{2}(\log n)^{2}$ for large $n$ and small $\rho$.

To give a sense of scale for human population genetics (species effective population size $N_{e} \approx 10^{4}$, recombination rate $c \approx 10^{-} 8 \mathrm{bp}^{-1}$ ), in a sample of 125 individuals ( $n=250$ haploid genomes), to ensure that disallowed interactions occur with probability less than 0.1 $(P(n, \rho)>0.9)$, the requirement shown in Fig. 17 is $\rho<0.11$. Using $\rho=4 N_{e} c L$, where $L$ is the length of the genome segment analyzed, this requirement becomes $L<275 \mathrm{bp}$. In

> Probability
a

$p_{S}(k, s)=\frac{\rho(k-2 s)}{k(k+\rho-1)}$
$p_{R}(k, s)=\frac{2 s}{k(k+\rho-1)}$
$p_{J}(k, s)=\frac{k(k-1)-2 s(2 s-1)}{k(k+\rho-1)}$
$(k, s) \mapsto(k-1, s)$


Process terminates (genealogy is not a galled tree.)
$(k, s) \mapsto(k+1, s+1)$
d

$(k, s) \mapsto(k-1, s-1)$

Figure 16. Possible events in the coalescent process described in the text. Light gray lines indicate ordinary lineages, medium blue and dark red lines indicate pairs of parental lineages that split from their recombinant child lineages, and the black line indicates a join between the red and blue lineages. Each diagram is read upwards (back in time). To the right of each diagram are the corresponding transition probability and the resulting change in state of the Markov chain, in terms of population-scaled recombination rate $\rho$, number of lineages $k$, and number of unresolved splits $s$. Where multiple events are depicted, the probability is given only for the topmost (most ancient) event. (a) A disallowed split occurs when a lineage that has already split (light blue) splits again (dark red). (b) A disallowed join occurs when parental lineages from two separate splitting events (blue, red) join together (black). (c) An allowed split. (d) A join that resolves a split. (e) An allowed join that does not resolve a split.


Figure 17. Probability that the coalescent with recombination yields a galled tree, for several values of the recombination rate parameter $\rho$ and the number of sampled genomes $n$. We see that for fixed $n$, the probability of obtaining a galled tree tends to 1 as $\rho$ tends to 0 .
a sample of one thousand ( $n=2000$ in Fig. 17), the requirement is instead $\rho<0.08$, or $L<200 \mathrm{bp}$. Extrapolating from the apparent $\rho^{2}(\log n)^{2}$ scaling, in a sample of one million ( $n=2 \times 10^{6}$ ), the requirement is $\rho<0.04$, or $L<100 \mathrm{bp}$.

## Appendix B. Subsamples Rarely Violate our Theoretical Bounds for Complete Samples

Example 6.26 makes clear that Theorem 6.25 (i), our persistent homology lower bound on the novelty profile of an evolutionary history, does not hold for arbitrary samples $S$ of the history $\mathcal{E}$. Nevertheless, one might hope that violations of Theorem 6.25 are relatively rare. Here we use simulations to explore this question in the coalescent model with recombination.

Our computations focus on how often the number of intervals in $\mathcal{B}_{1}(S)$ exceeds the number of recombinants in $\mathcal{E}$. We find that in our simulations, this happens quite rarely, though it does occur. We did not consider the frequency of other kinds of violations of Theorem 6.25 for subsamples, though it would be interesting to do so.

We simulated over 42,000 evolutionary histories, assuming a constant population size and using parameters $n=10,15$, or $20 ; \rho=1,2,3$, or 4 ; and $\theta=5,10$, or 30 , where $\theta$ is the mutation rate parameter for the coalescent [52]. We used rejection sampling, retaining only those histories that were indexed by galled trees. Consistent with Section 2, we used an infinite-sites model of mutation. Each genetic site in a recombinant offspring inherited the state of a parent with probability one-half. For each simulated history, we counted the number of detectable recombination events that took place - defined as the number of events giving rise to a recombinant that generates an incompatibility according to the four-gamete test [31]. We then computed the maximum $\left|\mathcal{B}_{1}(S)\right|$ among 2500 random subsamples $S$ of
the history, with 5 to 30 genomes in each subsample. Among all simulations, nine had a maximum $\left|\mathcal{B}_{1}(S)\right|$ greater than the true number of detectable recombination events (Table 1). The rarity of this violation suggests that counterexamples such as Example 6.26 may be uncommon in actual population-genetic data.

Table 1. Counts of coalescent simulations, by number of detectable recombination events and maximum number of intervals in $\mathcal{B}_{1}(S)$ among all subsamples $S$. For nine of the simulations (shown in red), the maximum number of intervals in $\mathcal{B}_{1}(S)$ exceeds the number of detectable recombinations.

|  | Max. $\left\|\mathcal{B}_{1}(S)\right\|$ |  |  |  |  |  | among all subsamples $S$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | 0 | 1 | 2 | 3 | 4 | 5 |
|  | 0 | 17953 | 0 | 0 | 0 | 0 | 0 |
| Number of | 1 | 9346 | 10049 | 9 | 0 | 0 | 0 |
| detectable | 2 | 1159 | 2632 | 714 | 0 | 0 | 0 |
| recombination | 3 | 126 | 395 | 182 | 14 | 0 | 0 |
| events in history | 4 | 16 | 48 | 39 | 3 | 0 | 0 |
|  | 5 | 0 | 5 | 2 | 0 | 0 | 0 |

## References

[1] M. Adamaszek and H. Adams, The Vietoris-Rips complexes of a circle, Pacific Journal of Mathematics, 290 (2017), pp. 1-40.
[2] M. Adamaszek, H. Adams, E. Gasparovic, M. Gommel, E. Purvine, R. Sazdanovic, B. Wang, Y. Wang, and L. Ziegelmeier, Vietoris-rips and cech complexes of metric gluings, in 34th International Symposium on Computational Geometry (SoCG 2018), Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
[3] M. Adamaszek, H. Adams, and S. Reddy, On Vietoris-Rips complexes of ellipses, Journal of Topology and Analysis, (2017), pp. 1-30.
[4] M. Arenas, G. Valiente, and D. Posada, Characterization of reticulate networks based on the coalescent with recombination., Molecular biology and evolution, 25 (2008), pp. 2517-2520.
[5] U. Bauer and M. Lesnick, Induced matchings and the algebraic stability of persistence barcodes, Journal of Computational Geometry, 6 (2015), pp. 162-191.
[6] D. Bezemer, A. van Sighem, V. V. Lukashov, L. van der Hoek, N. Back, R. Schuurman, C. A. B. Boucher, E. C. J. Claas, M. C. Boerlijst, R. A. Coutinho, F. de Wolf, and ATHENA observational cohort, Transmission networks of HIV-1 among men having sex with men in the Netherlands., AIDS (London, England), 24 (2010), pp. 271-282.
[7] A. J. Blumberg and M. Lesnick, Universality of the homotopy interleaving distance, arXiv preprint arXiv:1705.01690, (2017).
[8] P. G. Cámara, A. J. Levine, and R. Rabadán, Inference of ancestral recombination graphs through topological data analysis, PLoS Comput Biol, 12 (2016), p. e1005071.
[9] P. G. Cámara, D. I. S. Rosenbloom, K. J. Emmett, A. J. Levine, and R. Rabadán, Topological Data Analysis Generates High-Resolution, Genome-wide Maps of Human Recombination., Cell systems, 3 (2016), pp. 83-94.
[10] G. Carlsson, Topology and data, Bulletin of the American Mathematical Society, 46 (2009), pp. 255-308.
[11] G. Carlsson, Topological pattern recognition for point cloud data, Acta Numerica, 23 (2014), pp. 289-368, https://doi.org/10.1017/s0962492914000051.
[12] A. H. Chan, P. A. Jenkins, and Y. S. Song, Genome-wide fine-scale recombination rate variation in drosophila melanogaster, PLoS genetics, 8 (2012), p. e1003090.
[13] J. Chan, G. Carlsson, and R. Rabadán, Topology of viral evolution, Proceedings of the National Academy of Sciences, 110 (2013).
[14] F. Chazal, D. Cohen-Steiner, M. Glisse, L. Guibas, and S. Oudot, Proximity of persistence modules and their diagrams, in Proceedings of the $25^{\text {th }}$ annual symposium on Computational geometry, ACM, 2009, pp. 237-246.
[15] F. Chazal, D. Cohen-Steiner, L. Guibas, F. Mémoli, and S. Oudot, Gromov-Hausdorff stable signatures for shapes using persistence, in Proceedings of the Symposium on Geometry Processing, Eurographics Association, 2009, pp. 1393-1403.
[16] F. Chazal, V. de Silva, M. Glisse, and S. Oudot, The Structure and Stability of Persistence Modules, Springer International Publishing, 2016, https://doi.org/10.1007/ 978-3-319-42545-0, https://doi.org/10.1007\%2F978-3-319-42545-0.
[17] F. Chazal, V. De Silva, and S. Oudot, Persistence stability for geometric complexes, Geometriae Dedicata, 173 (2014), pp. 193-214.
[18] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer, Stability of persistence diagrams, Discrete and Computational Geometry, 37 (2007), pp. 103-120.
[19] W. Crawley-Boevey, Decomposition of pointwise finite-dimensional persistence modules, Journal of Algebra and Its Applications, 14 (2015), p. 1550066.
[20] J. Davies and D. Davies, Origins and evolution of antibiotic resistance, Microbiology and Molecular Biology Reviews, 74 (2010), pp. 417-433.
[21] R. Durrett, Probability: theory and examples, Cambridge University Press, 2010.
[22] H. Edelsbrunner and J. Harer, Computational topology: an introduction, American Mathematical Society, 2010.
[23] J. S. Eden, M. M. Tanaka, M. F. Boni, W. D. Rawlinson, and P. A. White, Recombination within the pandemic norovirus gii. 4 lineage, Journal of Virology, 87 (2013), pp. 6270-6282.
[24] K. Emmett, D. Rosenbloom, P. Cámara, and R. Rabadán, Parametric inference using persistence diagrams: A case study in population genetics, Proc. 31st Intl. Conf. Machine Learning, (2014), https://arxiv.org/abs/C21B8DD4-D6F5-4057-ADCC-BAF5DCC28A7B.
[25] K. J. Emmett and R. Rabadán, Characterizing Scales of Genetic Recombination and Antibiotic Resistance in Pathogenic Bacteria Using Topological Data Analysis, Lecture Notes in Computer Science, 8609 (2014), pp. 540-551.
[26] R. Forman, A user's guide to discrete Morse theory, Sém. Lothar. Combin., 48 (2002), pp. Art. B48c, 35.
[27] D. Gusfield, ReCombinatorics: the algorithmics of ancestral recombination graphs and explicit phylogenetic networks, MIT Press, 2014.
[28] D. Gusfield, S. Eddhu, and C. Langley, Efficient reconstruction of phylogenetic networks with constrained recombination, in Bioinformatics Conference, 2003. CSB 2003. Proceedings of the 2003 IEEE, IEEE, 2003, pp. 363-374.
[29] S. Harker, M. Kramár, R. Levanger, and K. Mischaikow, A comparison framework for interleaved persistence modules, Journal of Applied and Computational Topology, 3 (2019), pp. 85-118, https://doi.org/10.1007/s41468-019-00026-x, https://doi.org/10.1007/ s41468-019-00026-x.
[30] A. Hatcher, Algebraic topology, Cambridge University Press, 2002.
[31] R. R. Hudson and N. L. Kaplan, Statistical properties of the number of recombination
events in the history of a sample of DNA sequences., Genetics, 111 (1985), pp. 147-164.
[32] D. P. Humphreys, M. R. McGuirl, M. Miyagi, and A. J. Blumberg, Fast estimation of recombination rates using topological data analysis, Genetics, (2019), pp. genetics-301565.
[33] D. H. Huson, R. Rupp, and C. Scornavacca, Phylogenetic networks: concepts, algorithms and applications, Cambridge University Press, 2010.
[34] T. Ito, J. N. S. S. Couceiro, S. Kelm, L. G. Baum, S. Krauss, M. R. Castrucci, I. Donatelli, H. Kida, J. C. Paulson, R. G. Webster, and Y. Kawaoka, Molecular basis for the generation in pigs of influenza A viruses with pandemic potential, Journal of Virology, 72 (1998), pp. 7367-7373.
[35] M. Kahle, Random geometric complexes, Discrete \& Computational Geometry, 45 (2011), pp. 553-573.
[36] O. Kallenberg, Foundations of modern probability, Springer Science \& Business Media, 2006.
[37] D. Kozlov, Combinatorial algebraic topology, vol. 21 of Algorithms and Computation in Mathematics, Springer, Berlin, 2008, https://doi.org/10.1007/978-3-540-71962-5, http: //dx.doi.org/10.1007/978-3-540-71962-5.
[38] S. L. Lauritzen, Graphical Models, Oxford University Press, 1996.
[39] J. Munkres, Elements of Algebraic Topology, Prentice Hall, 1984, http://www.worldcat. org/isbn/0131816292.
[40] S. R. Myers and R. C. Griffiths, Bounds on the minimum number of recombination events in a sample history., Genetics, 163 (2003), pp. 375-394.
[41] T. Nora, C. Charpentier, O. Tenaillon, C. Hoede, F. Clavel, and A. J. Hance, Contribution of recombination to the evolution of human immunodeficiency viruses expressing resistance to antiretroviral treatment, Journal of virology, 81 (2007), pp. 7620-7628.
[42] S. Y. Oudot, Persistence Theory: From Quiver Representations to Data Analysis, no. 209 in AMS Mathematical Surveys and Monographs, American Mathematical Society, 2015.
[43] L. Parida, F. Utro, D. Yorukoglu, A. P. Carrieri, D. Kuhn, and S. Basu, Topological signatures for population admixture, in International Conference on Research in Computational Molecular Biology, Springer, 2015, pp. 261-275.
[44] R. Rabadán and A. J. Blumberg, Topological Data Analysis for Genomics and Evolution: Topology in Biology, Cambridge University Press, 2019.
[45] M. D. Rasmussen, M. J. Hubisz, I. Gronau, and A. Siepel, Genome-Wide Inference of Ancestral Recombination Graphs, PLoS Genetics, 10 (2014), p. e1004342.
[46] H. Rohde, J. Qin, Y. Cui, D. Li, N. J. Loman, M. Hentschke, W. Chen, F. Pu, Y. Peng, J. Li, et al., Open-source genomic analysis of shiga-toxin-producing e. coli o104: H4, New England Journal of Medicine, 365 (2011), pp. 718-724.
[47] A. Solovyov, G. Palacios, T. Briese, W. I. Lipkin, and R. Rabadán, Cluster analysis of the origins of the new influenza $A$ (H1N1) virus, Euro surveillance: bulletin Europeen sur les maladies transmissibles $=$ European communicable disease bulletin, 14 (2009).
[48] Y. S. Song, Y. Wu, and D. Gusfield, Efficient computation of close lower and upper bounds on the minimum number of recombinations in biological sequence evolution., Bioinformatics (Oxford, England), 21 Suppl 1 (2005), pp. i413-22.
[49] T. Stadler, R. Kouyos, V. von Wyl, S. Yerly, J. Boni, P. Burgisser, T. Klimkait, B. Joos, P. Rieder, D. Xie, H. F. Gunthard, A. J. Drummond, S. Bonhoeffer, and the Swiss HIV Cohort Study, Estimating the Basic Reproductive Number from Viral Sequence Data, Molecular biology and evolution, 29 (2011), pp. 347-357.
[50] The Swiss HIV Cohort Study, Cohort profile: the Swiss HIV Cohort study, International Journal of Epidemiology, 39 (2010), pp. 1179-1189.
[51] V. Trifonov, H. Khiabanian, R. Rabadán, et al., Geographic dependence, surveillance,
and origins of the 2009 influenza a (h1n1) virus., New England Journal of Medicine, 361 (2009), pp. 115-119.
[52] J. Wakeley, Coalescent Theory, An Introduction, Roberts \& Co., 2007.
[53] L. Wang, K. Zhang, and L. Zhang, Perfect phylogenetic networks with recombination, Journal of Computational Biology, 8 (2001), pp. 69-78.

Lesnick: SUNY Albany, NY, USA
E-mail address: mlesnick@albany.edu
Rabadán: Columbia University, New York, NY, USA
E-mail address: rr2579@cumc.columbia.edu
Rosenbloom: Merck Research Laboratories, Rahway, nJ, USA
E-mail address: daniel.rosenbloom@merck.com

