

Max-Planck-Institut  
für Mathematik  
in den Naturwissenschaften  
Leipzig

Tropical Sufficient Statistics for  
Persistent Homology

by

*Anthea Monod, Sara Kališnik Verovšek,  
Juan Ángel Patiño-Galindo, and Lorin Crawford*

Preprint no.: 66

2017





# Tropical Sufficient Statistics for Persistent Homology

Anthea Monod<sup>1,†</sup>, Sara Kališnik Verovšek<sup>2</sup>, Juan Ángel Patiño-Galindo<sup>1</sup>, and Lorin Crawford<sup>3,4,5</sup>

**1 Department of Systems Biology, Columbia University, New York, NY, USA**

**2 Research Group in Nonlinear Algebra, Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany**

**3 Department of Biostatistics, Brown University, Providence, RI, USA**

**4 Center for Statistical Sciences, Brown University, Providence, RI, USA**

**5 Center for Computational Molecular Biology, Brown University, Providence, RI, USA**

† Corresponding e-mail: am4691@cumc.columbia.edu

## Abstract

In this paper, we show that an embedding in Euclidean space based on tropical algebraic geometry generates stable sufficient statistics for barcodes. Conventionally, barcodes are multiscale summaries of topological characteristics that capture the “shape” of data; however, in practice, they have complex structures which make them difficult to use in statistical settings. The sufficiency result presented in this work allows for classical probability distributions to be assumed on the tropicalized representation of barcodes. This makes a variety of parametric statistical inference methods amenable to barcodes, all while maintaining their initial interpretations. More specifically, we show that exponential family distributions may be assumed, and that likelihood functions for persistent homology may be constructed. We conceptually demonstrate sufficiency and illustrate its utility in persistent homology dimensions 0 and 1 with concrete parametric applications to HIV and avian influenza data.

## 1 Introduction

In this paper, we provide statistical sufficiency in Euclidean space for persistent homology — an important concept in the field of topological data analysis (TDA) that summarizes the “shape” and “size” of data. Our result is based on a topological embedding given by functions defined in tropical algebraic geometry. With these sufficient summaries, probability distributions from classical statistics may be applied to persistent homology. More importantly, extensive statistical methodology and parametric inference become accessible to TDA.

Recently, TDA has become particularly relevant due to its theoretical foundations that allow for dimensionality reduction with qualitative and robust summaries of observed data. To this end, it is applicable to a wide range of complex data structures that arise in various domains of data science. Persistent homology is an important topological invariant upon which many TDA methods depend when applied in practice (Carlsson and Zomorodian, 2005). Specifically, this method has been used to address problems in fields ranging from sensor networks (Ghrist and de Silva, 2006; Adams and Carlsson, 2014), medicine (Ferri and Stanganelli, 2010; Adcock et al., 2014), neuroscience (Chung et al., 2009; Curto et al., 2013; Giusti et al., 2015), as well as imaging analysis (Perea and Carlsson, 2014).

Persistent homology defines a proper metric space, but since it produces a collection of intervals (known as barcodes) rather than numerical quantities, it has been difficult to apply the method to parametric data analysis. Recently, there have been substantial advancements in topological methodology and theory to bypass this issue. One natural approach to integrate the shape information of data (via barcodes) into existing computational machinery has been to form vectors using metrics defined on the space of barcodes. Other vectorization techniques have also been developed with specific properties that aim to accomplish slightly different objectives. For example, Carrière et al. (2015) construct feature representations by rearranging entries of a distance matrix between points in a persistence diagram. Bubenik (2015) developed persistence landscapes that produce elements in Banach space. Though there are concerns of stability, methods proposed by Bendich et al. (2016) use a binning approach to obtain a vector representation of features. Reininghaus et al. (2015), Kwitt et al. (2015), and Adams et al. (2017) present methods that are similar in nature where

they construct stable surfaces of persistence diagrams based on kernel methods with the aim to incorporate persistence information into support vector machines and other machine learning algorithms. Alternatively, Crawford et al. (2017) integrate persistence information as a collection of smoothed Euler characteristic summaries which are then used as infinite-dimensional statistics in nonlinear functional regression.

Summary statistics are relevant to the notion of sufficiency, which is a desirable property in studies centered around inference. Sufficiency allows a given sample of data to be mapped to a less complex or lower dimensional space (i.e. harboring less computational burden) without the loss of information (Bartlett, 1937). Moreover, sufficient statistics provide the functional form of probability distributions via a classical factorization criterion. In other words, they represent the key basis for complete parametric inference in statistics. Summary statistics were initially explored in the context of persistent homology to study shapes and surfaces (closed, compact subsets) with  $S$  in  $\mathbf{R}^2$  and  $\mathbf{R}^3$  (Turner et al., 2014). These quantities were constructed by mapping from  $\{S \subset \mathcal{M}_d\} \times \mathcal{S}^{d-1}$  to  $\mathcal{D}^d$  — where  $\mathcal{M}_d$ ,  $d = 2, 3$ , denotes the set of all closed, compact subsets of  $\mathbf{R}^d$  with a finite simplicial complex representation,  $\mathcal{S}^{d-1}$  is the unit sphere, and  $\mathcal{D}^d$  is the space of all  $d$ -dimensional persistence diagrams. This construction provided a notion of summary statistics for the family of probability distributions on shapes and surfaces in 2 and 3-dimensions. Perhaps more importantly, this effort also shed light on the possibility of proving statistical sufficiency for topological quantities.

The main contribution of this work is that we provide sufficiency for the general family of probability measures on persistent homology. A fundamental observation that makes this result possible is the formal probabilistic characterization of both the domain (i.e. the space of barcodes) and codomain (i.e. Euclidean space) of our mapping (Mileyko et al., 2011). In principle, inference should proceed from a probability model defined directly on the barcodes; however, this type of specification is complicated. Given sufficient statistics for barcodes, we may use the likelihood principle to proceed with parametric inferences directly. As has been previously demonstrated, this suggests that a generative or sampling model on barcodes is possible in persistent homology (Adler et al., 2017).

The remainder of this paper is organized as follows. In Section 2, we give a formal definition of persistent homology and outline the approach for vectorizing persistence barcodes using tropical algebraic geometry. We also provide the mathematical characteristics of this method. In Section 3, we provide our main result of a sufficient statistic for persistence barcodes. Here, we give the definition of an exponential family of probability distributions and likelihood functions based on this result for persistent homology. Section 4 gives two concrete applications to data in dimensions 0 and 1. The first is a practical demonstration of sufficiency based on HIV data, while the second shows how a parametric assumption statistically quantifies the biological distinction between intra- and intersubtype reassortment in avian influenza. We close with a discussion in Section 5 with directions for future research.

## 2 Persistent Homology & Vectorization

In this section, we give the mathematical background to our main result. Specifically, we provide details on the space of barcodes arising from persistent homology and describe the construction of tropical coordinates on the space of persistence barcodes.

### 2.1 The Space of Persistence Barcodes

Homology groups were developed in classical topology to “measure” the shape of spaces by abstractly counting the occurrences of patterns (e.g. the number of connected components, loops, and voids in three dimensions). *Persistent homology* or *persistence* (Frosini and Landi, 1999; Edelsbrunner et al., 2002; Carlsson and Zomorodian, 2005) is the adaptation of classical homology to point clouds (i.e. finite metric spaces), which is often the form in which data are collected. Persistence, as a topological construct, produces a robust summary of the data that is invariant to noisy perturbations (Cohen-Steiner et al., 2007). Moreover, it also captures integral geometric features, or “size”, of the data. We now describe the construction of homology and its extension to persistent homology.

**Simplicial Homology & Persistence.** Simplicial homology studies the shape of simplicial complexes, which can be seen as skeletal representations of data types, such as images. It is an important building block in computing persistent homology. A *simplicial complex* is a collection  $K$  of non-empty subsets of a set  $K_0$  such that  $\tau \subset \sigma$  and  $\sigma \in K$  guarantees that  $\tau \in K$ . The elements of  $K_0$  are called *vertices* of  $K$ , and the elements of  $K$  are called *simplices*. A simplex has dimension  $k$ , or is a *k-simplex*, if it has a cardinality of  $k + 1$ . Simplicial homology is constructed by considering simplicial *k-chains*, which are linear combinations over a field  $\mathbf{F}$  of  $k$ -simplices in finite  $K$ . A set of  $k$ -chains defines a vector space  $C_k(K)$ . The *boundary map* is denoted as

$$\begin{aligned} \partial_k : C_k(K) &\rightarrow C_{k-1}(K) \\ \partial_k([v_0, v_1, \dots, v_k]) &= \sum_{i=0}^k [v_0, \dots, v_{-i}, \dots, v_k] \end{aligned}$$

with linear extension, where  $v_{-i}$  indicates that the  $i^{\text{th}}$  element is dropped. *Boundaries* are elements of  $B_k(K) = \text{im } \partial_{k+1}$ , and *cycles* are elements of  $Z_k(K) = \text{ker } \partial_k$ .

**Definition 1.** The  $k^{\text{th}}$  *homology group* of  $K$  is given by the quotient group

$$H_k(K) := Z_k(K)/B_k(K).$$

The motivating idea underlying homology is to account for the structure of  $K$ , which is the finite simplicial complex representation of a topological space  $X$ . In dimension zero, elements representing connected components of  $X$ , with finite simplicial complex representation  $K$ , generate the zeroth homology group  $H_0(X)$ . If, for example,  $X$  has two connected components, then  $H_0(X) \cong \mathbf{F} \oplus \mathbf{F}$  ( $\cong$  here denotes group isomorphism). For higher dimensions  $k \geq 1$ ,  $k$ -dimensional holes are the result of considering the boundary of a  $(k + 1)$ -dimensional object.  $H_k(X)$  is generated by elements that represent  $k$ -dimensional holes in  $X$ .

Intuitively, persistent homology is computed by tracking the progression of connected components, loops, and higher dimensional voids, with respect to a *filtration* assigned to the observed point cloud — that is, a finite nested sequence of simplicial complexes  $\mathcal{K} := \{K_r\}_{r=a}^b$  indexed by a parameter  $r \in \mathbf{R}^+$  such that  $K_{r_1} \subseteq K_{r_2}$  if  $r_1 < r_2$ .

**Definition 2.** Let  $K$  be a filtered simplicial complex  $K_1 \subset K_2 \subset \dots \subset K_t = K$ . The  $k^{\text{th}}$  *persistence module derived in homology*, or simply  $k^{\text{th}}$  *persistent homology* (for short), of  $K$  is given by

$$\text{PH}_k(K) := \{H_k(K_r)\}_{1 \leq r \leq t}$$

together with the collection of linear maps  $\{\varphi_{r,s}\}_{1 \leq r \leq s \leq t}$ , where  $\varphi_{r,s} : H_k(K_r) \rightarrow H_k(K_s)$  is induced by the inclusion  $K_r \hookrightarrow K_s$  for all  $r, s \in \{1, 2, \dots, t\}$  such that  $r \leq s$ .

Persistent homology contains the homology information on individual spaces  $\{K_r\}$  and the mappings between their homologies for every  $K_r$  and  $K_s$  such that  $r \leq s$ . Note that the motivating idea underlying persistent homology is to account for the homology of  $X$  simultaneously across multiple scales. Rather than restricting the analysis to one static instance, persistence tracks how the topological structure evolves over the indexed filtration. The final output after computing persistence is a *barcode* (i.e. a collection of intervals). Barcodes are random objects when the point cloud that generates them are randomly sampled data points. Each interval (or bar) corresponds to a topological feature that appears (is born) at the value of a parameter given by the left endpoint of the interval, and disappears or merges with another existing feature (dies) at the value given by the right endpoint.

Barcodes can alternatively be seen as summary statistics of the data generating process because they appropriately allow for a reduction in dimensionality of the ambient space. Therefore, in a broader sense, topological approaches can be used to provide sets of principal variables in data science problems where the total number of features greatly exceeds the number of collected observations.

**Barcodes & Persistence Diagrams.** A barcode with  $\ell$  intervals can be written as  $(x_1, d_1, x_2, d_2, \dots, x_\ell, d_\ell)$  where  $x_i$  is the left endpoint of the  $i^{\text{th}}$  interval and  $d_i$  denotes the length. In this paper, we assume that

$x_i \geq 0$  such that the birth times of bars are positive. Notice that the order in which coordinates  $(x_i, d_i)$  appear within the collection  $(x_1, d_1, x_2, d_2, \dots, x_\ell, d_\ell)$  does not affect the content of the topological information encoded in the barcode. Hence, the indices  $i = 1, 2, \dots, \ell$  are essentially dummy variables. This means that all permutations of coordinates  $(x_i, d_i)$  across  $\ell$  positions define the same collection. Algebraically, this amounts to taking the orbit space of the action of the symmetric group on  $\ell$  letters on the product  $([0, \infty) \times [0, \infty))^\ell$  given by permuting the coordinates. Here, we denote this orbit space by  $B_\ell$ .

**Definition 3.** The *barcode space*  $\mathcal{B}_{\leq n}$  consisting of bars with at most  $n$  intervals is given by the quotient

$$\coprod_{\ell \in \mathbf{N}_{\leq n}} B_\ell / \sim,$$

where  $\sim$  is generated by equivalences of the form

$$\{(x_1, d_1), (x_2, d_2), \dots, (x_\ell, d_\ell)\} \sim \{(x_1, d_1), (x_2, d_2), \dots, (x_{\ell-1}, d_{\ell-1})\},$$

whenever  $d_\ell = 0$ .

Notice that we may set  $y_i := x_i + d_i$  to be the right endpoint of the  $i^{\text{th}}$  interval, and equivalently consider the collection of ordered pairs  $(x_i, y_i)$  together with the diagonal  $\Delta = \{(x, y) \mid x = y\}$ , where each point is taken with infinite multiplicity. Here,  $\Delta$  contains the bars of length zero. These correspond to topological noise since they are features that are born and then die immediately. By plotting all ordered pairs on a scatterplot, we obtain an alternative representation of the barcode, known as a *persistence diagram*.

**Metrics on Barcode Space.** The space of barcodes is a metric space. Distances between barcodes are defined by specifying the distance between any pair of intervals, including the set of zero length intervals  $\Delta = \{(x, x) \mid -\infty < x < \infty\}$ . Let

$$d_\infty((x_i, d_i), (x_j, d_j)) = \max\{|x_i - x_j|, |d_i - d_j + x_i - x_j|\}$$

be the distance between two intervals. The distance between an interval and the set  $\Delta$  is then

$$d_\infty((x, d), \Delta) = \frac{d}{2}.$$

The factor of  $1/2$  comes from the idea that points close to the diagonal on a persistence diagram will tend to merge with the diagonal. For instance, when considering the  $\ell^2$  or  $\ell^\infty$  norm, the closest point to a diagonal is the intersection of a line through (i) the point that is perpendicular to the diagonal, and (ii) the diagonal itself (e.g. Chazal et al., 2016).

Now let  $\mathcal{B}_i = \{I_\alpha\}_{\alpha \in A}$  and  $\mathcal{B}_j = \{J_\beta\}_{\beta \in B}$  be two barcodes. We give the definition of two important metrics frequently used for computations on barcode space below.

**Definition 4.** For a fixed  $p$ , finite sets  $A$  and  $B$ , any bijection  $\vartheta$  from a subset  $A' \subseteq A$  to  $B' \subseteq B$ , and a penalty on  $\vartheta$  set to

$$P_p(\vartheta) := \sum_{a \in A'} d_\infty(I_a, J_{\vartheta(a)})^p + \sum_{a \in A \setminus A'} d_\infty(I_a, \Delta)^p + \sum_{b \in B \setminus B'} d_\infty(I_b, \Delta)^p,$$

then

$$d_p^W(\mathcal{B}_i, \mathcal{B}_j) := \left( \min_{\vartheta} P_p(\vartheta) \right)^{\frac{1}{p}},$$

yields the *Wasserstein  $p$ -distance* ( $p \geq 1$ ) between  $\mathcal{B}_i$  and  $\mathcal{B}_j$ .

**Definition 5.** For finite subsets and sets  $A' \subseteq A$  and  $B' \subseteq B$ , bijection  $\vartheta$  as previously defined above, and a penalty set to

$$P_\infty(\vartheta) := \max \left\{ \max_{a \in A'} \{d_\infty(I_a, J_{\vartheta(a)})\}, \max_{a \in A \setminus A'} d_\infty(I_a, \Delta), \max_{b \in B \setminus B'} d_\infty(I_b, \Delta) \right\},$$

then

$$d_\infty^B(\mathcal{B}_i, \mathcal{B}_j) := \min_{\vartheta} P_\infty(\vartheta),$$

yields the *bottleneck distance* between  $\mathcal{B}_i$  and  $\mathcal{B}_j$ , where the minimum is taken over all bijections from subsets  $A'$  to  $B'$ .

**Regularizing Subsets of Barcode Space.** In this paper, for regularity conditions, we consider the following subsets of barcode space. For a fixed  $m > 0$ , denote  $\mathcal{B}_{\leq n}^m$  to be the subset of  $\mathcal{B}_{\leq n}$  consisting of those  $(x_1, d_1, x_2, d_2, \dots, x_\ell, d_\ell)$  with  $d_i > 0$  such that

$$x_i \leq md_i \tag{1}$$

for  $i = 1, 2, \dots, \ell$ . Given a finite set of barcodes, the value of  $m$  is straightforward to determine. The motivation for considering this subset is to avoid computational difficulties with the infinite multiplicity of computations directly involving the diagonal set  $\Delta$ . While this might appear restrictive at first, it does not really pose any limitations in practice. In fact, for data generated by some finite process (e.g. meshes have a finite number of vertices/faces, images have limited resolution, etc.), establishing  $m$  requires minimal consideration.

## 2.2 Coordinates on the Space of Persistence Barcodes

In an effort to integrate shape information of data (via barcodes) into existing computational machinery and methodology, recent vectorization results have been developed. Related to the approach that we adopt in this paper, persistence diagrams have been vectorized via the construction of complex polynomials where the points of the persistence diagrams are the roots (Ferri and Landi, 1999; Di Fabio and Ferri, 2015). In a more relevant approach, vectors have been assigned directly to barcodes by defining a ring of algebraic functions on the space of barcodes satisfying certain properties (Adcock et al., 2016). However, these functions (polynomials) are not Lipschitz with respect to the Wasserstein  $p$ - and bottleneck distances (see Appendix A1 for a formal proof). Lipschitz continuity is a desirable property for the transformation of barcodes, since it guarantees stability in the target space under small perturbations in the domain, therefore making the transformation computationally robust. As an alternative solution to address the shortcomings of formerly defined polynomials, tropical functions on the space of barcodes have been shown to be Lipschitz (and therefore continuous) with respect to the Wasserstein  $p$ -distance and bottleneck distance (Kališnik Verovšek, 2016). We now detail the construction of tropical functions as coordinates on barcode space.

**Fundamentals of Tropical Algebra.** Tropical algebra is a branch of mathematics based on the study of particular semirings, as defined below.

**Definition 6.** The tropical (equivalently, min-plus) semiring is given by  $(\mathbf{R} \cup \{+\infty\}, \oplus, \odot)$ , with addition and multiplication being defined as follows:

$$a \oplus b := \min \{a, b\} \quad \text{and} \quad a \odot b := a + b,$$

where both operations are commutative and associative, and multiplication distributes over addition. Similarly, there also exists the arctic (equivalently, max-plus) semiring  $(\mathbf{R} \cup \{-\infty\}, \boxplus, \odot)$ , where multiplication of two elements is defined as in the case of the tropical semiring, but addition amounts to taking the maximum instead of the minimum. Namely:

$$a \boxplus b := \max \{a, b\} \quad \text{and} \quad a \odot b := a + b,$$

where again both operations are associative, commutative, and distributive.

As in the case of ordinary polynomials that are formed by multiplying and adding real variables, max-plus polynomials can be formed by multiplying and adding variables in the max-plus semiring as follows. Let  $x_1, x_2, \dots, x_n$  be variables that are elements in the max-plus semiring. A *max-plus monomial expression* is any product of these variables, where repetition is permitted. By commutativity, we can sort these products and write monomial expressions with the variables raised to exponents:

$$p(x_1, x_2, \dots, x_n) = a_1 \odot x_1^{a_1^1} x_2^{a_2^1} \dots x_n^{a_n^1} \boxplus a_2 \odot x_1^{a_1^2} x_2^{a_2^2} \dots x_n^{a_n^2} \boxplus \dots \boxplus a_m \odot x_1^{a_1^m} x_2^{a_2^m} \dots x_n^{a_n^m}.$$

Here the coefficients  $a_1, a_2, \dots, a_m$  are in  $\mathbf{R}$ , and the exponents  $a_j^i$  for  $1 \leq j \leq n$  and  $1 \leq i \leq m$  are in  $\mathbf{Z}^+$ .

It is easy to see that max-plus polynomial expressions do not uniquely define functions (e.g. Kališnik Verovšek and Carlsson, 2014). Thus, for  $p$  and  $q$  max-plus polynomial expressions, if

$$p(x_1, x_2, \dots, x_n) = q(x_1, x_2, \dots, x_n)$$

for all  $(x_1, x_2, \dots, x_n) \in (\mathbf{R} \cup \infty)^n$ , then  $p$  and  $q$  are said to be *functionally equivalent*, and we write  $p \sim q$ . Max-plus polynomials are the semiring of equivalence classes of max-plus polynomial expressions with respect to  $\sim$ .

**Identifying Tropical Functions for Barcodes.** We can apply the preceding observation to the setting of barcodes as follows. Fix  $n$  and let the symmetric group  $S_n$  act on the matrix of indeterminates

$$X = \begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \\ \vdots & \vdots \\ x_{n1} & x_{n2} \end{pmatrix}$$

by left multiplication. The matrix  $X$  represents a barcode with  $n$  bars, where each bar is represented by a row entry. As mentioned previously, the action of  $S_n$  on  $X$  amounts to permuting the order of the bars (as coordinates  $x_i, d_i$ ) within  $(x_1, d_1, x_2, d_2, \dots, x_n, d_n)$ .

Consider the collection of *exponent* matrices

$$\mathcal{E}_n = \left\{ \begin{pmatrix} e_{11} & e_{12} \\ e_{21} & e_{22} \\ \vdots & \vdots \\ e_{n1} & e_{n2} \end{pmatrix} \neq [0]_n^2 \mid e_{ij} \in \{0, 1\} \text{ for } i = 1, 2, \dots, n, \text{ and } j = 1, 2 \right\}.$$

A matrix  $E \in \mathcal{E}_n$ , together with  $X$ , determines a max-plus monomial by  $P(E) = x_{11}^{e_{11}} x_{12}^{e_{12}} \dots x_{n1}^{e_{n1}} x_{n2}^{e_{n2}}$ . Notice that  $P(E)$  defines a max-plus monomial by singling out individual elements  $x_{ij}$  in  $X$ , since elements  $e_{ij}$  of  $E$  exponentiate  $x_{ij} \in X$  to the power of either 0 or 1. Again, recall that the multiplication and exponentiation are calculated according to the operations in the max-plus semiring.

Denote the set of orbits under the row permutation action on  $\mathcal{E}_n$  by  $\mathcal{E}_n/S_n$ . The orbits  $\{E_1, E_2, \dots, E_m\}$  determine a *2-symmetric* max-plus polynomial by

$$P(E_1) \boxplus P(E_2) \boxplus \dots \boxplus P(E_m).$$

Let  $E_{(e_{11}, e_{12}), (e_{21}, e_{22}), \dots, (e_{n1}, e_{n2})}$  denote the polynomial that arises from the orbit

$$\left[ \begin{pmatrix} e_{11} & e_{12} \\ e_{21} & e_{22} \\ \vdots & \vdots \\ e_{n1} & e_{n2} \end{pmatrix} \right].$$

**Proposition 1.** Let  $[(x_1, d_1, x_2, d_2, \dots, x_n, d_n)]$  and  $[(x'_1, d'_1, x'_2, d'_2, \dots, x'_n, d'_n)]$  be two orbits under the row permutation action on  $\mathbf{R}^{2n}$ . If

$$E_{(0,1)^i(1,1)^j} [(x_1, d_1, x_2, d_2, \dots, x_n, d_n)] = E_{(0,1)^i(1,1)^j} [(x'_1, d'_1, x'_2, d'_2, \dots, x'_n, d'_n)]$$

for all  $i, j \leq n$ , then

$$[(x_1, d_1, x_2, d_2, \dots, x_n, d_n)] = [(x'_1, d'_1, x'_2, d'_2, \dots, x'_n, d'_n)].$$

*Proof.* Assume without loss of generality that  $d_1 \leq d_2 \leq \dots \leq d_n$  and  $d'_1 \leq d'_2 \leq \dots \leq d'_n$ . Since

$$E_{(0,1)} [(x_1, d_1, x_2, d_2, \dots, x_n, d_n)] = E_{(0,1)} [(x'_1, d'_1, x'_2, d'_2, \dots, x'_n, d'_n)],$$



it follows that  $d_n = d'_n$ . Since

$$\begin{aligned} E_{(0,1)^2}[(x_1, d_1, x_2, d_2, \dots, x_n, d_n)] - E_{(0,1)^1}[(x_1, d_1, x_2, d_2, \dots, x_n, d_n)] &= \\ &= E_{(0,1)^2}[(x_1, d'_1, x_2, d'_2, \dots, x_n, d'_n)] - E_{(0,1)^1}[(x_1, d'_1, x_2, d'_2, \dots, x_n, d'_n)], \end{aligned}$$

it follows that  $d_{n-1} = d'_{n-1}$ . We continue and get  $d_i = d'_i$  for all  $i$ .

The fact that the greatest element of  $\{x_1, x_2, \dots, x_n\}$  equals the greatest element of  $\{x'_1, x'_2, \dots, x'_n\}$  follows by applying  $E_{(0,1)^{n-1}(1,1)}$ . Computing  $E_{(0,1)^{n-2}(1,1)^2} - E_{(0,1)^{n-1}(1,1)}$  yields the equality of second greatest elements in respective sets. We continue in this manner and get that  $\{x_1, x_2, \dots, x_n\}$  is a permutation of  $\{x'_1, x'_2, \dots, x'_n\}$ .

Showing that  $\{(x_1, d_1), (x_2, d_2), \dots, (x_n, d_n)\}$  is a permutation of  $\{(x'_1, d'_1), (x'_2, d'_2), \dots, (x'_n, d'_n)\}$  follows from the proof of Proposition 5.8 in Kališnik Verovšek and Carlsson (2014).  $\square$

A corollary of this result is that not all orbits under the row permutation action are needed to define separating polynomials for nonequivalent barcodes.

**Tropical Coordinates on Barcode Space.** The above Proposition, together with Kališnik Verovšek (2016), shows that defining a vector with elements given by

$$E_{m,(1,1)^i,(0,1)^j}(x_1, d_1, x_2, d_2, \dots, x_n, d_n) := E_{(1,1)^i,(0,1)^j}(x_1 \oplus d_1^m, d_1, x_2 \oplus d_2^m, d_2, \dots, x_n \oplus d_n^m, d_n) \quad (2)$$

induces an injective map on  $\mathcal{B}_{\leq n}^m$ . Hence, this collection of functions separates nonequivalent barcodes. These functions are also Lipschitz with respect to the Wasserstein  $p$ - and bottleneck distances (Kališnik Verovšek, 2016). We now give an example for calculating these functions and evaluating them on barcodes.

*Example.* Fix  $n = 2$ . The set of orbits under the  $S_2$  action is

$$\mathcal{E}_2/S_2 = \left\{ \begin{array}{l} \left[ \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \right], \left[ \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \right], \left[ \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \right], \left[ \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix} \right], \\ \left[ \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \right], \left[ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right], \left[ \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix} \right], \left[ \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \right], \left[ \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \right] \end{array} \right\}.$$

According to Proposition 1, we need only consider the functions defined by a subset of  $\mathcal{E}_2/S_2$  to separate barcodes when  $n = 2$ . We may consider

$$\left[ \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \right], \left[ \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix} \right], \left[ \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix} \right], \left[ \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \right], \left[ \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \right]. \quad (3)$$

Suppose that we have two barcodes  $\mathcal{B}_1 = \{(1, 2), (3, 1)\}$  and  $\mathcal{B}_2 = \{(2, 2)\}$ , where  $\mathcal{B}_1, \mathcal{B}_2 \in \mathcal{B}_{\leq 2}$ .

1. Compute  $m$  from (1): For intervals  $(1, 2)$ ,  $(3, 1)$ , and  $(2, 2)$ , find the smallest  $m$  such that  $x_i \leq md_i$ . The quotients are  $\frac{1}{2}, \frac{3}{1}, 1$ , so take  $m = 3$  with  $\mathcal{B}_1, \mathcal{B}_2 \in \mathcal{B}_{\leq 2}^3$ .
2. Determine the 2-symmetric max-plus polynomials from the subcollection of orbits (3):

$$\begin{aligned} E_{3,(0,1),(0,0)}(x_1, d_1, x_2, d_2) &= d_1 \boxplus d_2 \\ &= \max\{d_1, d_2\} \\ E_{3,(0,1),(0,1)}(x_1, d_1, x_2, d_2) &= d_1 d_2 \\ &= d_1 + d_2 \\ E_{3,(0,0),(1,1)}(x_1, d_1, x_2, d_2) &= (x_2 \oplus d_2^3) d_2 \boxplus (x_1 \oplus d_1^3) d_1 \\ &= \max\{\min\{x_2, 3d_2\} + d_2, \min\{x_1, 3d_1\} + d_1\} \\ E_{3,(1,1),(0,1)}(x_1, d_1, x_2, d_2) &= (x_1 \oplus d_1^3) d_1 d_2 \boxplus (x_2 \oplus d_2^3) d_2 d_1 \\ &= \max\{\min\{x_1, 3d_1\} + d_1 + d_2, \min\{x_2, 3d_2\} + d_2 + d_1\} \\ E_{3,(1,1),(1,1)}(x_1, d_1, x_2, d_2) &= (x_1 \oplus d_1^3) d_1 (x_2 \oplus d_2^3) d_2 \\ &= \min\{x_1, 3d_1\} + d_1 + \min\{x_2, 3d_2\} + d_2 \end{aligned}$$

3. Evaluate on  $\mathcal{B}_1$ :

$$\begin{aligned} \max\{2, 1\} &= 2 \\ 2 + 1 &= 3 \\ \max\{\min\{1, 6\} + 2, \min\{3, 3\} + 1\} &= \max\{1 + 2, 3 + 1\} = 4 \\ \max\{\min\{1, 6\} + 2 + 1, \min\{3, 3\} + 2 + 1\} &= \max\{4, 6\} = 6 \\ \min\{1, 6\} + 2 + \min\{3, 3\} + 1 &= 7 \end{aligned}$$

4. Evaluate on  $\mathcal{B}_2$ :

$$\begin{aligned} \max\{2, 2\} &= 2 \\ 2 + 0 &= 2 \\ \max\{\min\{2, 6\} + 2\} &= 4 \\ \max\{\min\{2, 6\} + 2\} &= 4 \\ \min\{2, 6\} + 2 &= 4 \end{aligned}$$

The Euclidean-space vector representation of  $\mathcal{B}_1$  and  $\mathcal{B}_2$  is  $(2, 3, 4, 6, 7)$  and  $(2, 2, 4, 4, 4)$ , respectively.

### 3 Sufficient Statistics & Likelihoods for Barcodes

In this section, we will motivate the need for sufficient statistics for persistent homology. We also present our main result that tropical functions, defined according to the max-plus and min-plus semirings, are sufficient statistics for persistence barcodes. Lastly, we discuss the implications of our result on statistical parametric inference by imposing the structure of exponential family distributions and defining likelihood functions for the tropicalized representation of barcodes.

#### 3.1 Sufficient Statistics & Persistent Homology

Although a well-defined metric space, the geometry of the space of barcodes is known to be a complex Alexandrov space with curvature bounded from below (Turner et al., 2014). In the more conventional sense, this structure is theoretically not restrictive since it is also a Polish space and admits well-defined characterizations of expectations, variances, conditional probabilities. This means that formal probability measures and distributions may be formulated for persistence diagrams and, equivalently, barcodes (Mileyko et al., 2011). Probability measures on the space of persistence diagrams have been of key interest for conducting statistical inference in topological data analysis (e.g. Adler and Taylor, 2011; Carlsson, 2014; Blumberg et al., 2014). However, formulating explicitly parametric probability distributions directly on the space of persistence diagrams remains a significant challenge. The existence of Fréchet means is proven, but these are not unique. It is therefore difficult to perform statistical analyses, particularly parametric inference, directly on the space of persistence diagrams due to its complicated geometry (Crawford et al., 2017).

**Statistical Sufficiency.** The difficulty of defining explicit probability distributions directly on the space of barcodes is our main motivation for the present work. Our strategy to solve this problem is to achieve sufficiency for barcodes via coordinatization into Euclidean space. Once in Euclidean space, we can establish distributional forms for the coordinatized equivalents of barcodes, so long as our mapping is a sufficient statistic. Since the coordinatization is a topological embedding, and the mappings are continuous and injective between barcode and Euclidean space, any parametric inference performed on these coordinatized equivalents naturally preserves all key properties in barcode space. The property of injectivity ensures the existence of an inversion. Therefore, the results of parametric analyses on Euclidean space can theoretically be translated back into barcode space by the inverse mappings. However, note that this is not necessary because sufficiency guarantees that, for analytical purposes, all of the information of the original barcodes is retained by their coordinates.

Intuitively speaking, a statistic is said to be *sufficient* for a family of probability distributions if the original sample from which the statistic was calculated provides no more additional information on the underlying probability distribution, than does the statistic itself. This is true whether the statistic in question is a function of the data mapping to a codomain, or whether the statistic is an evaluation of such a function. In the former case, sufficiency refers to a property of the statistic with respect to probability measures on the corresponding domain and codomain, when the exact distribution is unknown but the spaces nevertheless admit well-defined probability measures. In the latter case, sufficiency refers to the purpose of determining the exact probability distribution that generated the observed data (a key aim of parametric inference). In this case, knowing a sufficient statistic contains all of the information needed to compute any estimator of the parameters for the assumed underlying distribution.

**Definition 7.** Let  $X$  be a vector of observations of size  $n$  whose components  $X_i$  are independent and identically distributed (*i.i.d.*). Let  $\vartheta$  be the parameter that characterizes the underlying probability distribution that generates  $X_i$ . A statistic  $T(X)$  is *sufficient* for the parameter  $\vartheta$  if the conditional probability of observing  $X$ , given  $T(X)$ , is independent of  $\vartheta$ . Equivalently,

$$\mathbb{P}(X = x | T(X) = t, \vartheta) = \mathbb{P}(X = x | T(X) = t).$$

In the measure-theoretic sense, a statistic  $T(\cdot)$  is sufficient for a given set  $\mathcal{M}$  of probability measures on a  $\sigma$ -algebra  $\mathbf{S}$ , if the conditional probability of some subset  $E \subseteq \mathbf{S}$ , given a value  $y$  of  $T(\cdot)$ , is invariant for every probability measure in  $\mu \in \mathcal{M}$ .

We now recall an important result in the mathematical statistics literature for assessing statistical sufficiency: the factorization criterion for sufficient statistics (Fisher, 1922; Neyman, 1935). The following theorem is classically used as a definition for checking sufficiency.

**Theorem 1** (Fisher–Neyman Factorization Criterion). *If the probability density function for the observed data is  $f(x; \vartheta)$ , then the statistic  $T = T(x)$  is sufficient for  $\vartheta \in \Theta$  if and only if nonnegative functions  $g$  and  $h$  exist such that*

$$f(x; \vartheta) = h(x)g(T(x); \vartheta). \quad (4)$$

*Namely, the density  $f$  can be factored into a product where one factor  $h$  is independent of  $\vartheta$ , and the other factor  $g$  depends on  $\vartheta$  and only depends on  $x$  only through  $T$ .*

This factorization need not be unique for any given distribution — meaning multiple factorizations may result in the same distribution. To see that the identity function  $T(X) = X$  is a sufficient statistic, we may factorize by setting  $h(x) = 1$  and  $g(T(x); \vartheta) = f(x; \vartheta)$  for all  $x$ . This is the trivial sufficient statistic, and is interpreted as the data themselves being sufficient for the data.

This theorem is valid for arbitrary families of distributions. The following measure-theoretic reinterpretation asserts the generality of the theorem, particularly in the case where sufficiency refers to statistics on probability measures (Halmos and Savage, 1949). Essentially, the result depends on the assumption that the family of probability measures  $\mathcal{P}_\vartheta$  is dominated by a single  $\sigma$ -finite measure. In the case of discrete variables, this amounts to the existence of a countable set  $S$ , which does not depend on  $\vartheta$ , for which  $\mathcal{P}_\vartheta(S) = 1$  for all  $\vartheta \in \Theta$ . In the case of absolutely continuous variables, this amounts to the measurability of the statistic.

**Theorem 2.** *A necessary and sufficient condition that the statistic  $T(\cdot)$  be sufficient for a dominated set  $\mathcal{M}$  of measures on a  $\sigma$ -algebra  $\mathbf{S}$  is that for every  $\mu \in \mathcal{M}$ , the density  $f_\mu := d\mu/d\lambda$  (where  $d\mu/d\lambda$  is the Radon–Nikodym derivative) admits the factorization*

$$f_\mu(x) = h(x)g_\mu(T(x)),$$

*where  $h(\cdot) \neq 0$  is a measurable function with respect to  $\mathbf{S}$ , and  $g_\mu(\cdot)$  is a measurable function with respect to the codomain  $\mathcal{T}$  of  $T(\cdot)$ .*

As will be shown subsequently, in our case where the statistic is given by tropical polynomials (2),  $\mathbf{S}$  corresponds to the space of persistence barcodes  $\mathcal{B}_{\leq n}^m$ , while  $\mathcal{T}$  is Euclidean space. Although the above theorem holds in broader generality, the case of absolutely continuous variables is sufficient for our setting since the tropical polynomials are Lipschitz continuous, and are therefore absolutely continuous as well.

**Sufficient Statistics for Persistence Barcodes.** Our main theorem ascertains that the map from  $\mathcal{B}_{\leq n}^m$ , as an embedding into Euclidean space via tropical functions, is a sufficient statistic. Sufficiency for statistics is classically asserted by the injectivity of continuous mappings (e.g. Halmos and Savage, 1949; Diaconis, 1992). We follow this same approach in our proof.

**Theorem 3.** Consider the subset  $\mathcal{B}_{\leq n}^m$  of  $\mathcal{B}_{\leq n}$  consisting of barcodes  $(x_1, d_1, x_2, d_2, \dots, x_n, d_n)$  with  $d_i > 0$  such that  $x_i \leq md_i$ , for some fixed  $m$  and  $i = 1, 2, \dots, n$ . Let  $\mathcal{P}$  be the family of probability measures on  $\mathcal{B}_{\leq n}^m$ . The map defined by

$$T : \mathcal{B}_{\leq n}^m \rightarrow \mathbf{R}^d$$

$$\mathcal{B} \mapsto (E_{m,(1,1)^i,(0,1)^j}(x_1, d_1, x_2, d_2, \dots, x_n, d_n))_{i+j \in \mathbf{N}_{\leq n}}(\mathcal{B}), \quad (5)$$

where  $d$  is the number of orbits used to define separating functions, is a sufficient statistic for all  $P \in \mathcal{P}$ .

*Proof.* Denote the image of  $\mathcal{B}_{\leq n}^m$  under  $(E_{m,(1,1)^i,(0,1)^j}(x_1, d_1, x_2, d_2, \dots, x_n, d_n))_{i+j \in \mathbf{N}_{\leq n}}$  by  $\mathcal{T} \subseteq \mathbf{R}^d$ . Both  $\mathcal{B}_{\leq n}^m$  and  $\mathcal{T}$  are metric spaces, where a natural metric on  $\mathcal{B}_{\leq n}^m$  is the Wasserstein  $p$ -distance, and  $\mathcal{T}$  is the standard Euclidean distance. Both can be viewed as measurable spaces by taking Borel  $\sigma$ -algebras. Both metric spaces are separable, so the Borel  $\sigma$ -algebras equal the  $\sigma$ -algebras generated by the open balls of respective spaces.

To apply Theorem 2 and show that  $T$  generates sufficient statistics for barcodes on  $\mathcal{B}_{\leq n}^m$ , we use the property that the map  $T$  is injective, proven in Proposition 1. By definition, this means that there exists some  $\varphi : \mathcal{T} \rightarrow \mathcal{B}_{\leq n}^m$  such that  $\varphi \circ T = \text{Id}_{\mathcal{B}_{\leq n}^m}$  and  $T \circ \varphi = \text{Id}_{\mathcal{T}}$ . For notational simplicity, we denote  $f_{\vartheta}(\mathcal{B}) := f(\mathcal{B}; \vartheta)$ , and  $g_{\vartheta}(T(\mathcal{B})) := g(T(\mathcal{B}); \vartheta)$ . The condition of Theorem 2 holds for  $T$  where  $g_{\vartheta} = f_{\vartheta} \circ \varphi$ :

$$\begin{aligned} f_{\vartheta}(\mathcal{B}) &= f_{\vartheta}(\varphi(T(\mathcal{B}))), \\ &= f_{\vartheta} \circ \varphi(T(\mathcal{B})), \\ &= g_{\vartheta}(T(\mathcal{B})). \end{aligned}$$

We now proceed to verify that all the relevant functions are measurable. The function  $h(\mathcal{B}) \equiv 1$  is a constant, and so is automatically measurable. To show that  $f_{\vartheta}$  is measurable, notice that  $g_{\vartheta} = f_{\vartheta} \circ \varphi$ . Since  $g_{\vartheta}$  is measurable by assumption, it will be sufficient to show  $\varphi : \mathcal{T} \rightarrow \mathcal{B}_{\leq n}^m$  is also measurable. Since the map  $T$  is Lipschitz continuous with respect to the the Wasserstein  $p$ - and standard Euclidean distances, it is measurable with respect to Borel  $\sigma$ -algebras. An inverse of an injective Borel continuous map is Borel continuous, which leads to the conclusion that  $\varphi$  is Borel measurable.  $\square$

As previously mentioned, the dimension  $d$  of the codomain  $\mathbf{R}^d$  of the tropical mappings (5) need not equal the number of all orbits given by the symmetric group action, as a result of Proposition 1. The dimension  $d$  is determined by the restriction that  $i + j \in \mathbf{N}_{\leq n}$ , where  $n$  is the number of bars in the largest barcode, which simplifies to  $2d = 2n + n(n + 1)$ .

### 3.2 Exponential Families & Likelihood Functions for Persistent Homology

Sufficient statistics are strictly parametric by nature, and thus are most relevant in studies centered around inference. Parametric inference is carried out under the assumption that the parameter space is necessarily finite and therefore bounded with respect to sample size. The only class of distributions that satisfies this restriction is known as the *exponential family* (Darmois, 1935; Koopman, 1936; Pitman, 1936), generated by a particular case of Theorem 1.

**Definition 8.** If the following restriction is applied for equation (4) in Theorem 1,

$$g(T(x); \vartheta) = \exp \{ \langle \eta(\vartheta), T(x) \rangle - A(\vartheta) \}, \quad (6)$$

where  $\langle \cdot, \cdot \rangle$  denotes the usual inner product in Euclidean space, and  $\eta(\vartheta)$ ,  $T(x)$ , and  $A(\vartheta)$  are known functions, then the resulting class of probability distributions spanned by (4), with  $g$  given by (6), is known as the *exponential family* of probability densities.

Though less general than the form in (4), the exponential family is nevertheless an important class of distributions that is comprised of many of the most common distributions in statistics (e.g. Gaussian, Poisson,  $\chi^2$ , Bernoulli, exponential, and gamma distributions). Given an observation  $x$ , the *likelihood function*  $\mathcal{L}$  is a function of the parameter  $\vartheta$  from the probability distribution  $f_\vartheta$  that generates the observation within the sample. Namely,

$$\begin{aligned}\mathcal{L}(\vartheta | x) &= f_\vartheta(x) \\ &= h(x)g(T(x); \vartheta),\end{aligned}\tag{7}$$

for any given sufficient statistic. For an exponential family model, where  $g$  is given by (6), if we denote  $a(\vartheta) := e^{-A(\vartheta)}$ , then we obtain

$$\mathcal{L}(\vartheta | x) = h(x)a(\vartheta) \exp \{ \langle \eta(\vartheta), T(x) \rangle \}.\tag{8}$$

**Likelihoods for Persistence Barcodes.** For  $K$  *i.i.d.* barcodes distributed according to  $f_\vartheta$  as defined by (7), the joint likelihood for the  $K$  observations may be written as

$$\mathcal{L}(\vartheta | \mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_K) = \prod_{k=1}^K f_\vartheta(T(\mathcal{B}_k)).\tag{9}$$

When  $f_\vartheta$  belongs to the exponential family, as denoted in (8), the joint likelihood for the  $K$  observations may be written as

$$\begin{aligned}\mathcal{L}(\vartheta | \mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_K) &= a(\vartheta)^K \prod_{k=1}^K \exp \{ \langle \eta(\vartheta), T(\mathcal{B}_k) \rangle \} \\ &= a(\vartheta)^K \exp \left\{ \sum_{i=1}^K \langle \eta(\vartheta), T(\mathcal{B}_k) \rangle \right\}.\end{aligned}\tag{10}$$

Since the form of our tropical sufficient statistic (5) is a vector in Euclidean space for each barcode, this quantity is well-defined. Most importantly, the inner product is computationally tractable.

## 4 Application: Studying Evolutionary Behavior in Viral Datasets

We now demonstrate the utility of our sufficiency result in applications regarding two infectious viral diseases: the human immunodeficiency virus (HIV), and avian influenza. More specifically, we study both dimensions 0 and 1 of persistent homology. In dimension 0, we concretely demonstrate sufficiency in the preservation of biological characteristics of HIV. In dimension 1, we impose the parametric structure of a particular exponential distributive family to carry out pairwise comparisons between the marginal distributions of intra- and inter-subtype reassortment in avian influenza.

### 4.1 Demonstrating Sufficiency in HIV Transmission Clustering

HIV transmission clusters represent groups of epidemiologically-related individuals who share a common recent viral ancestor (Hué et al., 2005; Hassan et al., 2017). Recently, Patiño-Galindo et al. (2017) reported the existence of a prominent HIV transmission cluster affecting more than 100 people from an MSM population in Valencia, Spain.

Since epidemiologically-related groups are often characterized by the identification of small pairwise genetic distances, sequences derived from a transmission cluster can be easily differentiated from those derived from unrelated patients (see Figure 1). Similarly, dimension 0 in persistent homology ( $\text{PH}_0$ ) may also be used to capture the clustering of sequences based on vertical evolution (Chan et al., 2013). This is shown via bottleneck distances in Figure 2(a). Note that considering  $\text{PH}_0$  persistence poses an advantage over simply considering genetic distance alone. This is because  $\text{PH}_0$  persistence removes recombination

events which are known to affect phylogenetic inference — a practice which is often seen as desirable prior to phylogenetic analysis (Posada and Crandall, 2002).

In order to demonstrate sufficiency, and verify that the tropical representations of the barcodes retain the same biologically relevant information in dimension 0, we generated 10 random subsets of 30 polymerase sequences from (i) the MSM transmission cluster, (ii) the reference dataset (i.e. the unrelated patients), and (iii) sequences from both the transmission cluster and the reference dataset. Matrices of pairwise genetic distances were obtained from each subset based on evolutionary models. We then obtained their dimension 0 barcodes and evaluated them via the tropical functions described in Section 2.2. The resulting barcodes (and their Euclidean representations after evaluating them via tropical functions) from subsets (i) and (ii) provide information on intra-group distances, while those from subset (iii) provide information on inter-group distances (i.e. transmission cluster versus unrelated patients). Figures 1 and 2 show that the three subsets can be easily differentiated, thus demonstrating that the tropicalization of the barcodes are sufficient statistics for epidemiologically-related individuals both under, as well as in the absence of, recombination events.

## 4.2 Analyzing Intra- and Intersubtype Reassortment in Avian Influenza

The influenza virus presents a genome that consists of eight different RNA molecules (segments). Consequently, genetic reassortment — that is, the process of swapping gene segments — is an important evolutionary process that contributes to the virus’s diversity. Subtype classification of influenza is based on the analysis of two surface proteins: hemagglutinin (HA) and neuraminidase (NA). There exists 18 types of HA (indexed from 1 to 18) and 11 types of NA (indexed from 1 to 11). Gene reassortment can occur either between viruses from the same subtype (i.e. intrasubtype reassortment, Maljkovic Berry et al., 2016), or between viruses from different subtypes (i.e. intersubtype reassortment, Nguyen et al., 2016). From a broad perspective, detecting gene reassortment is key to understanding mutations within the evolutionary dynamics of viruses.

Traditionally, the detection of reassortment events is based on molecular phylogenetic analyses, in which different genomic regions (i.e. segments) are analyzed independently. Specific details on this procedure given in Appendix A2. If no reassortment has occurred, phylogenetic trees obtained for each independent segment show a similar clustering pattern of the viral samples that are analyzed. This indicates that all segments share the same evolutionary history. Alternatively, when a reassortment event occurs, it is detected by a segment generating a phylogenetic tree with a differing clustering pattern. This indicates the existence of segment swapping (Pérez-Losada et al., 2015). The detection and quantification of reassortment events with phylogenetic approaches is a slow and tedious process, partly due to the fact that obtaining accurate phylogenetic trees can take days, or even weeks, when hundreds or thousands of samples are analyzed (Maljkovic Berry et al., 2016). The use of persistent homology in detecting genetic reassortment events poses a clear advantage over phylogenetic analysis, not only due to the increased computational speed, but also because of the enhanced interpretability. Dimension 1 persistence intervals provide explicit information on the genetic divergence of the sequences involved in the reassortment event (Chan et al., 2013; Emmett et al., 2014).

Given that viruses from the same subtype are more closely related to each other, than they are to viruses from different subtypes, it is expected that lengths of intervals in  $\text{PH}_1$  barcodes derived from events of intrasubtype reassortment will be shorter than those from intersubtype reassortment (Emmett et al., 2014). We will demonstrate this parametrically using our sufficiency result. In this context, sufficiency means preserving the information that intra- and intersubtype samples should not cluster together.

Since we are interested in analyzing reassortment events in this example, we focus on avian influenza type A and study dimension 1 persistence. We generated 100 random subsets of 56 concatenated HA and NA sequences representing: (i) only H5N1 sequences, and (ii) sequences from all subtypes H5N $u$  and H $v$ N1 where  $u = 1, 2, \dots, 12$  and  $v = 1, 2, \dots, 9$ , respectively. The resulting barcodes from subsets (i) provide information on intrasubtype reassortment events, while those from subsets (ii) provided information on intersubtype reassortment events. Because intrasubtype reassortment events are significantly scarcer than intersubtype reassortment events, the intersubtype reassortment events resulted in higher-dimensional (or longer) vectors (i.e.  $p = 100$  for intersubtypes and  $p = 10$  for intrasubtypes).

Since we have sufficient statistics for the family of probability measures on the space of barcodes, we can

assume a parametric distribution. In particular, we assume that the Euclidean representation of barcodes, from each of the intra- and intersubtype populations, jointly come from a multivariate normal distribution as in (10):

$$T(\mathcal{B}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}),$$

with mean vector  $\boldsymbol{\mu}$  and covariance matrix  $\boldsymbol{\Sigma}$ . This implies that marginally, we may assume that each barcode is normally distributed, as denoted in (8):

$$T(\mathcal{B}_k) \sim \mathcal{N}(\mu_k, \sigma_k^2), \quad k = 1, 2, \dots, K, \quad (11)$$

with scalar mean and variance  $\vartheta = \{\mu_k, \sigma_k^2\}$ , respectively. This is empirically seen by the distribution plots of intra- and intersubtype reassortment events from the generated data shown in Figure 3.

In this parametric setting, we may study pairwise comparisons between the marginal distributions of each  $T(\mathcal{B}_k)$  in order to assess the following clustering behavior: samples that cluster together will tend to have marginal distributions that are more similar or, in other words, have shorter distances. A function commonly used to measure differences between two probability distributions is an  $f$ -divergence. A particular instance of an  $f$ -divergence is the Hellinger distance (Hellinger, 1909), which we now briefly define.

**Definition 9.** Assume that  $T(\mathcal{B}_i)$  and  $T(\mathcal{B}_j)$  are probability measures that are absolutely continuous with respect to a third probability measure,  $\lambda$ . The *Hellinger distance* is then computed by solving the following

$$H^2(T(\mathcal{B}_i), T(\mathcal{B}_j)) = \frac{1}{2} \int \left( \sqrt{\frac{dT(\mathcal{B}_i)}{d\lambda}} - \sqrt{\frac{dT(\mathcal{B}_j)}{d\lambda}} \right)^2 d\lambda, \quad (12)$$

where  $dT(\mathcal{B}_i)/d\lambda$  and  $dT(\mathcal{B}_j)/d\lambda$  are the Radon–Nikodym derivatives of  $T(\mathcal{B}_i)$  and  $T(\mathcal{B}_j)$ , respectively.

The Hellinger distance is a well-defined metric, and thus satisfies the triangle inequality. Derivable from the Cauchy–Schwarz inequality, the Hellinger distance also satisfies the following bound for all probability distributions:

$$0 \leq H(T(\mathcal{B}_i), T(\mathcal{B}_j)) \leq 1.$$

It can be shown that when the probability measures  $T(\mathcal{B}_i)$  and  $T(\mathcal{B}_j)$  stem from an the exponential family, (12) has a closed form (e.g. Van der Vaart, 2000). More specifically, for two normal random variables  $T(\mathcal{B}_i) \sim \mathcal{N}(\mu_i, \sigma_i^2)$  and  $T(\mathcal{B}_j) \sim \mathcal{N}(\mu_j, \sigma_j^2)$ , we have the following:

$$H^2(T(\mathcal{B}_i), T(\mathcal{B}_j)) = 1 - \sqrt{\frac{2\sigma_i\sigma_j}{\sigma_i^2 + \sigma_j^2}} \exp \left\{ -\frac{(\mu_i - \mu_j)^2}{4(\sigma_i^2 + \sigma_j^2)} \right\}. \quad (13)$$

Based on the marginal normality assumption in (11) with parameter values  $\mu_i, \mu_j$  and  $\sigma_i, \sigma_j$  estimated empirically from the data, we show that the tropicalization of persistence barcodes sufficiently preserves the fact that intra- and intersubtype samples do not cluster together. This was done by calculating (13) for every pair  $T(\mathcal{B}_i)$  and  $T(\mathcal{B}_j)$  and encoding this information as a similarity matrix  $\mathbf{H}$ . Figure 4 plots a scaled version of the matrix  $\mathbf{H}^* = \mathbf{1}\mathbf{1}^\top - \mathbf{H}$ , where  $\mathbf{1}$  is a vector of ones. These results can be interpreted as follows: a value of 1 represents exact likeness, while  $-1$  represents complete dissimilarity. As expected, the intersubtypes are seen to be more homogeneous and demonstrate a higher level of similarity, while the intrasubtypes are comparatively much more random. Moreover, there is very little overlap between the two groups.

## 5 Discussion

In this paper, we explored a functional vectorization method for persistence barcodes that is a topological embedding into Euclidean space based on tropical algebraic geometry. We proved that it generates sufficient statistics for the family of all probability measures on the space of barcodes, under mild regularity conditions. The statistical and data analytic utility of this result lies in the fact that sufficiency allows parametric probabilistic assumptions to be imposed, which previously has been difficult due to the prohibitive geometry

of the space of barcodes. This allows for the application of classical parametric inference methodologies to persistence barcodes, which we demonstrated in a concrete example in dimension 1 persistence applied to avian influenza data.

Our sufficiency also result provides the foundations for future research towards a better understanding of parametric probabilistic behavior on the space of barcodes. Since the tropical vectorization method was shown to be an injective function, we may take a class of parametric probability distributions (e.g. the exponential family) and calculate its image via the inverse of the tropical functions in order to explore what functional form the exponential family assumes on the space of barcodes. Such a study would be algorithmic in nature, since it would entail exploring collections of maps given by subsets of the orbits under the symmetric group action.

In terms of an increased utility in statistical analyses, the tropical functions are limited since the vector representation is unique up to barcodes only, and not on the level of bars. This is restrictive in questions motivated by the application of our paper. For example, when predicting or modeling genetic outbreaks of infectious diseases, individual sequences (or at least the persistence generators) would need to be identifiable. Future research towards these efforts may begin with the development of a vectorization method that traces back to the individual bars of a barcode.

## Software and Data Availability

Software to compute and evaluate the tropical functions on barcodes is publicly available in C++ code, co-authored by Melissa McGuirl (Brown University) and Steve Oudot (INRIA Saclay), and located on the Tropix GitHub repository at <https://github.com/lorinanthony/Tropix>. The persistence barcodes were calculated using Ripser (Bauer, 2017), which is written in C++ and is freely available at <https://github.com/Ripser/ripser>. The bottleneck distances were computed using Hera (Kerber et al., 2017), which is also written in C++ and freely available at [https://bitbucket.org/grey\\_narn/hera](https://bitbucket.org/grey_narn/hera). The data used in this paper were obtained from publicly available sources and preprocessed, as detailed in Appendix A3. The final version used in the analyses in this paper are also publicly available on the Tropix GitHub repository.

## Acknowledgements

The authors are very grateful to Karen Gomez-Inguanzo (Columbia University) and Melissa McGuirl (Brown University) for help with formulation of the code. We also wish to thank Robert Adler (Technion), Ulrich Bauer (Technische Universität München), Omer Bobrowski (Technion), Joseph Minhow Chan (Memorial Sloan Kettering Cancer Center), Paweł Dłotko (Swansea University), Hossein Khiabani (Rutgers University), Albert Lee (Columbia University), Sayan Mukherjee (Duke University), and Raúl Rabadán (Columbia University) for helpful discussions. We are especially indebted to Steve Oudot (INRIA Saclay) for his extensive input and support in code formulation and content throughout the course of this project.

A.M. and J.A.P.-G. are supported by the National Institutes of General Medical Sciences (NIGMS) of the National Institutes of Health (NIH) under award number R01GM117591. L.C. would like to acknowledge the support of start up funds from Brown University. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of any of the funders. The authors would also like to acknowledge GenBank, the HIV Sequence Database operated by Los Alamos National Security, LLC, and the National Center for Biotechnology Information (NCBI) Influenza Virus Database initiatives for making the data in this study publicly available.

## Appendix

### A1 Mathematical Supplement

**Claim A1.** *The polynomials that are a ring of algebraic functions on the space of barcodes, identified by Adcock et al. (2016), are not Lipschitz with respect to the Wasserstein  $p$ - and bottleneck distances.*



*Proof.* Consider the following counterexample. Take one barcode to be the diagonal  $\Delta$  (i.e. intervals of length 0) denoted by  $\mathcal{B}_0$ , and consider the family of barcodes with a single interval  $\mathcal{B}_x = \{[x, x + 1]\}$ . The Wasserstein  $p$ - and bottleneck distances between these  $\mathcal{B}_0$  and  $\mathcal{B}_x$  is the same for all  $x$ , namely  $1/2$ . Consider a polynomial given by Adcock et al. (2016), e.g.

$$p_{2,1}(x_1, y_1, x_2, y_2, \dots) = \sum_i (x_i + y_i)^2 (y_i - x_i)^1,$$

then

$$p_{2,1}(\mathcal{B}_0) = 0 \quad \text{and} \quad p_{2,1}(\mathcal{B}_x) = (2x + 1)^2.$$

Now, the difference  $p_{2,1}(\mathcal{B}_x) - p_{2,1}(\mathcal{B}_0)$  tends to infinity as  $x$  tends to infinity. The bottleneck (and Wasserstein  $p$ -) distance, meanwhile, is constant at value  $1/2$  the entire time. Therefore, by definition, this function is not Lipschitz. This particular case is problematic and shows that other functions proposed in Adcock et al. (2016) are also not Lipschitz with respect to the Wasserstein  $p$ - and bottleneck distances.  $\square$

## A2 Molecular Phylogenetic Analysis

A phylogenetic analysis consists of studying the evolutionary relationships of a set of individuals (i.e. viral sequences), with the aim of extracting and analyzing their diversification history. These type of analyses are usually performed by using an alignment of sequences as input. Given this input, a *phylogenetic tree* is then produced — that is, a graphical representation of the evolutionary relationships among the individuals. The most widely used methods of phylogenetic inference are maximum likelihood and Bayesian phylogenetics, which aim to maximize a statistical parameter (either the likelihood of the alignment or its posterior probability) by modeling the molecular evolutionary processes considered in the phylogenetic reconstruction using stochastic processes (e.g Guindon et al., 2010; Ronquist et al., 2012).

## A3 Data Sourcing & Preprocessing

The respective virus datasets in this paper were obtained from public sources. We give the details on their sources and preliminary data processing procedures below.

**HIV.** HIV polymerase sequences derived from patients included in the MSM HIV transmission cluster (Patiño-Galindo et al., 2017; Patiño-Galindo et al., 2017) were retrieved from supplementary material made public on GenBank. Sequences derived from unrelated patients were obtained from the Los Alamos HIV database in October 2016. Only sequences from the same subtype (HIV subtype B) spanning the polymerase region were considered. In order to ensure that these sequences were not epidemiologically related, redundant sequences were removed after conducting an initial clustering analysis with a specified genetic distance threshold of 5%, using CD-HIT (Huang et al., 2010). All sequences were aligned using MAFFTv7 (Katoh and Standley, 2013).

**Avian Influenza.** Hemagglutinin (HA) and Neuraminidase (NA) genes of avian influenza A were downloaded from the Influenza Virus Database of the National Center for Biotechnology Information (NCBI). The resulting gene datasets were aligned with MAFFTv7 (Katoh and Standley, 2013). Concatenated sequences of both genes (derived from the same sample) were generated with the package *ape* written in R (Paradis et al., 2004). The multiple sequence alignments were trimmed with TrimAl (Capella-Gutierrez et al., 2009) in order to remove regions of sparse homology (i.e. biologically shared ancestry).

In both viral examples, pairwise distances were obtained using PAUP (Swofford, 2001) and were calculated by using the GTR + GAMMA (4 CAT) model, which is commonly used for studying HIV and Influenza data (Tian et al., 2015; Worobey et al., 2016). Briefly, the GTR model is a time reversible model that considers variable base frequencies, where each pair of nucleotide substitutions occur at different rates (Donnelly and Tavaré, 1995). Combined with a gamma distribution, it also accounts for rate variation among sites (Yang, 1995). The use of a substitution model when calculating genetic distances, as carried out according to these procedures, leads to estimates that are assumed to be more biologically accurate.

## Figures




Figure 1: **PCA plot of genetic clustering for HIV sequences.** Principal components calculated from pairwise genetic distances for HIV transmission cluster versus unrelated reference cluster.




Figure 2: **Projection plots for pairwise bottleneck distances and Euclidean distances for HIV sequences.** (a) Metric MDS calculated from pairwise bottleneck distances calculated from dimension 0 persistence barcodes. (b) PCA calculated from pairwise Euclidean distances of tropicalized barcodes.




Figure 3: **Marginal distribution plot of intra- and intersubtype reassortment events for avian influenza.** The marginal distributions for square-root transformations of the Euclidean barcode representations (via tropicalization) were calculated for both intra- and intersubtype reassortment for avian influenza, and then fitted with a smooth density function.




Figure 4: **Pairwise comparisons of scaled Hellinger distances between Euclidean barcode representations.** The entries are color-coded according to values of the matrix  $\mathbf{1}\mathbf{1}^\top - \mathbf{H}$ , where  $\mathbf{1}$  is a vector of ones. A value of 1 represents exact likeness, while  $-1$  represents complete dissimilarity. Blue values represent those that are more similar, while red values are more dissimilar. This plot shows that intersubtypes are more homogeneous (a higher level of similarity, i.e. more blue), while the intrasubtypes are comparatively much more random. There is very little overlap between the two groups.

## References

- Adams, H. and G. Carlsson (2014). Evasion paths in mobile sensor networks. *The International Journal of Robotics Research* 34, 90–104.
- Adams, H., T. Emerson, M. Kirby, R. Neville, C. Peterson, P. Shipman, S. Chepushtanova, E. Hanson, F. Motta, and L. Ziegelmeier (2017). Persistence images: A stable vector representation of persistent homology. *J. Mach. Learn. Res.* 18(1), 218–252.
- Adcock, A., E. Carlsson, and G. Carlsson (2016). The ring of algebraic functions on persistence bar codes. *Homology, Homotopy and Applications* 18, 381–402.
- Adcock, A., D. Rubin, and G. Carlsson (2014). Classification of hepatic lesions using the matching metric. *Computer Vision and Image Understanding* 121, 36 – 42.
- Adler, R. and J. Taylor (2011). *Topological Complexity of Smooth Random Functions: École d’Été de Probabilités de Saint-Flour XXXIX-2009*. Lecture Notes in Mathematics. Springer Berlin Heidelberg.
- Adler, R. J., S. Agami, and P. Pranav (2017). Modeling and replicating statistical topology, and evidence for cmb non-homogeneity.
- Bartlett, M. S. (1937). Properties of sufficiency and statistical tests. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* 160(901), 268–282.
- Bauer, U. (2015-2017). © Ripser [github.com/Ripser/ripser](https://github.com/Ripser/ripser).
- Bendich, P., S. P. Chin, J. Clark, J. Desena, J. Harer, E. Munch, A. Newman, D. Porter, D. Rouse, N. Strawn, and A. Watkins (2016). Topological and statistical behavior classifiers for tracking applications. *IEEE Transactions on Aerospace and Electronic Systems* 52(6), 2644–2661.
- Blumberg, A. J., I. Gal, M. A. Mandell, and M. Pancia (2014). Robust statistics, hypothesis testing, and confidence intervals for persistent homology on metric measure spaces. *Foundations of Computational Mathematics* 14(4), 745–789.
- Bubenik, P. (2015). Statistical topological data analysis using persistence landscapes. *Journal of Machine Learning Research* 16(1), 77–102.
- Capella-Gutierrez, S., J. M. Silla-Martinez, and T. Gabaldon (2009). trimAl: a tool for automated alignment trimming in large-scale phylogenetic analyses. *Bioinformatics* 25(15), 1972–1973.
- Carlsson, G. (2014). Topological pattern recognition for point cloud data. *Acta Numerica* 23, 289–368.
- Carlsson, G. and A. J. Zomorodian (2005). Computing persistent homology. *Discrete and Computational Geometry* 33, 249–274.
- Carrière, M., S. Y. Oudot, and M. Ovsjanikov (2015). Stable topological signatures for points on 3d shapes. *Eurographics Symposium on Geometry Processing 2015* 34, 77–102.
- Chan, J. M., G. Carlsson, and R. Rabadán (2013). Topology of viral evolution. *Proceedings of the National Academy of Sciences* 110(46), 18566–18571.
- Chazal, F., V. De Silva, M. Glisse, and S. Oudot (2016). *The structure and stability of persistence modules*. Springer.
- Chung, M. K., P. Bubenik, and P. T. Kim (2009). Persistence diagrams of cortical surface data. In *International Conference on Information Processing in Medical Imaging*, pp. 386–397. Springer.
- Cohen-Steiner, D., H. Edelsbrunner, and J. Harer (2007). Stability of persistence diagrams. *Discrete & Computational Geometry* 37(1), 103–120.

- Crawford, L., A. Monod, A. X. Chen, S. Mukherjee, and R. Rabadán (2017). Functional Data Analysis using a Topological Summary Statistic: the Smooth Euler Characteristic Transform. arXiv.1611.06818.
- Curto, C., V. Itskov, A. Veliz-Cuba, and N. Youngs (2013). The neural ring: An algebraic tool for analyzing the intrinsic structure of neural codes. *Bulletin of Mathematical Biology* 75(9), 1571–1611.
- Darmois, G. (1935). Sur les lois de probabilités à estimation exhaustive. *C.R. Acad. Sci. Paris (in French)* 200, 1265–1266.
- Di Fabio, B. and M. Ferri (2015). Comparing persistence diagrams through complex vectors. In S. I. Publishing (Ed.), *International Conference on Image Analysis and Processing*, pp. 294–305. Springer.
- Diaconis, P. (1992). Sufficiency as statistical symmetry. *Proc. 100th Anniversary American Mathematical Society*, 15–26.
- Donnelly, P. and S. Tavaré (1995). Coalescents and genealogical structure under neutrality. *Annual Review of Genetics* 29(1), 401–421. PMID: 8825481.
- Edelsbrunner, H., D. Letscher, and A. J. Zomorodian (2002). Topological persistence and simplification. *Discrete and Computational Geometry* 28, 511–533.
- Emmett, K., D. Rosenbloom, P. Camara, and R. Rabadan (2014). Parametric inference using persistence diagrams: A case study in population genetics.
- Ferri, M. and C. Landi (1999). Representing size functions by complex polynomials. *Proc. Math. Met. in Pattern Recognition* 9, 16–19.
- Ferri, M. and I. Stanganelli (2010). Size functions for the morphological analysis of melanocytic lesions. *International Journal of Biomedical Imaging* 2010, 1–5.
- Fisher, R. A. (1922). On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 222(594-604), 309–368.
- Frosini, P. and C. Landi (1999). Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis* 9(4), 596–603.
- Ghrist, R. and V. de Silva (2006). Coordinate-free coverage in sensor networks with controlled boundaries via homology. *International Journal of Robotics Research* 25, 1205–1222.
- Giusti, C., E. Pastalkova, C. Curto, and V. Itskov (2015). Clique topology reveals intrinsic geometric structure in neural correlations. *Proceedings of the National Academy of Sciences* 112(44), 13455–13460.
- Guindon, S., J.-F. Dufayard, V. Lefort, M. Anisimova, W. Hordijk, and O. Gascuel (2010). New algorithms and methods to estimate maximum-likelihood phylogenies: Assessing the performance of phylml 3.0. *Systematic Biology* 59(3), 307–321.
- Halmos, P. R. and L. J. Savage (1949). Application of the radon-nikodym theorem to the theory of sufficient statistics. *Ann. Math. Statist.* 20(2), 225–241.
- Hassan, A. S., O. G. Pybus, E. J. Sanders, J. Albert, and J. Esbjornsson (2017). Defining HIV-1 transmission clusters based on sequence data. *AIDS* 31(9), 1211–1222.
- Hellinger, E. (1909). Neue Begründung der Theorie quadratischer Formen von unendlichvielen Veränderlichen. *J. Reine Angew. Math.* 136, 210–271.
- Huang, Y., B. Niu, Y. Gao, L. Fu, and W. Li (2010). CD-HIT Suite: a web server for clustering and comparing biological sequences. *Bioinformatics* 26(5), 680–682.
- Hué, S., D. Pillay, J. P. Clewley, and O. G. Pybus (2005). Genetic analysis reveals the complex structure of hiv-1 transmission within defined risk groups. *Proceedings of the National Academy of Sciences of the United States of America* 102(12), 4425–4429.

- Kališnik Verovšek, S. (2016). Tropical coordinates on the space of persistence barcodes. arXiv.1604.00113.
- Kališnik Verovšek, S. and G. Carlsson (2014). Symmetric and r-symmetric tropical polynomials and rational functions. *Journal of Pure and Applied Algebra* 220, 3610–3627.
- Katoh, K. and D. M. Standley (2013). MAFFT multiple sequence alignment software version 7: improvements in performance and usability. *Mol. Biol. Evol.* 30(4), 772–780.
- Kerber, M., D. Morozov, and A. Nigmatov (2017). Geometry helps to compare persistence diagrams. *Journal of Experimental Algorithmics (JEA)* 22(1), 1–4.
- Koopman, B. O. (1936). On distributions admitting a sufficient statistic. *Transactions of the American Mathematical Society* 39(3), 399–409.
- Kwitt, R., S. Huber, M. Niethammer, W. Lin, and U. Bauer (2015). Statistical topological data analysis – a kernel perspective. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems 28*, pp. 3070–3078. Curran Associates, Inc.
- Maljkovic Berry, I., M. C. Melendrez, T. Li, A. W. Hawksworth, G. T. Brice, P. J. Blair, E. S. Halsey, M. Williams, S. Fernandez, I. K. Yoon, L. D. Edwards, R. Kuschner, X. Lin, S. J. Thomas, and R. G. Jarman (2016). Frequency of influenza H3N2 intra-subtype reassortment: attributes and implications of reassortant spread. *BMC Biol.* 14(1), 117.
- Mileyko, Y., S. Mukherjee, and J. Harer (2011). Probability measures on the space of persistence diagrams. *Inverse Problems* 27(12), 124007.
- Neyman, J. (1935). Su un teorema concernente le cosiddette statistiche sufficienti. *Giornale Dell’Istituto Italiano degli Attuari* 6, 320–334.
- Nguyen, T. H., V. T. Than, H. D. Thanh, V.-K. Hung, D. T. Nguyen, and W. Kim (2016). Intersubtype reassortments of h5n1 highly pathogenic avian influenza viruses isolated from quail. *PLOS ONE* 11(2), 1–15.
- Paradis, E., J. Claude, and K. Strimmer (2004). Ape: Analyses of phylogenetics and evolution in R language. *Bioinformatics* 20(2), 289–290.
- Patiño-Galindo, J. A., M. Torres-Puente, M. A. Bracho, I. Alastrué, A. Juan, D. Navarro, M. J. Galindo, C. Gimeno, E. Ortega, and F. González-Candelas (2017). Identification of a large, fast-expanding HIV-1 subtype B transmission cluster among MSM in Valencia, Spain. *PLoS ONE* 12(2), e0171062.
- Patiño-Galindo, J. Á., M. Torres-Puente, M. A. Bracho, I. Alastrué, A. Juan, D. Navarro, M. J. Galindo, D. Ocete, E. Ortega, C. Gimeno, J. Belda, V. Domínguez, R. Moreno, and F. González-Candelas (2017). The molecular epidemiology of hiv-1 in the comunidad valenciana (spain): analysis of transmission clusters. *Scientific Reports* 7, 11584.
- Perea, J. A. and G. Carlsson (2014). A klein-bottle-based dictionary for texture representation. *International Journal of Computer Vision* 1, 75–97.
- Pérez-Losada, M., M. Arenas, J. C. Galán, F. Palero, and F. González-Candelas (2015). Recombination in viruses: Mechanisms, methods of study, and evolutionary consequences. *Infection, Genetics and Evolution* 30(Supplement C), 296 – 307.
- Pitman, E. J. G. (1936). Sufficient statistics and intrinsic accuracy. *Mathematical Proceedings of the Cambridge Philosophical Society* 32(4), 567–579.
- Posada, D. and K. A. Crandall (2002). The effect of recombination on the accuracy of phylogeny estimation. *J. Mol. Evol.* 54(3), 396–402.
- Reininghaus, J., S. Huber, U. Bauer, and R. Kwitt (2015). A stable multi-scale kernel for topological machine learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

- Ronquist, F., M. Teslenko, P. van der Mark, D. L. Ayres, A. Darling, S. Höhna, B. Larget, L. Liu, M. A. Suchard, and J. P. Huelsenbeck (2012). Mrbayes 3.2: Efficient bayesian phylogenetic inference and model choice across a large model space. *Systematic Biology* 61(3), 539–542.
- Swofford, D. L. (2001). Paup\*: Phylogenetic analysis using parsimony (and other methods) 4.0.b5.
- Tian, H., S. Zhou, L. Dong, T. P. Van Boeckel, Y. Cui, S. H. Newman, J. Y. Takekawa, D. J. Prosser, X. Xiao, Y. Wu, B. Cazelles, S. Huang, R. Yang, B. T. Grenfell, and B. Xu (2015). Avian influenza h5n1 viral and bird migration networks in asia. *Proceedings of the National Academy of Sciences* 112(1), 172–177.
- Turner, K., Y. Mileyko, S. Mukherjee, and J. Harer (2014). Fréchet means for distributions of persistence diagrams. *Discrete & Computational Geometry* 52(1), 44–70.
- Turner, K., S. Mukherjee, and D. M. Boyer (2014). Persistent homology transform for modeling shapes and surfaces. *Information and Inference* 3(4), 310–344.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, Volume 3. Cambridge university press.
- Worobey, M., T. D. Watts, R. A. McKay, M. A. Suchard, T. Granade, D. E. Teuwen, B. A. Koblin, W. Heneine, P. Lemey, and H. W. Jaffe (2016). 1970s and 'Patient 0' HIV-1 genomes illuminate early HIV/AIDS history in North America. *Nature* 539(7627), 98–101.
- Yang, Z. (1995). A space-time process model for the evolution of DNA sequences. *Genetics* 139(2), 993–1005.