

Persistent Homology Lower Bounds on High-Order Network Distances

WeiYu Huang, *Student Member, IEEE*, and Alejandro Ribeiro, *Member, IEEE*

Abstract—High-order networks are weighted hypergraphs collecting relationships between elements of tuples, not necessarily pairs. Valid metric distances between high-order networks have been defined but they are difficult to compute when the number of nodes is large. The goal here is to find tractable approximations of these network distances. The paper does so by mapping high-order networks to filtrations of simplicial complexes and showing that the distance between networks can be lower bounded by the difference between the homological features of their respective filtrations. Practical implications are explored by classifying weighted pairwise networks constructed from different generative processes and by comparing the coauthorship networks of engineering and mathematics academic journals. The persistent homology methods succeed in identifying different generative models, in discriminating engineering and mathematics communities, as well as in differentiating engineering communities with different research interests.

Index Terms—Network theory, networked data, high order networks, computational topology, pattern recognition.

I. INTRODUCTION

HIGH order networks describe relationships between elements of tuples that go beyond the pairwise relationship that are more often considered [4], [5]. High order networks are important in applications in which relationships between triplets, quadruplets, and generic n -tuples are important. For example, coauthorship networks can be constructed by detailing the number of joint publications by pairs of scholars, but adding information on publications of individual authors and joint publications by specific triplets of authors generates a more accurate description. The relevance of expressing tuple relationships other than pairwise has been noted in various domains, including sensor networks [6], [7], cognitive learning [8], wireless networking [9], image ranking [10], object recognition [11], and social networks [12].

The goal of this paper is to develop efficient methods to quantify differences between high order networks of possibly

Manuscript received March 2, 2016; revised June 23, 2016; accepted October 9, 2016. Date of publication October 25, 2016; date of current version November 17, 2016. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Gustau Camps-Valls. This work was supported in part by the National Science Foundation under Grant CCF-1217963 and in part by the Multidisciplinary University Research Initiative under Grant AFOSR FA9550-10-1-0567. This paper was presented in part at the Global Conference on Signal Information Processing, Orlando, FL, USA, December 2015, in part at the IEEE International Conference on Acoustics, Speech and Signal Processing, Shanghai, China, August 2016, and in part at the Asilomar Conference on Signals Systems Computers, Pacific Grove, CA, USA, November 6–9, 2016.

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: whuang@seas.upenn.edu; aribeiro@seas.upenn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2016.2620963

different sizes. Ideally, differences should be quantified using distances derived from an underlying metric in the space of high order networks. However, while distances are not difficult to define, they can be intractable to compute. This difficulty has motivated the use of feature comparisons in which the difference between specific properties of the network is used as a more manageable alternative. Examples of feature comparisons are clustering coefficients [13], neighborhood topologies [14], betweenness [15], motifs [16], wavelets [17], and graphlet degree distributions [18]. While computationally simpler, the use of features is application dependent, utilizes only a small portion of the information conveyed by the networks, and may yield conflicting comparative judgements because the triangle inequality is not necessarily valid. A proper distance between pairwise and high order networks overcomes these drawbacks. This is the motivation for the definition of network metrics that generalize the Gromov-Hausdorff distance between metric spaces [19] to pairwise [20] and high order networks [21], [22].

The metric distances between high order networks defined in [21] have been applied to compare networks with small number of nodes and have succeeded in identifying collaboration patterns of coauthorship networks. However, because they have to consider all possible node correspondences (Definition 1), network distances are difficult to compute when the number of nodes in the networks is large. The goal of this paper is to develop network discrimination methods that are computable in networks with large numbers of nodes. These discrimination methods are constructed by drawing a parallel between high order networks and algebraic topology filtrations. Homological features of filtrations are then used to compare high order networks and shown to provide *lower bounds* for the actual network distances. Although distance lower bounds suffer from some of the same problems associated with feature comparisons, they nonetheless have important properties. Among them, we know that a large lower bound entails a large distance and that we can use lower bounds to estimate distance intervals because upper bounds are easy to determine. The idea of using persistent homology to get lower bounds is related to the development of size theory – e.g. [23] – and the consideration of multidimensional persistent Betti numbers, see e.g. [24]. Besides, the discrete approach examined in this manuscript is not restrictive but closely related with continuous approaches considered in papers such as [25].

The paper begins with a review of high order networks and high order network distances (Section II and [21]). We then introduce concepts from algebraic topology [26] and discuss the representation of high order networks in terms of filtrations (Section III). Specifically, we revisit the definition of a homological feature as a formalization of the notion of a hole without

an interior in an unweighted network and the definition of a filtration as a nested collection of unweighted networks indexed by a time parameter. The most important observation in this preliminary discussion is that a class of high order networks can be represented as a filtration that starts empty and to which nodes, edges, and high order tuples are added at a time given by their weight (Section III-A). We then introduce persistence diagrams to monitor the appearance and disappearance of homological features from the unweighted networks that compose a filtration (Section III-B).

We propose to compare networks by considering their respective filtrations and using the difference between the persistence diagrams of the respective filtrations as a proxy for their distance; this is related with the universality property studied in [27]. (Section IV). This proposed methodology is substantiated by the fact that we can lower bound the computationally intractable distance between two high order networks with a tractable distance between their respective persistence diagrams (Theorems 1 and 2 in Section IV). These lower bounds are tight such that there are examples in which the lower bound and the actual distances coincide. Since persistent homologies can be computed efficiently for large networks (Section VI and [28]), we can use these lower bounds in, e.g., network classification problems. We do so for artificial networks created with different models – random networks, Gaussian kernel proximity networks, and Euclidean feature networks – and for coauthorship networks constructed from the publications of a number of journals from engineering and mathematics communities. The proposed methods succeed in distinguishing networks with different generative models (Section VII-A) and are also effective in discriminating engineering journals from mathematics journals (Section VII-B). We also attempt a more challenging classification problem of three different engineering communities where we achieve a moderate success (Section VII-C). We close with concluding remarks (Section VIII).

II. HIGH ORDER NETWORKS

A network N_X^K of order K over the node set X is defined as a collection of $K + 1$ normalized relationship functions $\{r_X^k : X^{k+1} \rightarrow [0, 1]\}_{k=0}^K$ from the space X^{k+1} of $(k + 1)$ -tuples to the closed unit interval $[0, 1]$. For tuples $x_{0:k} := x_0, x_1, \dots, x_k$, the values $r_X^k(x_{0:k})$ of their k -order relationship functions are intended to represent a measure of similarity or dissimilarity for members of the group. We restrict attention to networks abiding to symmetry and identity properties such that the relationship function depends on the elements of the tuple but not on ordering nor repetition. This implies, e.g., that $r_X^2(x, x, x') = r_X^1(x, x')$ because x is repeated in the tuple x, x, x' and that $r_X^2(x, x', x'') = r_X^2(x', x, x'')$ because the tuples x, x', x'' and x', x, x'' contain the same elements in different order. Formally, the symmetry property requires $r_X^k(x_{[0:k]}) = r_X^k(x_{0:k})$ for any k and tuple $x_{0:k}$, where $x_{[0:k]}$ is an arbitrary reordering of $x_{0:k}$. The identity property requires $r_X^k(x_{0:k}) = r_X^{\bar{k}}(x_{l_0:l_{\bar{k}}})$ for any k and tuple $x_{0:k}$, where $x_{l_0:l_{\bar{k}}}$ is a subtuple of $x_{0:k}$ such that the number of unique elements is the same in $x_{l_0:l_{\bar{k}}}$ and $x_{0:k}$. Observe that the network

N_X^K can be considered as a weighted complete hypergraph [29] whose weights for all possible hyperedges are defined. Since relationships are not explicitly defined for all hyperedges in many cases, we adopt the convention that the hyperedges with undefined relationships take the value 0 or 1, depending on whether the relationships encode dissimilarities or similarities (see Section II-A and Section V). The value of K is no greater than 2 in most applications. The set of all high order networks of order K is denoted as \mathcal{N}^K .

Two networks N_X^K and N_Y^K are said k -isomorphic if there exists a bijection $\pi : X \rightarrow Y$ such that for all $x_{0:k} \in X^{k+1}$ we have

$$r_Y^k(\pi(x_{0:k})) = r_X^k(x_{0:k}), \quad (1)$$

where we use the shorthand notation $r_Y^k(\pi(x_{0:k})) := r_Y^k(\pi(x_0), \pi(x_1), \dots, \pi(x_k))$. Since the map π is bijective, (1) can only be satisfied when X is a permutation of Y . When networks N_X^K and N_Y^K are k -isomorphic we write $N_X^K \cong_k N_Y^K$. The space of K -order networks where k -isomorphic networks are represented by the same element is termed the set of K -networks modulo k -isomorphism and denoted by $\mathcal{N}^K \text{ mod } \cong_k$. For each $0 \leq k \leq K$, the space $\mathcal{N}^K \text{ mod } \cong_k$ can be endowed with a pseudometric [21]. The definition of this distance requires introducing the notion of correspondence [30, Def. 7.3.17]:

Definition 1: A correspondence between two sets X and Y is a subset $C \subseteq X \times Y$ such that $\forall x \in X$, there exists $y \in Y$ such that $(x, y) \in C$ and $\forall y \in Y$ there exists $x \in X$ such that $(x, y) \in C$. The set of all correspondences is denoted as $\mathcal{C}(X, Y)$.

A correspondence in the sense of Definition 1 connects node sets X and Y so that every element of each set has at least one correspondent in the other set. We can now define the distance between two networks by selecting the correspondence that makes them the most similar as we formally define next.

Definition 2: Given networks N_X^K and N_Y^K , a correspondence C between the node sets X and Y , and an integer $0 \leq k \leq K$ define the k -order network difference with respect to C as

$$\Gamma_{X,Y}^k(C) := \max_{(x_{0:k}, y_{0:k}) \in C} |r_X^k(x_{0:k}) - r_Y^k(y_{0:k})|, \quad (2)$$

where $(x_{0:k}, y_{0:k})$ stands for $(x_0, y_0), (x_1, y_1), \dots, (x_k, y_k)$. The k -order network distance between N_X^K and N_Y^K is defined as

$$d_{\mathcal{N}}^k(N_X^K, N_Y^K) := \min_{C \in \mathcal{C}(X,Y)} \{\Gamma_{X,Y}^k(C)\}. \quad (3)$$

For a given correspondence $C \in \mathcal{C}(X, Y)$ the network difference $\Gamma_{X,Y}^k(C)$ selects the maximum distance difference $|r_X^k(x_{0:k}) - r_Y^k(y_{0:k})|$ among all pairs of correspondents – we compare $r_X^k(x_{0:k})$ with $r_Y^k(y_{0:k})$ when all the points $x_l \in x_{0:k}$ and $y_l \in y_{0:k}$ are correspondents. The distance in (3) is defined by selecting the correspondence that minimizes these maximal differences. The distance in Definition 2 is a pseudometric in $\mathcal{N}^K \text{ mod } \cong_k$ [21]. Observe that since correspondences may be between networks with different number of elements,

$d_{\mathcal{N}}^k(N_X^K, N_Y^K)$ is defined when the node cardinalities $|X|$ and $|Y|$ are different.

The distances in (2) compare relationship functions of a given order k . To compare all orders $0 \leq k \leq K$ we begin by defining two K -order networks N_X^K and N_Y^K as isomorphic if there exists a bijection $\pi : X \rightarrow Y$ such that (1) holds for all $0 \leq k \leq K$ and $x_{0:k} \in X^{k+1}$. When networks N_X^K and N_Y^K are isomorphic we write $N_X^K \cong N_Y^K$. The space of K -order networks modulo isomorphism is denoted as $\mathcal{N}^K \text{ mod } \cong$. A family of pseudo-metrics measuring the difference over all order functions can be endowed to $\mathcal{N}^K \text{ mod } \cong$ as we formally state next.

Definition 3: Given networks N_X^K and N_Y^K , a correspondence C between the node sets X and Y , and some vector p -norm $\|\cdot\|_p$, define the network difference with respect to C as

$$\|\Gamma_{X,Y}^K(C)\|_p := \left\| \left(\Gamma_{X,Y}^0(C), \Gamma_{X,Y}^1(C), \dots, \Gamma_{X,Y}^K(C) \right)^T \right\|_p, \quad (4)$$

where $\Gamma_{X,Y}^k(C)$ is the k -order difference as per C defined in (2). The p -norm network distance between N_X^K and N_Y^K is defined as

$$d_{\mathcal{N},p}(N_X^K, N_Y^K) := \min_{C \in \mathcal{C}(X,Y)} \left\{ \|\Gamma_{X,Y}^K(C)\|_p \right\}. \quad (5)$$

In Definition 3 we compare relationship functions of all orders no larger than K . The norm $\|\Gamma_{X,Y}^K(C)\|_p$ is assigned as the difference between N_X^K and N_Y^K measured by the correspondence C . The distance $d_{\mathcal{N},p}(N_X^K, N_Y^K)$ is then the minimum of these differences achieved by some correspondence.

The k -order function r_X^k of a network N_X^K does not impose constraints on the l -order function r_X^l of the same network. In practical situations, however, it is common that adding nodes to a tuple results in increasing or decreasing relationships for the extended tuple. This motivates dissimilarity networks, which we do next, and proximity networks, which we discuss in Section V.

A. Dissimilarity Networks

In dissimilarity networks, $r_X^k(x_{0:k})$ encodes a level of dissimilarity between elements. In this scenario it is reasonable that adding elements to a tuple makes the group more dissimilar and results in a higher value in the relationship function. This restriction makes up the formal definition that we introduce next.

Definition 4: We say that the K -order network $D_X^K = (X, r_X^0, \dots, r_X^K)$ is a dissimilarity network if the order increasing property holds, i.e. for any order $1 \leq k \leq K$ and tuples $x_{0:k} \in X^{k+1}$ we have

$$r_X^k(x_{0:k}) \geq r_X^{k-1}(x_{0:k-1}), \quad (6)$$

and the inequality (6) equalizes if and only if all the elements in $x_{0:k}$ also appear in $x_{0:k-1}$. Denote the set of all dissimilarity networks of order K as \mathcal{D}^K .

As per Definition 4 dissimilarities $r_X^k(x_{0:k})$ grow as we add elements to a tuple. Notice that due to symmetry, a relationship as in (6) holds if we remove an arbitrary node from the tuple $x_{0:k}$, not necessarily the last. We further emphasize that Definition 4 includes the condition of having the inequality in (6) be strict

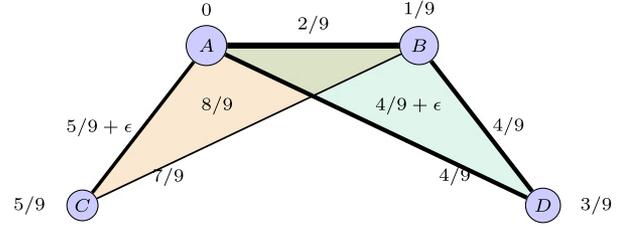


Fig. 1. Temporal dynamics for the formation of a research community. The k -order relationship in this 2-order dissimilarity network [cf. Definition 4] denotes the normalized time instant when members of a $(k+1)$ -tuple write their first joint paper. E.g., A writes her first paper at time 0, and coauthors with B , D , C at times $2/9$, $4/9$, $5/9$, respectively. The ϵ in the relationship function is a technical modification to differentiate relationship between two authors from that of a single author. Author A also writes jointly with B and D at time $4/9$.

unless all the elements of $x_{0:k}$ appear in $x_{0:k-1}$. This implies that to have $r_X^k(x_{0:k}) = r_X^{k-1}(x_{0:k-1})$ it must be that at least one element of $x_{0:k}$ is repeated in the tuple. Conversely, when we add a *distinct* element to a tuple, the dissimilarity must increase, i.e. $r_X^k(x_{0:k-1}, x_k) > r_X^{k-1}(x_{0:k-1})$ if x_k is not a repetition of an element of $x_{0:k-1}$. This condition is technical and of little consequence for the results presented later – see Remark 2.

To see why the order increasing property is reasonable, consider a network describing the temporal dynamics of the formation of a research community – see Figure 1. The relationship function in this network marks the normalized time instant at which members of a given $(k+1)$ -tuple write their first joint paper. For $k=0$, the zeroth order dissimilarities r_X^0 are the normalized time instants when authors publish their first paper. In Figure 1 authors A , B , C , and D publish their first papers at times 0, $1/9$, $5/9$, and $3/9$. The first order dissimilarities r_X^1 between pairs denote times at which nodes become coauthors. Since a coauthored paper is also a paper by each of the authors, it is certain that $r_X^1(x, x') \geq r_X^0(x)$. In Figure 1, A and B become coauthors at time $2/9$, which occurs after they publish their respective first papers at times 0 and $1/9$. Authors A and D become coauthors at time $4/9$ and A and C become coauthors at time $5/9$. Second order dissimilarities r_X^2 denote the time at which a paper is coauthored by a triplet. Since a paper can't be coauthored by three people without being at the same time coauthored by each of the three possible pairs of authors we must have $r_X^2(x, x', x'') \geq r_X^1(x, x')$. In Figure 1, authors A , B , and D publish a joint paper at time $4/9$, which is no smaller than the pairwise coauthorship times between each two of the individual authors. Undefined hyperedge relationships in dissimilarity networks take the value 1 indicating the maximum dissimilarities for the tuple. An additive constant ϵ is used to make $r_X^1(A, C)$ strictly larger than $r_X^0(C)$ and $r_X^2(A, B, D)$ strictly larger than $r_X^1(A, D)$ and $r_X^1(B, D)$. This is needed because of the requirement in Definition 4 that the equality in (6) be achieved only if the elements in $x_{0:k}$ are all elements of $x_{0:k-1}$.

The requirement considered in dissimilarity networks can also be justified from a probability point of view. In specific, we assume that each node x in the network has an underlying random observation X . Consider $\mathbb{E}[X_1^2]$ and $\mathbb{E}[X_2^2]$ as the 0-order proximities for nodes x_1 and x_2 (higher value

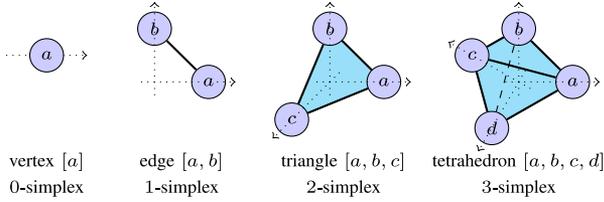


Fig. 2. Elementary k -simplices for $0 \leq k \leq 3$.

means more important node), and consider $\mathbb{E}[X_1 X_2]$ as the 1-order proximity between the pair of nodes (higher value means that the observations at the pair are more likely to be correlated). Cauchy-Schwarz inequality then implies that $\mathbb{E}[X_1 X_2] \leq \sqrt{\mathbb{E}[X_1^2] \mathbb{E}[X_2^2]}$. Therefore, the level of proximity between nodes decreases as we include more nodes into consideration, or equivalently, dissimilarity increases as we consider more nodes, which is same as the requirement for relationships of dissimilarity networks. Notice that the inequality takes equality if and only if the observation X_1 and X_2 are same with probability 1, which is akin to the condition in dissimilarity networks that two nodes are essentially represented by the same node.

When specialized to dissimilarity networks, the network distances in definitions 2 and 3 are termed the k -order dissimilarity network distance d_D^k and the p -norm dissimilarity network distance $d_{D,p}$. The restriction to dissimilarities makes the distances well-defined *metrics* in their respective spaces [21]. Although these metrics provide well-founded methods to compare high order networks, the search for the optimal correspondences in (3) and (5) is combinatorial. This makes it impossible to evaluate distances when the number of nodes in the networks is large and motivates the development of tractable lower bounds. These bounds are obtained by relating dissimilarity networks to filtrations of simplicial complexes [26], [31] as we explain next.

III. NETWORKS AND SIMPLICIAL COMPLEXES

We introduce elemental notions of computational topology as they apply to the study of high order networks. We refer authors to [32] for more background and more formal notions on computational homology. Begin by defining a k -simplex $\phi = [x_{0:k}]$ as the convex hull of the set of points $x_{0:k}$ (see Figure 2) and a simplicial complex L as the collection of simplices such that for any simplex $[x_{0:k}]$, the convex hull of any subset of $x_{0:k}$ also belongs to L . Figure 3(a) exemplifies two triangles connected together as a simplicial complex of dimension 2. It is a simplicial complex because for any of its simplices, say, $[a, b, c]$, the convex hulls of the subsets of $\{a, b, c\}$ – which include the points $[a]$, $[b]$, and $[c]$ as well as the edges $[a, b]$, $[b, c]$, and $[a, c]$ – all belong to the simplicial complex as well.

An important concept in simplicial complexes is that of a hole without interior. For the simplicial complex in Figure 3(a), the area enclosed by $[b, d]$, $[d, c]$, $[c, b]$ is a hole without interior because the area is not filled. The area enclosed by $[a, b]$, $[b, d]$, $[d, a]$ is not because the interior is filled by the 2-simplex $[a, b, d]$. Homologies are defined to formalize this intuition and rely on the definitions of chains, cycles,

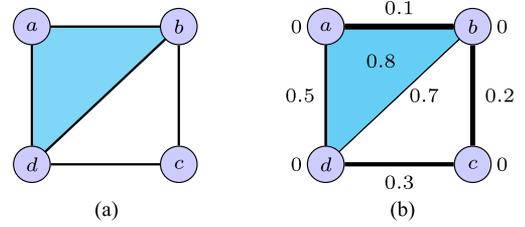


Fig. 3. (a) An example of simplicial complex L which consists of four 0-simplices, five 1-simplices, and one 2-simplex. (b) A dissimilarity network can be represented as a simplicial complex with weights. When the weight of a simplex denotes the time instant the simplex appears in the nested sequence of simplicial complexes, it yields a valid filtration \mathcal{L} .

and boundaries. The k -chain $\Phi_k = \sum_i \beta_i \phi_i$ is a summation of k -simplices ϕ_i modulated by coefficients β_i whose signs denote orientation. This definition is a generalization of the familiar definition of chains in graphs. E.g., in Figure 3, $[a, b] + [b, d]$ is a 1-chain that we can equivalently represent as $[a, b] - [d, b]$. Further consider a given k -simplex $\phi = [x_{0:k}]$ and its border $(k-1)$ -simplices defined as the ordered set of elements $[x_{0:\hat{l}:k}] = \text{conv}\{x_{0:k} \setminus x_l\}$, in which each of the elements is removed in order. The *boundary* $\partial_k \phi$ of the simplex ϕ is the chain formed by its borders using alternating orientations,

$$\partial_k \phi = \sum_{l=0}^k (-1)^l [x_{0:\hat{l}:k}]. \quad (7)$$

According to (7), the boundary $\partial_k \phi$ of a k -simplex is the collection of $(k-1)$ -simplices. For the simplices in Figure 2, the boundaries are $\partial_0[a] = 0$, $\partial_1[a, b] = [b] - [a]$, $\partial_2[a, b, c] = [b, c] - [a, c] + [a, b]$ and $\partial_3[a, b, c, d] = [b, c, d] - [a, c, d] + [a, b, d] - [a, b, c]$.

Having defined chains and boundaries we consider a k -chain $\Phi_k = \sum_i \beta_i \phi_i$ and define the *chain boundary* $\partial_k \Phi_k$ as the summation of the boundaries of its component simplices,

$$\partial_k \Phi_k = \sum_i \beta_i (\partial_k \phi_i). \quad (8)$$

As per this definition, the boundary of the chain $\Phi = [a, b] + [b, d]$ in Figure 3 is $\partial_k \Phi = [a] - [b] + [b] - [d] = [a] - [d]$, which matches our intuition of what the chain's boundary should be. We can now formally define a k -cycle Ψ_k as a chain whose boundary is null, i.e., a k -chain for which $\partial_k \Psi = 0$. In Figure 3 the chain $\Psi = [a, b] + [b, d] + [d, a]$ is a cycle because its boundary is $\partial_k \Psi = [a] - [b] + [b] - [d] + [d] - [a] = 0$.

A k -cycle Ψ_k is a k -homological feature if it cannot be represented as the boundary of a $(k+1)$ -chain Φ_{k+1} ; otherwise Ψ_k is a k -boundary. We refer readers to [32] for the more formal definition. The group $\mathcal{H}_k(L)$ of k -homological features of a simplicial complex L represents all k -cycles that are not k -boundaries. I.e., a homological feature appears when a k -cycle Ψ_k is not among the group of boundaries $\partial_{k+1} \Phi_{k+1}$ of $(k+1)$ -chains.

For the simplicial complex in Figure 3(a), the chain $\Psi_1 = [a, b] + [b, d] + [d, a]$ is a 1-cycle because $\partial_1 \Psi_1 = 0$; simultaneously, Ψ_1 is also a 1-boundary because $\partial_2[a, b, d] = [b, d] - [a, d] + [a, b]$, which is identical to Ψ_1 . Therefore, Ψ_1 does not

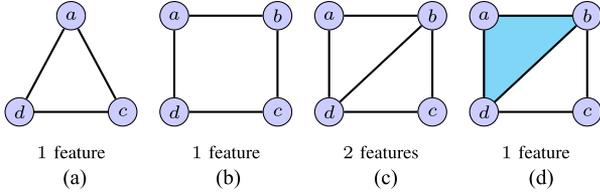


Fig. 4. More examples of simplicial complexes and homologies. The number of 2-homological features are described below each of the simplicial complex.

represent a homological feature. On the other hand, the chain $\Psi'_1 = [b, c] + [c, d] + [d, b]$ describes a homological feature because it is a cycle with $\partial_1 \Psi'_1 = 0$, and not a boundary due to the fact that the 2-simplex $[b, c, d]$ is absent from the complex. One could also say that $\Psi''_1 = [a, b] + [b, c] + [c, d] + [d, a]$ describes a 1-cycle that is not a 1-boundary. Nonetheless, because the difference between Ψ''_1 and Ψ'_1 is $[a, b] + [b, d] + [d, a]$ which is a 1-boundary, the cycles Ψ''_1 and Ψ'_1 are essentially denote the same homological feature. Therefore, there only exists one 1-homological feature, which may be represented by chain Ψ'_1 or Ψ''_1 .

Other illustrative examples are shown in Figure 4. The complex (a) has one 1-homological feature whose one representation is the chain $[a, b] + [b, c] + [c, a]$ and the complex (b) has 1 feature which can be represented by the chain $[a, b] + [b, c] + [c, d] + [d, a]$. Complex (c) has 2 features that may be denoted by chains $[a, b] + [b, d] + [d, a]$ and $[b, c] + [c, d] + [d, b]$ but complex (d) has one feature because the chain $[a, b] + [b, d] + [d, a]$ is the boundary of the simplex $[a, b, d]$.

A. Representation of High Order Networks as Filtrations

In the same manner in which a graph represents an unweighted network, a simplicial complex represents a high order unweighted network. To represent weighted high order networks we assign a weight to each simplex, but instead of thinking the network as a weighted simplicial complex we think of weights as parameters that indicates the time at which the simplex comes into existence. Formally, for parameters $\alpha \in [0, 1]$ we define a filtration \mathcal{L} as a collection of simplicial complexes L_α such that for any ordered sequence $0 = \alpha_0 < \alpha_1 < \dots < \alpha_m = 1$ it holds $\emptyset = L_{\alpha_0} \subseteq L_{\alpha_1} \subseteq \dots \subseteq L_{\alpha_m} = L$. The minimum time α at which a simplex becomes an element of L_α is the birth time of the simplex.

Given a dissimilarity network D_X^K , we construct $\mathcal{L}(D_X^K)$ by assigning the appearing time of simplex $[x_{0:k}]$ as the relationship $r_X^k(x_{0:k})$ of the corresponding tuple,

$$[x_{0:k}] \in L_\alpha \iff r_X^k(x_{0:k}) \leq \alpha. \quad (9)$$

We show next that (9) defines a valid filtration.

Proposition 1: The construction $\mathcal{L}(D_X^K)$ with L_α established from a dissimilarity network D_X^K via the relationship in (9) is a well defined filtration.

Proof: For a given α , (9) defines a set of simplices L_α . It is clear that the collection of sets L_α is nested, i.e., that $\emptyset = L_{\alpha_0} \subseteq \dots \subseteq L_{\alpha_m} = L$ holds for any set of birth times $0 = \alpha_0 < \dots < \alpha_m = 1$. To prove the statement we need to show

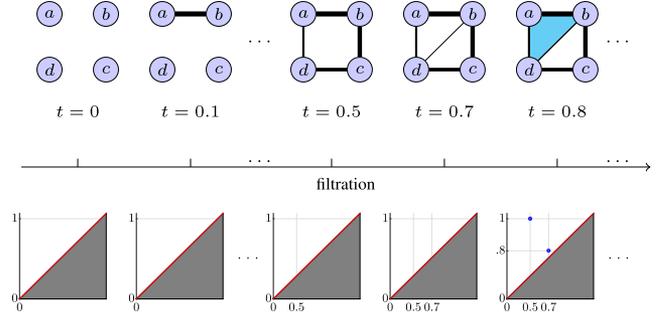


Fig. 5. Top: filtration of nested simplicial complexes of the dissimilarity network exhibited in Figure 3(b). The simplicial complex at each time instant are the collections of all vertices appearing before or on that time. Bottom: 1-persistence diagrams describing 1-persistent homologies for each of the simplicial complexes detailed in Top. The horizontal axis denotes the birth time of homological features and the vertical axis represents the death time.

that the set of simplices L_α is a valid simplicial complex. To show this it suffices to verify that for any α , all faces of each simplex in L_α also appear no later than α . Suppose simplex $[x_{0:k}]$ appears before time α for some chosen α , then it must be true that $r_X^k(x_{0:k}) \leq \alpha$. For any faces of $[x_{0:k}]$, say $[x_{0:\bar{k}}]$ with $\bar{k} < k$, the order increasing property implies $r_X^{\bar{k}}(x_{0:\bar{k}}) \leq r_X^k(x_{0:k})$. Therefore,

$$r_X^{\bar{k}}(x_{0:\bar{k}}) \leq r_X^k(x_{0:k}) \leq \alpha, \quad (10)$$

which shows that the face $[x_{0:\bar{k}}]$ appears before time or on time α . This means that L_α is a valid simplicial complex. ■

We emphasize that filtrations can't be defined for an arbitrary high order networks because the order increasing property is necessary for the proof. For dissimilarity networks that do satisfy the order increasing property, Proposition 1 establishes the existence of an equivalent filtration.

Figure 5 (top) shows the construction of the filtration associated with the dissimilarity network in Figure 3(b). At time $\alpha = 0$, there are four nodes because $r_X^0(a) = r_X^0(b) = r_X^0(c) = r_X^0(d) = 0$. At time $\alpha = 0.1$, the edge $[a, b]$ starts to appear in the simplicial complex because the relationship $r_X^1(a, b) = 0.1$. As time gradually increases, more simplices get included in the simplicial complex at each time instant. At the end of the filtration with $\alpha = 1$, all simplices in the simplicial complex, whose corresponding relationships in the dissimilarity networks are either defined or undefined, are involved.

B. Persistent Homologies and Persistence Diagrams

As time advances in a filtration, holes and interiors appear and disappear. Persistent homologies examine when these homological features appear for the sequence of simplicial complexes in the filtration. Formally, given a filtration \mathcal{L} , its k -dimensional persistence diagram $\mathcal{P}_k \mathcal{L}$ is a collection of points of the form $\mathbf{q} = [q_b, q_d]$ where q_b and $q_d > q_b$ represent the birth and death time of a homological feature. I.e., the times q_b represent resolutions $\alpha = q_b$ at which a feature is added to the group of k -homological features $\mathcal{H}_k(L_\alpha) = \mathcal{H}_k(L_{q_b})$ and the times q_d are resolutions $\alpha = q_d$ at which a feature is removed from the group of k -homological features $\mathcal{H}_k(L_\alpha) = \mathcal{H}_k(L_{q_d})$. We refer

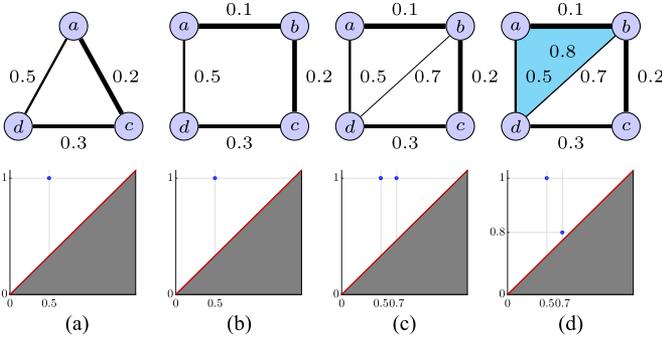


Fig. 6. The 1-persistence diagrams for the filtration induced by each of the dissimilarity networks. The maximum death time of any homological features is 1 because hyperedges with undefined relationships take value 1 implicitly.

authors to [32] for more background and more formal definitions on persistent homology, and refer to [33] for a more comprehensive treatment of persistence and its connection to application.

The persistence diagram for the filtration in Figure 5-(top) is shown in Figure 5-(bottom). In this diagram the horizontal axis denotes the birth time of 1-homological features (when holes appear) and the vertical axis represents the death time of 1-homological features (when holes are filled). At time $\alpha = 0$, there are no 1-simplices and consequently no 1-holes. The first 1-cycle appears at time $\alpha = 0.5$ when the 1-chain $[a, b] + [b, c] + [c, d] + [d, a]$ appears. We mark this event by the addition of a vertical line in the persistence diagram. A second homological feature appears in the form of the cycle $[a, b] + [b, d] + [d, a]$ at time $\alpha = 0.7$. This event is marked by the addition of a second vertical line in the persistence diagram. This homological feature dies at time $\alpha = 0.8$ when the 2-simplex $[a, b, d]$ appears and makes the 1-cycle $[a, b] + [b, d] + [d, a]$ a 1-boundary and therefore no longer a homological feature. This is marked by the addition of the point $\mathbf{q}_1 = (0.7, 0.8)$ to the persistence diagram. At this time all simplices with defined relationship functions have been added to the filtration. This means that the homological feature $[a, b] + [b, c] + [c, d] + [d, a]$ never becomes a boundary of a 2-simplex. This is marked by adding the point $\mathbf{q}_2 = (0.5, 1)$ to the persistence diagram – recall that dissimilarities must be between 0 and 1 by definition. All undefined simplices, e.g. $[a, c]$, $[a, c, d]$, $[a, b, c, d]$, are considered to appear at 1 and would not affect the persistence diagram.

A few more examples are shown in Figure 6. Network (a) contains one feature that appears at $\alpha = 0.5$ and stays alive until $\alpha = 1$. Network (b) also has one feature born at $\alpha = 0.5$ that never gets trivialized until $\alpha = 1$. Network (c) involves two features, born at $\alpha = 0.5$ and $\alpha = 0.7$ that also stay alive until $\alpha = 1$. Network (d) is the same as the filtration in Figure 5-(top).

We close this section by noting all dissimilarity values between tuples with non-repeating elements of any dissimilarity network appear in the homological features of the induced filtration.

Proposition 2: Given a dissimilarity network D_X^K , any of its k -order dissimilarities between tuples with unique elements appear either in the death time of the $(k - 1)$ -th dimensional

homological features or the birth time of the k -th dimensional features.

Proof: See Appendix A. ■

Proposition 2 is specific to filtrations induced from dissimilarity networks since it follows from the fact that adding elements to a tuple results in *strictly increasing* dissimilarities [cf. Definition 4]. The proposition implies that different networks result in different persistence diagrams except in the rare cases when a single point in the diagram represents multiple homological features. Thus, persistence diagrams retain almost all of the dissimilarity values of a given network and are therefore not unreasonable proxies for network discrimination. We cement these observations in the next section by proving that differences between persistence diagrams yield lower bounds on the network distances defined in Section II. We do so after an important remark.

Remark 1: The maximum death time of any homological features in filtrations induced from dissimilarity networks is 1. This is specific for filtrations constructed from dissimilarity networks because any hyperedges with undefined relationships in dissimilarity networks all take value 1 implicitly. In computing the persistence diagram, we can either (i) set the appearance time of all hyperedges with undefined relationships as 1 before evaluating the persistence diagrams, or (ii) determine the persistence diagram with no additional care of undefined edges, and then set the death time of any undead features to 1. These two procedures yield the same persistence diagrams, except for the first 0-dimensional persistent interval which dies at ∞ and represents the connected component. We use the latter method in practice because it is computational simpler.

IV. PERSISTENCE BOUNDS ON NETWORK DISTANCES

In this section we use differences between persistence diagrams to compute lower bounds of network distances. Recall that a persistence diagram $\mathcal{P}_k \mathcal{L}(D_X^K)$ is a collection of points of the form $\mathbf{q} = [q_b, q_d]$ where q_b and $q_d > q_b$ represent the birth and death time of a k -homological feature [cf. Figure 6]. To compare persistence diagrams $\mathcal{P}_k \mathcal{L}(D_X^K)$ and $\mathcal{P}_k \tilde{\mathcal{L}}(D_Y^K)$ of networks D_X^K and D_Y^K we begin by defining the cost of matching features $\mathbf{q} \in \mathcal{Q}$ and $\tilde{\mathbf{q}} \in \tilde{\mathcal{Q}}$ through the infinity norm of their difference,

$$\|\mathbf{q} - \tilde{\mathbf{q}}\|_\infty := \max[|q_b - \tilde{q}_b|, |q_d - \tilde{q}_d|]. \quad (11)$$

Further observe that the diagonal of a persistence diagram can be construed to represent an uncountable number of features with equal birth and death times. Thus, any feature $\mathbf{q} \in \mathcal{Q}$ can be compared to any of the artificial features of the form $\tilde{\mathbf{q}} = [\tilde{q}, \tilde{q}]$. Since this feature $\tilde{\mathbf{q}}$ can be placed anywhere in the diagonal of the persistence diagram, we choose to place it in the point that makes the infinity norm difference smallest. This point is $\tilde{\mathbf{q}} = [(q_b + q_d)/2, (q_b + q_d)/2]$ which yields the difference $\|\mathbf{q} - \tilde{\mathbf{q}}\|_\infty = |q_d - q_b|/2$. Likewise, artificial features $\tilde{\mathbf{q}} = [(\tilde{q}_b + \tilde{q}_d)/2, (\tilde{q}_b + \tilde{q}_d)/2]$ can be added for any point $\tilde{\mathbf{q}} \in \tilde{\mathcal{Q}}$. We can then rephrase the comparison in (11) so that it is more advantageous to compare with artificial diagonal features. This is formally accomplished by defining the matching

cost $c(\mathbf{q}, \tilde{\mathbf{q}})$ between \mathbf{q} and $\tilde{\mathbf{q}}$ as

$$c(\mathbf{q}, \tilde{\mathbf{q}}) := \min \left[\|\mathbf{q} - \tilde{\mathbf{q}}\|_\infty, (1/2) \max [|q_d - q_b|, |\tilde{q}_d - \tilde{q}_b|] \right], \quad (12)$$

The norm $\|\mathbf{q} - \tilde{\mathbf{q}}\|_\infty$ is the cost of directly matching features \mathbf{q} and $\tilde{\mathbf{q}}$. The term $(1/2) \max [|q_d - q_b|, |\tilde{q}_d - \tilde{q}_b|]$ is the cost of matching both, \mathbf{q} and $\tilde{\mathbf{q}}$ to artificial diagonal features. The cost $c(\mathbf{q}, \tilde{\mathbf{q}})$ of matching \mathbf{q} and $\tilde{\mathbf{q}}$ is the smaller of these two.

If the respective collections of features \mathcal{Q} and $\tilde{\mathcal{Q}}$ contain the same number of elements $m = \tilde{m}$ we can consider bijections $\pi : \mathcal{Q} \rightarrow \tilde{\mathcal{Q}}$ from \mathcal{Q} to $\tilde{\mathcal{Q}}$. The bottleneck distance between the persistence diagrams $\mathcal{P}_k \mathcal{L}(D_X^K)$ and $\mathcal{P}_k \tilde{\mathcal{L}}(D_Y^K)$ of networks D_X^K and D_Y^K is then defined as

$$b^k(D_X^K, D_Y^K) := \min_{\pi: \mathcal{Q} \rightarrow \tilde{\mathcal{Q}}} \max_{\mathbf{q} \in \mathcal{Q}} c(\mathbf{q}, \pi(\mathbf{q})). \quad (13)$$

The distance in (13) is set to the pair of features \mathbf{q} and $\tilde{\mathbf{q}}$ that are most difficult to match across all possible bijections π .

When the number of points in the persistence diagrams are different we assume without loss of generality that $\tilde{m} < m$. In this case we extend $\tilde{\mathcal{Q}}$ by adding $\tilde{m} - m$ diagonal features to define the set $\tilde{\mathcal{Q}}_e := \tilde{\mathcal{Q}} \cup \{(\tilde{q}_i, \tilde{q}_i)\}_{i=1}^{\tilde{m}-m}$. Since the sets \mathcal{Q} and $\tilde{\mathcal{Q}}_e$ contain the same number of elements bijections $\pi : \mathcal{Q} \rightarrow \tilde{\mathcal{Q}}_e$ are well defined. We can then modify (13) to define a valid comparison between $\mathcal{P}_k \mathcal{L}$ and $\mathcal{P}_k \tilde{\mathcal{L}}$ when the number of points in the diagrams are possibly different. We do so in the following formal definition.

Definition 5: Given persistence diagrams $\mathcal{P}_k \mathcal{L}(D_X^K)$ and $\mathcal{P}_k \tilde{\mathcal{L}}(D_Y^K)$ with sets of points \mathcal{Q} and $\tilde{\mathcal{Q}}$ having cardinalities $\tilde{m} < m$, define the extended set $\tilde{\mathcal{Q}}_e := \tilde{\mathcal{Q}} \cup \{(\tilde{q}_i, \tilde{q}_i)\}_{i=1}^{\tilde{m}-m}$ by adding $\tilde{m} - m$ artificial diagonal features. The bottleneck distance between the persistence diagrams $\mathcal{P}_k \mathcal{L}(D_X^K)$ and $\mathcal{P}_k \tilde{\mathcal{L}}(D_Y^K)$ of networks D_X^K and D_Y^K is defined as

$$b^k(D_X^K, D_Y^K) := \min_{\pi: \mathcal{Q} \rightarrow \tilde{\mathcal{Q}}_e} \max_{\mathbf{q} \in \mathcal{Q}} c(\mathbf{q}, \pi(\mathbf{q})), \quad (14)$$

where π ranges over all bijections from \mathcal{Q} to $\tilde{\mathcal{Q}}_e$ and the cost $c(\mathbf{q}, \pi(\mathbf{q}))$ is defined in (12).

The number of bijections between two sets of points is factorial and it appears that the problem as in (14) is as difficult as the problem of finding the correspondence in evaluating the network distance. However, the problem in (14) is a instantiation of the Linear Bottleneck Assignment Problem (LBAP) that can be solved efficiently – see Section VI and [34, Algorithm 6.1]. We emphasize that $c(\mathbf{q}, \tilde{\mathbf{q}}) = c(\mathbf{q}, \tilde{\mathbf{q}}')$ for any \mathbf{q} whenever $\tilde{\mathbf{q}}$ and $\tilde{\mathbf{q}}'$ are on the diagonal. Therefore, the locations of diagonal points added to construct $\tilde{\mathcal{Q}}_e$ are unsubstantial as per their cost. The definition of b^k has been used in [35] to compare three-dimensional surfaces using persistent homology.

We prove now that the bottleneck distance between the persistence diagrams of the filtrations induced by two dissimilarity networks is a lower bound of their dissimilarity network distance.

Theorem 1: Let D_X^K and D_Y^K be two K -order dissimilarity networks. The bottleneck distance between the k -th dimensional persistence diagrams of the filtrations $\mathcal{L}(D_X^K)$ and $\mathcal{L}(D_Y^K)$ is at

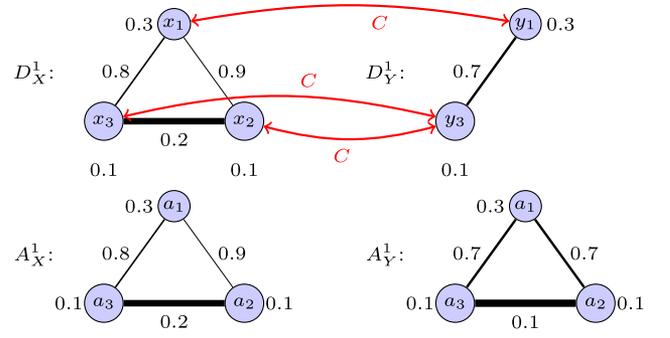


Fig. 7. An example to construct augmented networks. D_X^1 and D_Y^1 are 1-order dissimilarity networks. The augmented network using correspondence C induces $A_X^1 \cong D_X^1$. An additional node a_2 and two edges are augmented in A_Y^1 .

most $d_{\mathcal{D},\infty}(D_X^K, D_Y^K)$ for any $0 \leq k \leq K$, i.e.

$$b^k(D_X^K, D_Y^K) \leq d_{\mathcal{D},\infty}(D_X^K, D_Y^K). \quad (15)$$

Proof: The proof of Theorem 1 relies on the stability theorem for persistence diagrams (also termed Interleaving Theorem) [36], which state that if the birth times of each simplex in two filtrations differ by no more than δ , the bottleneck distance between the corresponding persistence diagrams is no greater than δ as well. In order to prove Theorem 1, we introduce the following notion of augmented networks. This definition solves the issue when a single node in one network has multiple correspondents in the other network. In such cases, say, if a in X , has both c and d in Y as correspondents, we examine the difference between $r_X^1(a, a)$ and $r_Y^1(c, d)$ in evaluating the network difference. However $r_X^1(a, a) = r_X^0(a)$ represents the birth time of $[a]$ while $r_Y^1(c, d)$ denotes the time of appearance of an edge $[c, d]$. The augmented networks are introduced to resolve this discrepancy.

Definition 6: Given two K -order dissimilarity networks D_X^K and D_Y^K and a correspondence C between their node sets X and Y , the augmented networks $A_{X,C}^K$ and $A_{Y,C}^K$ are a pair of K -order networks defined on C . Each node a_i in $A_{X,C}^K$ and $A_{Y,C}^K$ represents a correspondent pair (x_{C_i}, y_{C_i}) in C . Relationship functions for $a_{0:k}$ with $0 \leq k \leq K$ are defined as

$$r_{A_X}^k(a_{0:k}) = r_X^k(x_{C_0:C_k}), \quad r_{A_Y}^k(a_{0:k}) = r_Y^k(y_{C_0:C_k}). \quad (16)$$

When the underlying correspondence is clear, C is omitted in the subscripts. For a pair of dissimilarity networks D_X^K and D_Y^K and a correspondence C between X and Y , the augmented networks A_X^K and A_Y^K have identical $|C|$ nodes where $|C|$ denotes the number of correspondence pairs in C . Each node a_i in both A_X^K and A_Y^K represents a correspondent pair (x_{C_i}, y_{C_i}) in C . For each tuple $a_{0:k}$, its relationship $r_{A_X}^k(a_{0:k})$ for the network A_X^K is the same as the relationship $r_X^k(x_{C_0:C_k})$ between the tuple $x_{C_0:C_k}$ in D_X^K . An example to construct augmented networks is illustrated in Figure 7. Augmented networks satisfy the following properties.

Fact 1: Given a K -order dissimilarity network D_X^K and a correspondence C between the node sets X and Y , the augmented network A_X^K constructed by Definition 6 is a valid

K -order network. The induced filtration $\mathcal{L}(A_X^K)$ is a valid filtration and has identical persistence diagrams as $\mathcal{L}(D_X^K)$, i.e. for any $0 \leq k \leq K$,

$$\mathcal{P}_k \mathcal{L}(A_X^K) = \mathcal{P}_k \mathcal{L}(D_X^K). \quad (17)$$

Proof: See Appendix B. \blacksquare

Back to the proof of Theorem 1. Let $\delta = d_{\mathcal{D},\infty}(D_X^K, D_Y^K)$. From Definitions 3, there exists a correspondence C between X and Y such that $|r_X^k(x_{0:k}) - r_Y^k(y_{0:k})| \leq \delta$ for any k and any pairs of correspondents $(x_{0:k}, y_{0:k}) = (x_0, y_0), \dots, (x_k, y_k) \in C$. Construct the pair of augmented networks A_X^K and A_Y^K defined in Definition 6 using this correspondence. It follows from (16) that given a tuple $a_{0:k} \in C$ with each node a_i representing the correspondent pair (x_{C_i}, y_{C_i}) in C ,

$$|r_{A_X}^k(a_{0:k}) - r_{A_Y}^k(a_{0:k})| = |r_X^k(x_{C_0:C_k}) - r_Y^k(y_{C_0:C_k})| \leq \delta. \quad (18)$$

Since A_X^K and A_Y^K have identical nodes, (18) implies that any simplices that appear at time α in the induced filtration $\mathcal{L}(A_X^K)$ will appear no earlier than $\alpha - \delta$ and no later than $\alpha + \delta$ in the induced filtration $\mathcal{L}(A_Y^K)$. Using the Interleaving Theorem, this yields the bound

$$b^k(A_X^K, A_Y^K) \leq \delta = d_{\mathcal{D},\infty}(D_X^K, D_Y^K). \quad (19)$$

From Fact 1, $b^k(A_X^K, A_Y^K) = b^k(D_X^K, D_Y^K)$. Substituting this into (19) concludes the proof. \blacksquare

When comparing high order networks via their induced persistence diagrams, Theorem 1 provides justification that the dissimilarity obtained from persistent homologies is a lower bound on their infinity norm network distance. In particular, Theorem 1 implies that: (i) A large difference in persistent homologies imply the networks to be highly different. (ii) Lower bounds can be used to estimate distance intervals because upper bounds are easy to determine using specific correspondences as per Definition 2.

We note that the differences between persistence diagrams do capture important differences between networks. E.g., in Figure 6 we can consider network (b) as a modification of network (a) in which we separate node a into the closely related nodes a and b that have a dissimilarity $r_X^1(a, b) = 0.1$. The persistence diagrams of these two networks are identical, which is consistent with the relative proximity of these two networks. Network (c) is more different, because the addition of the dissimilarity $r_X^1(b, d) = 0.7$ complicates the argument that nodes a and b can be well represented by a single node as in network (a). Network (d) can be argued to be in between networks (b) and (c) since the dissimilarity $r_X^1(a, b, d) = 0.8$ implies a sense of added proximity between nodes a , b , and d . The difference between the persistence diagrams of networks (c) and (b) is indeed larger than the difference between the diagrams of networks (d) and (b).

Further evidence for the usefulness of the bound in (15) follows from the fact that the bound is tight in some cases. An example of a tight bound is shown in Figure 8. The optimal correspondence C , shown in the figure, yields a network distance $d_{\mathcal{D},\infty}(D_X^1, D_Y^1) = 0.1$. The coordinates

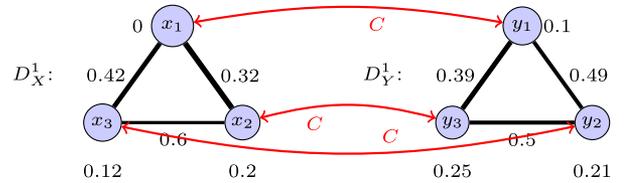


Fig. 8. An example where the bottleneck distance between the k -th dimensional persistence diagrams of the filtrations $\mathcal{L}(D_X^K)$ and $\mathcal{L}(D_Y^K)$ is the same as their ∞ -norm network distance for $k \in \{0, 1\}$. The optimal correspondence C yields $d_{\mathcal{D},\infty}(D_X^1, D_Y^1) = 0.1$. The points in the 0-dimension persistence diagrams for $\mathcal{L}(D_X^1)$ are with coordinates $(0, \infty)$, $(0.12, 0.42)$, $(0.2, 0.32)$ and for $\mathcal{L}(D_Y^1)$ are $(0.1, \infty)$, $(0.21, 0.51)$, $(0.25, 0.39)$. The point in the 1-dimension persistence diagrams for $\mathcal{L}(D_X^1)$ is with coordinate $(0.6, 1)$ and for $\mathcal{L}(D_Y^1)$ is $(0.5, 1)$. The bottleneck distances between the 0- as well as the 1- dimensional persistence diagrams of the filtrations induced from the two networks are 0.1.

of the points in the 0-dimension persistence diagrams for $\mathcal{L}(D_X^1)$ are $(0, \infty)$, $(0.12, 0.42)$, $(0.2, 0.32)$ and for $\mathcal{L}(D_Y^1)$ are $(0.1, \infty)$, $(0.21, 0.51)$, $(0.25, 0.39)$. The coordinates of the point in the 1-dimension persistence diagrams for $\mathcal{L}(D_X^1)$ are $(0.6, 1)$ and for $\mathcal{L}(D_Y^1)$ is $(0.5, 1)$. The bottleneck distances between the 0- and 1-dimensional persistence diagrams of the induced filtrations are 0.1; same as $d_{\mathcal{D},\infty}(D_X^K, D_Y^K)$.

A. Persistence Bounds on k -order Distances

Theorem 1 guarantees that the ∞ -norm dissimilarity network distance $d_{\mathcal{D},\infty}$ can be tightly bounded by persistence diagrams. In this section we build on Theorem 1 to show that the k -order network distance $d_{\mathcal{D}}^k$ can also be bounded by persistent methods.

Theorem 2: Given two dissimilarity networks D_X^K and D_Y^K and an integer $1 \leq k \leq K$, the bottleneck distance between the k' -th dimensional persistence diagrams of the filtrations $\mathcal{L}(D_X^K)$ and $\mathcal{L}(D_Y^K)$ is at most $d_{\mathcal{D}}^k(D_X^K, D_Y^K)$ for any $0 \leq k' < k$, i.e.

$$b^{k'}(D_X^K, D_Y^K) \leq d_{\mathcal{D}}^k(D_X^K, D_Y^K). \quad (20)$$

Proof: We prove Theorem 2 from Theorem 1. To do that, we leverage the relationship for network distances as we present next.

Fact 2: Given high order networks N_X^K and N_Y^K , we have

$$d_{\mathcal{N}}^{k'}(N_X^K, N_Y^K) \leq d_{\mathcal{N}}^k(N_X^K, N_Y^K), \quad (21)$$

for $0 \leq k' \leq k - 1 \leq K$, and

$$d_{\mathcal{N},\infty}(N_X^K, N_Y^K) = d_{\mathcal{N}}^K(N_X^K, N_Y^K). \quad (22)$$

Proof: See Appendix B. \blacksquare

Fact 2 implies that $d_{\mathcal{N}}^k$ increases as the order k becomes higher, and the ∞ -norm distance $d_{\mathcal{N},\infty}$ is the same as the K -order network distance. The result in Theorem 2 only considers relationship functions of order k , and to leverage the connection as in Fact 2, we introduce the following notion of truncated networks.

Definition 7: Given K -order network $N_X^K = (X, r_X^0, \dots, r_X^K)$, its k -order truncated network N_X^k is defined as $N_X^k = (X, r_X^0, \dots, r_X^k)$.

A k -order truncated network N_X^k has the same node set as its parent network N_X^K and collects the lowest $k + 1$ order

relationship functions of N_X^K . Back to the proof of Theorem 2, construct the k -order truncated networks D_X^k and D_Y^k from D_X^K and D_Y^K . It follows directly from Theorem 1 that for any $0 \leq k' \leq k - 1$,

$$b^{k'}(D_X^k, D_Y^k) \leq d_{\mathcal{D}, \infty}(D_X^k, D_Y^k). \quad (23)$$

Since D_X^k and D_Y^k are valid high order networks, Fact 2 implies that $d_{\mathcal{D}, \infty}(D_X^k, D_Y^k) = d_{\mathcal{D}}^k(D_X^k, D_Y^k)$. Meanwhile, $d_{\mathcal{D}}^k(D_X^k, D_Y^k) = d_{\mathcal{D}}^k(D_X^K, D_Y^K)$ follows from the facts that D_X^k and D_X^K have identical k -order relationship functions and D_Y^k and D_Y^K have same k -order relationship. Finally, for any $k' \leq k - 1$, the k' -th dimensional persistence diagram of $\mathcal{L}(D_X^k)$ is identical to that of $\mathcal{L}(D_X^K)$ and the k' -th dimensional persistence diagram of $\mathcal{L}(D_Y^k)$ is identical to that of $\mathcal{L}(D_Y^K)$, therefore $b^{k'}(D_X^k, D_Y^k) = b^{k'}(D_X^K, D_Y^K)$. Combining these observations, the proof concludes. ■

Theorem 2 bounds the k -order network distance by the bottleneck distance between persistence diagrams of any order $k' < k$. We note that best lower bound is not necessarily $k' = k - 1$. Observe that the result in Theorem 2 does not apply for $k = 0$. This is not a problem because the 0-order network distances $d_{\mathcal{D}}^0$ examine relationships between individual nodes only. Thus, we can compute the optimal correspondence in Definition 2 by matching each of the relationships $r_X^0(x)$ in the network \mathcal{D}_X^K to the closest relationship $r_Y^0(y)$ in the network \mathcal{D}_Y^K . We emphasize that the lower bound described in Theorem 2 is also tight since we can find dissimilarity networks D_X^K and D_Y^K such that $d_{\mathcal{D}}^k(D_X^K, D_Y^K)$ equals the bottleneck distance between the k' -th dimensional persistence diagrams of the filtration $\mathcal{L}(D_X^k)$ and $\mathcal{L}(D_Y^k)$ for any $0 \leq k' \leq k - 1 \leq K$. See Figure 8 for an illustration where bottleneck distances between the k -dimensional persistence diagrams coincide with $d_{\mathcal{D}}^k$.

One of the key arguments in the proof of Theorem 2 is the fact that the infinity norm network distance $d_{\mathcal{N}, \infty}(N_X^K, N_Y^K)$ and the K -order network distance $d_{\mathcal{N}}^K(N_X^K, N_Y^K)$ coincide [cf. Fact 2]. This equality implies that Theorem 1 follows as a particular case of Theorem 2 by setting $k = K$ in (20) and observing that $d_{\mathcal{N}, \infty}(N_X^K, N_Y^K) = d_{\mathcal{N}}^K(N_X^K, N_Y^K)$. Do notice that this doesn't mean that the proof of Theorem 1 is redundant because Theorem 1 is leveraged in the proof of Theorem 2. However, the result does imply that the infinity norm distance $d_{\mathcal{N}, \infty}(N_X^K, N_Y^K)$ does not contain information beside the one that is contained in the collection of k -order distance $d_{\mathcal{N}}^k(N_X^K, N_Y^K)$ for $k = 0, \dots, K$.

Another consequence of Theorem 2 is that we can bound the p -norm network distance $d_{\mathcal{D}, p}(D_X^K, D_Y^K)$ [cf. Definition 3] for arbitrary p . This can be done using the individual bounds for $d_{\mathcal{D}}^k(D_X^K, D_Y^K)$ established in (20) as we show next.

Corollary 1: Group the k -order dissimilarity network distances in the vector $\mathbf{d}_{\mathcal{D}}^K(D_X^K, D_Y^K) := [d_{\mathcal{D}}^0(D_X^K, D_Y^K), d_{\mathcal{D}}^1(D_X^K, D_Y^K), \dots, d_{\mathcal{D}}^K(D_X^K, D_Y^K)]^T$. Further define the vector $\mathbf{b}^K(D_X^K, D_Y^K) := [d_{\mathcal{D}}^0(D_X^K, D_Y^K), b^{k_1}(D_X^K, D_Y^K), \dots, b^{k_K}(D_X^K, D_Y^K)]^T$ whose first component is the zero order distance $d_{\mathcal{D}}^0(D_X^K, D_Y^K)$ and whose other components are bottleneck lower bounds with $k_l < l$. The p -norm network distance can be

lower bounded as

$$\|\mathbf{b}^K(D_X^K, D_Y^K)\|_p \leq \|\mathbf{d}_{\mathcal{D}}^K(D_X^K, D_Y^K)\|_p \leq d_{\mathcal{D}, p}(D_X^K, D_Y^K). \quad (24)$$

Proof: The second inequality follows from the fact that a single correspondence is used in $d_{\mathcal{D}, p}$ whereas an order-specific correspondence may be utilized in each k -order distance $d_{\mathcal{D}}^k$ [21]. For the first inequality, $b^{k_l}(D_X^K, D_Y^K) \leq d_{\mathcal{D}}^{k_l}(D_X^K, D_Y^K)$ comes from Theorem (2). This implies that the vector $\mathbf{b}^K(D_X^K, D_Y^K)$ is element-wise smaller or equal to the vector $\mathbf{d}_{\mathcal{D}}^K(D_X^K, D_Y^K)$. ■

Corollary 1 shows that the p -norm network distance $d_{\mathcal{D}, p}$ can be lower bounded using: (i) The exact value of the 0-order network distance. (ii) The persistence homology lower bounds in (20) of Theorem 2. The bounds in Corollary 1 are expected to be the most useful when they combine information gleaned from persistence diagrams of different orders.

Remark 2: The requirement of having the inequality in (6) be strict unless all the elements of $x_{0:k}$ appear in $x_{0:k-1}$ for the network to be a dissimilarity network [cf. Definition 4] is not always naturally satisfied in practice. As we have seen, its satisfaction may necessitate the addition of a small but arbitrary constant to some tuple relationships [cf. Figure 1]. This is not a problem in practice because Theorems 1 and 2 still holds if this technical condition is violated. The only modification in the statements is that the respective network distances are not metrics but pseudometrics. What this means is that it is possible to construct pathological examples of networks that satisfy the order decreasing property *not* strictly and that have null network distance while *not* being isomorphic. These networks are unlikely to appear in practice and, even if they do, the bounds in Theorems 1 and 2 hold, albeit with a different interpretation.

V. PROXIMITY NETWORKS

In dissimilarity networks relationships $r_X^k(x_{0:k})$ encode a level of dissimilarity between elements of the tuple. The relationship $r_X^k(x_{0:k})$ may, alternatively, describe similarity or proximity between elements. To account for this, we define proximity networks as those in which the relationship in (6) is flipped to

$$r_X^k(x_{0:k}) \leq r_X^{k-1}(x_{0:k-1}), \quad (25)$$

with the inequality equalized if and only if the element x_k also appear in $x_{0:k-1}$. We refer to (25) as the order decreasing property and denote the set of all proximity networks of order K as \mathcal{P}^K . When the input networks of Definition 2 or 3 are proximity networks, we refer to the network distances as the k -order proximity network distance $d_{\mathcal{P}}^k$ or the p -norm proximity network distances $d_{\mathcal{P}, p}$. The restrictions to proximities make the distances well-defined *metrics* in the respective spaces [21].

As is the case with the order increasing property in dissimilarity networks, the order decreasing property is well justified in proximity networks. E.g., consider a different aspect of coauthorship in which the k -order relationship function between authors is proportional to the number of publications coauthored by all members of a given $(k + 1)$ -tuple. In particular, the zeroth

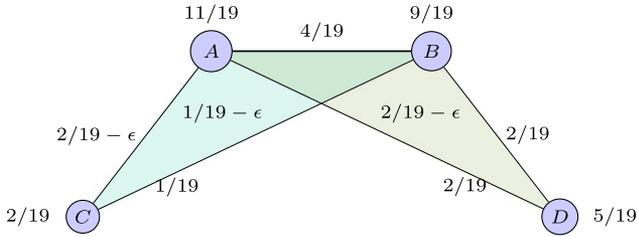


Fig. 9. Collaborations between authors in a research community. The k -order relationship in this 2-order network describes the number of publications between members of a $(k + 1)$ -tuple normalized by the total number of papers. E.g., A writes 11 papers in total, and coauthors 4, 2, and 2 papers respectively with B , D , and C . The ϵ in the relationship function is a technical modification to differentiate relationship between two authors from that of a single author. Author A also writes 1 paper jointly with B and C .

order proximities r_X^0 are the numbers of papers published by authors normalized by the total number of papers. In Figure 9 authors A, B, C, D publish 11, 9, 2, and 5 papers respectively and there are 19 papers in total which implies $r_X^0(A) = 11/19$, $r_X^0(B) = 9/19$, $r_X^0(C) = 2/19$, $r_X^0(D) = 5/19$. The first order proximities r_X^1 represent the number of papers co-published by nodes. Since a collaboration for a pair of authors is also a paper for each of the individuals it is certain that $r_X^1(x, x') \leq r_X^0(x)$. In Figure 9, A and B collaborate on 4 papers, which is less than the 11 and 9 papers written by each of the individuals. Authors A and C as well as A and D coauthor 2 papers in total. Second order proximities r_X^2 for triplets indicate the normalized number of papers coauthored by the three members. Since a paper with three authors is also a collaboration for the three pairs of authors we must have $r_X^2(x, x', x'') \leq r_X^1(x, x')$. In Figure 9, authors A, B , and D cowrite 2 papers, which is no more than the number of pairwise collaborations between each pair of the authors. An ϵ constant is subtracted from some proximities to satisfy the technical condition that $r_X^k(x_{0:k}) = r_X^{k-1}(x_{0:k-1})$ if and only the elements in $x_{0:k}$ are all elements of $x_{0:k-1}$. This is not a concern in practice either – see Remark 2. Undefined hyperedge relationships in proximity networks take the value 0 which is the minimum possible proximity for a tuple.

For any proximity network P_X^K with relationships $p_X^k(x_{0:k})$ we can define the dual network D_X^K with relationship functions $d_X^k(x_{0:k}) := 1 - p_X^k(x_{0:k})$. It is not difficult to see that D_X^K is a well defined dissimilarity network and any of the network distances between the dual dissimilarity networks coincides with the respective distance between the proximity networks [21]. It follows that the bounds in Theorems 1 and 2 apply for proximity networks with the caveat that the filtrations $\mathcal{P}_k \mathcal{L}(D_X^K)$ and $\mathcal{P}_k \mathcal{L}(D_Y^K)$ are filtrations of the dual networks associated with the given proximity networks P_X^K and P_Y^K .

VI. IMPLEMENTATION DETAILS

We discuss implementation issues regarding the procedure for evaluating lower bounds, its computational cost, and some heuristic simplifications.

Procedure: To compute distance lower bounds for a dissimilarity network we construct filtrations and evaluate persistence diagrams. Algorithmic procedures to compute filtrations and

persistence diagrams are discussed in [31], [37]. The bottleneck distances $b^k(D_X^K, D_Y^K)$ between the persistence diagrams of different orders $0 \leq k \leq K$ are then evaluated. These distances follow from Definition 5 and are found as solutions of LBAP with costs given by (12). If the interest is in the infinity norm network distances $d_{\mathcal{D}, \infty}$, any of the bottlenecks distances is a lower bound [cf. Theorem 1]. If we are interested in the k -order network distance $d_{\mathcal{D}}^k$, the bottleneck distance between persistence diagrams of dimension $k' \leq k - 1$ is a lower bound [cf. Theorem 2]. When given a proximity network we construct the dual network with relationships $r_X^k(x_{0:k}) = 1 - r_X^k(x_{0:k})$ and proceed as before.

Computational cost: The process above involves computation of the persistent homology and evaluation of the bottleneck distance. Given a K -order network with n nodes, denote as \tilde{n} the total number of simplices and m as the number of points in the persistence diagram. The complexity of computing a persistence diagram is akin to matrix multiplication of matrices with size \tilde{n} , i.e. the number of simplices considered [31]. Because the complexity for matrix completion is $O(\tilde{n}^{2.373})$ [37], computational complexity of evaluating a persistence diagram is also $O(\tilde{n}^{2.373})$. The complexity of evaluating the bottleneck distance using LBAP [cf. (14)] is of order $O(m^{2.5} \log m)$ [34, Algorithm 6.1]. The number of simplices of order k can be as large n^k and the number of points in a persistence diagram as large as the number of simplices. This means that we can have $m = \tilde{n} = n^K$, which yields impractical computational times for orders larger than $K = 2$. In the practical examples we consider in Sections VII-B and VII-C, the number of simplices \tilde{n} is much smaller than the maximum possible and scales linearly with the number of nodes. This yields complexity $O(n^{2.373})$ for computing the persistence diagram and $O(n^{2.5} \log n)$ for the bottleneck distance.

Undefined tuple relationships: Formally, dissimilarity networks require that all undefined tuple relationships take value 1. As defined above this is impractical because it would scale the total number of simplices as $\tilde{n} = n^k$. In computations we leave these simplices undefined and let features die at infinity. As it follows from Remark 1, the death time of these features can be set to $\alpha = 1$ to yield the same result that would be obtained from setting undefined tuple relationships to 1 before computing the persistence diagram.

Elimination of small homological features: We can reduce the number of points in a persistence diagram by removing points \mathbf{q} close to the diagonal. This is justified because: (i) These points represent ephemeral homological features and are likely to be generated by noise in observations. (ii) They contribute a small value to the cost in (12) and won't improve the lower bound much from the trivial nonnegative lower bounds $0 \leq d_{\mathcal{D}, \infty}(D_X^K, D_Y^K)$ or $0 \leq d_{\mathcal{D}}^k(D_X^K, D_Y^K)$ that hold for any distance.

VII. APPLICATIONS

We illustrate the usefulness of homology methods through experiments in both synthetic (Section VII-A) and real-world data (Sections VII-B and VII-C). The objective is to demonstrate

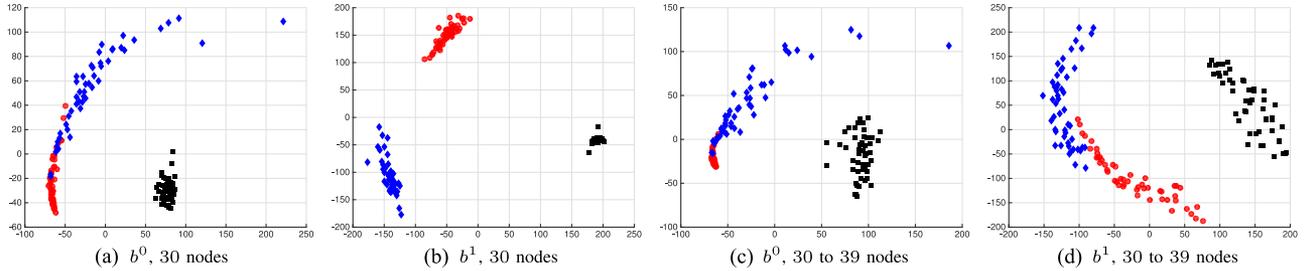


Fig. 10. Two dimensional Euclidean embeddings of the networks constructed from three different models with different number of nodes with respect to the network metric lower bounds b^0 and b^1 . In the embeddings, red circles denote networks constructed from the Erdős-Rényi model, blue diamonds represent networks constructed from the unit circle model, and black squares the networks from the correlation model.

that networks with similar structures and should be alike to each other are indeed similar in their persistent homologies.

A. Classification of Synthetic Networks

We consider three types of synthetic weighted pairwise networks. Edge weights in all three types of networks encode proximities. The first type of networks are with weighted Erdős-Rényi model [38], where the edge weight between any pair of nodes is a random number uniformly selected from the unit interval $[0, 1]$. In the second type of networks, the coordinates of the vertices are generated uniformly and randomly in the unit circle, and the edge weights are evaluated with the Gaussian radial basis function $\exp(-d(i, j)^2/2\sigma^2)$ where $d(i, j)$ is the distance between vertices i and j in the unit circle and σ is a kernel width parameter. In all simulations, we set σ to 0.5. The edge weight measures the proximity between the pair of vertices and takes value in the unit interval. In the third type of networks, we consider that each vertex i represents an underlying feature $\mathbf{u}_i \in \mathbb{R}^d$ of dimension d , and examine the Pearson's linear correlation coefficient ρ_{ij} between the corresponding features \mathbf{u}_i and \mathbf{u}_j for a given pair of nodes i and j . The weight for the edge connecting the pair is then set as $\rho_{ij}/2 + 0.5$, a proximity measure in the unit interval. The feature space dimension d is set as 30 in all simulations.

We start with networks of equal size $|X| = 30$ and construct 50 random networks for each aforementioned type. The edges with weights no greater than a threshold τ are removed to create sparsity. We set $\tau = 0.2$ in all simulations. In order to transform the constructed pairwise networks into high order proximity networks, the 0-order proximity for any node is set to 0, i.e. $r_X^0(x) := 0$ for any $x \in X$. We then use the persistent homology method described in Section VI to evaluate the dissimilarities between all networks. The persistent homologies are computed using JavaPlex [39]. Figure 10(a) and (b) plot the two dimensional Euclidean embeddings [40] of the network metric lower bounds b^0 and b^1 between the 0- and 1-dimensional persistence diagrams respectively. All embeddings in the paper are constructed with respect to minimizing the sum of squares of the interpoint distances; the other common choice to minimize the sum of four power of the internist distances yields similar results. Networks constructed with different models form clear separate clusters with respect to b^1 where networks with Erdős-Rényi model are denoted by red circles, networks with unit

circle model are described by blue diamonds, and correlation model represented as black squares. The clustering structure is not that clear in terms of b^0 but the dissimilarities between networks constructed from different models are in general much higher than that between networks from the same model, and an unsupervised classification with one linear boundary in the embedded space would yield 10 out of 150 errors (6.7%).

Next we consider networks with number of nodes ranging between 30 and 39. Ten networks are randomly generated for each network type and each number of nodes, resulting in 150 networks in total. Figure 10(c) and (d) illustrate the two dimensional Euclidean embeddings of the network metric lower bounds b^0 and b^1 for these networks. Despite the fact that networks with same model have different number of nodes, dissimilarities between persistent homologies are smaller when their underlying networks are from the same process. Besides, networks with the correlation model are highly similar to each other regarding their corresponding persistent homologies, irrespective of the sizes of the networks nor the dimension of persistence diagrams. An unsupervised classification with one linear boundary would yield 6 out of 150 errors (4.0%) for b^0 and 7 errors (4.7%) for b^1 . The performance of b^0 stays relatively unchanged when the number of nodes in the networks considered reside in a range of nodes. The simulations illustrate the applicability of persistent homology in identifying the patterns of networks of different processes. Similar results are obtained when different parameters are used in generating the networks.

Here we give interpretations of why persistent homologies succeed in network discrimination. Since only pairwise relationships are examined, the bottleneck distance b^0 is determined by the death time of the 0-dimensional persistent homologies. If we consider that nodes connected by an edge of weight w become members of the same community at time w , the death time in $\mathcal{P}_0\mathcal{L}$ can be interpreted as the time instant when isolated nodes join the main community of the network. We focus our attention on the 150 networks with size ranging between 30 and 39. Each row in Figure 11(a) represents a network, and plots the death time in the 0-dimensional persistence diagram $\mathcal{P}_0\mathcal{L}$ of the filtration induced from that network in ascending order from left to right. The top 50 networks represent those with Erdős-Rényi model, the middle 50 are with the unit circle model, and the bottom 50 are with the correlation model. Each row in Figure 11(b) plots the death time in $\mathcal{P}_0\mathcal{L}$, in descending order from left to right, of the induced filtration. It is not a direct

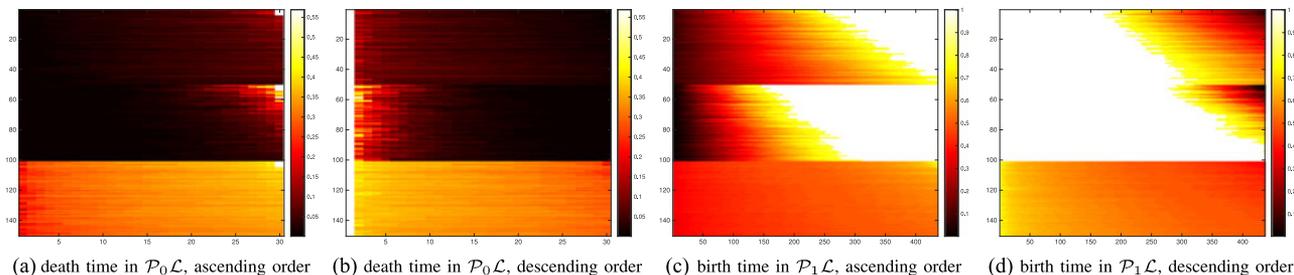


Fig. 11. Illustrations of the set of death time in 0-dimensional persistence diagrams $\mathcal{P}_0 \mathcal{L}$ and the set of birth time in 1-dimensional persistence diagrams $\mathcal{P}_1 \mathcal{L}$ for the filtrations constructed from the randomly generated networks with number of nodes ranging between 30 and 39. The top 50 networks represent those constructed from the Erdős-Rényi model, the middle 50 networks are with the unit circle model, and the bottom 50 are with the correlation model.

mirror of Figure 11(a) because the size of the networks ranges from 30 to 39 and so the number of points in the persistence diagrams are different. The time instants when isolated nodes join the main community in the networks with the correlation model are concentrated in the interval $[0.25, 0.4]$. This is due to the fact that the linear correlation coefficient between two randomly generated feature vectors cannot be too positive nor too negative. Also, networks with unit circle model are different from those constructed with the Erdős-Rényi model because in the latter case, the distribution of death time of points in $\mathcal{P}_0 \mathcal{L}$ has heavy tail towards 1, which results from the fact that some points in the unit circle may be far away from the main component and it requires larger distance for them to join the main component.

For higher order persistent homologies, b^1 is only affected by the birth time of the 1-dimensional persistent homologies, which can be interpreted as the maximum of the three pairwise connections between three nodes and therefore the time instant when a ‘closely-connected’ community is formed by the three nodes. Each row in Figure 11(c) and (d) exhibits the birth time in $\mathcal{P}_1 \mathcal{L}$, in descending and ascending order from left to right respectively, of the filtration induced from the corresponding network. The time instants when three nodes in the networks with correlation model form a closely connected triplet are highly focused in the interval of $[0.4, 0.6]$. Recall that edge weight is defined as $\rho_{ij}/2 + 0.5$, therefore a edge weight in the interval $[0.4, 0.6]$ implies very weak correlation ($|\rho_{ij}| \leq 0.1$). This is due to the fact that when the correlation coefficient ρ_{ij} between features \mathbf{u}_i and \mathbf{u}_j are small and ρ_{ik} between \mathbf{u}_i and \mathbf{u}_k is small, ρ_{jk} between \mathbf{u}_j and \mathbf{u}_k also tends to be small. For networks with unit circle model, the time points for many nodes in the network forming a closely connected triplet are either smaller or larger compared to the Erdős-Rényi model. This is because if three points in the unit circle are close to each other, all the pairwise distances between them would not be high; otherwise the minimum inscribed circle of the three points would possess large radius, resulting in a high value on the maximum of the pairwise distances. Admittedly, other methods to compare networks may also succeed in distinguishing networks, after some proper treatment towards the issue of different sizes. Nonetheless, persistent homology method would be more universal, not only for the reason that it establishes a lower bound to the actual network metrics, but also since it provides a systematic way to analyze the formation of communities in a given network.

B. Comparison of Coauthorship Networks

We apply the lower bounds to compare 2-order coauthorship networks where relationship functions denote the number of publications of single authors, pairs of authors, and triplets. These coauthorship networks are proximity networks because they satisfy the order decreasing property. We consider publications in 5 journals from mathematics community: Computational Geometry (CG), Discrete Computational Geometry (DCG), J. of Applied Probability, (JAP) J. of Mathematical Analysis and Applications (JMAA), SIAM J. on Numerical Analysis (SJNA), and 6 journals from engineering community, all from IEEE: Signal Processing Magazine (SPM), Trans. Automatic Control (TAC), Trans. Pattern Analysis and Machine Intelligence (TPAMI), Trans. Information Theory, Trans. Signal Processing (TSP), Trans. Wireless Communication (TWC). For each journal, we construct networks for the 2004-2008 and 2009-2013 quinquennia. For TAC, TSP, and TWC, we also construct networks for each annual from 2004 to 2013. Lists of publications are queried from [41].

For each of these journals we consider all publications in the period of interest and construct proximity networks where the node set X is formed by all authors of the publications. Zeroth order proximities are defined as the total number of publications of each member of the network, first order proximities as the number of papers coauthored by pairs, and second order proximities as the number of papers coauthored by triplets. To make networks with different numbers of papers comparable we normalize all relationships by the total number of papers in the network. The positive constant ϵ as in Figure 9 is set to $1/1000$. There are papers with more than three coauthors but we don’t record proximities of order higher than 2. By assuming that networks from the same community or constructed from the same journal have similar collaboration patterns, we show here that network metric lower bounds succeed in identifying these patterns and in distinguishing coauthorship networks from communities with different interests.

Figure 12 shows the two dimensional Euclidean embeddings of the network metric lower bounds b^0 , b^1 and b^2 . The 12 engineering networks (blue diamonds) separate clearly from the 10 mathematics networks (red circles) in b^1 and b^2 . The clustering is not that clear in b^0 but still networks from same community tend to be similar to each other. An unsupervised classification with one linear boundary running across the embeddings would

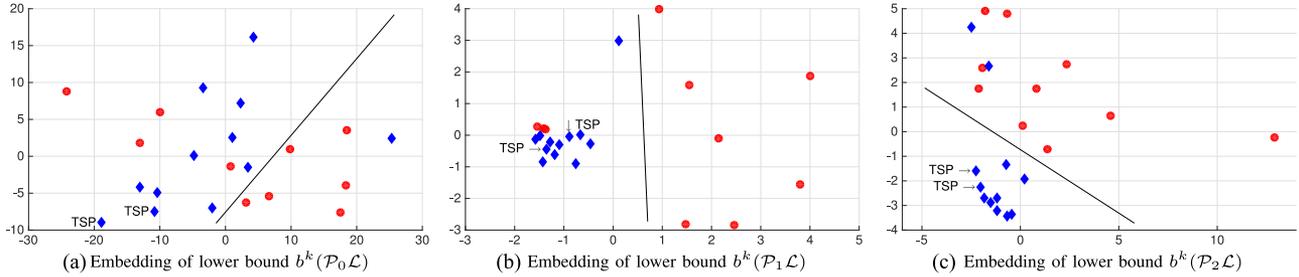


Fig. 12. Two dimensional Euclidean embeddings of the networks constructed from quinquennial publications in engineering and mathematics journals with respect to the network metric lower bounds $b^k(\mathcal{P}_0\mathcal{L})$, $b^k(\mathcal{P}_1\mathcal{L})$, and $b^k(\mathcal{P}_2\mathcal{L})$. In the embeddings, red circles denote networks constructed from mathematics journals and blue diamonds represent networks from engineering journals. Networks constructed from publications of TSP are labeled.

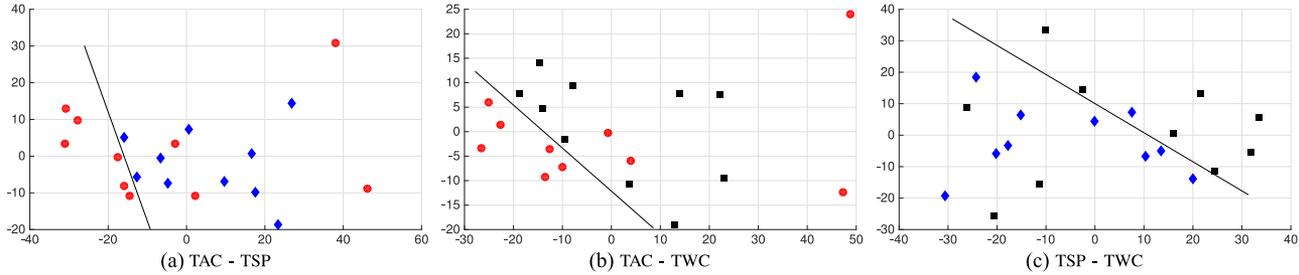


Fig. 13. Two dimensional Euclidean embeddings of the networks constructed from annual publications in TAC, TSP, and TWC with respect to the summation of the metric lower bounds $b^k(\mathcal{P}_0\mathcal{L})$, $b^k(\mathcal{P}_1\mathcal{L})$, and $b^k(\mathcal{P}_2\mathcal{L})$. In the embeddings, red circles represent TAC, blue diamonds TSP, and black squares TWC.

generate errors of 2 (9.09%) to 5 (22.73%) out of 22 networks. Networks constructed from the same journal tend to be close in the lower bounds. As an example, the networks of TSP with different quinquennia are marked in the embeddings and it is clear that their differences in homologies are considerably low. Such scenarios are observed for several other journals as well.

We analyze the persistent homology of each of the coauthorship networks to investigate the reason why persistent homologies succeed in network discrimination. Compared to networks from mathematics communities, networks from engineering communities in general would yield 0-dimensional persistent homologies with smaller birth time but larger death time, 1-dimensional homologies with larger birth and death time, and 2-dimensional homologies with larger birth time. An interpretation of such observations would be that in engineering, there exist more small communities that never collaborate with each other and it is uncommon to have a “club” of 3 to 5 authors in engineering that a strong collaboration exists between any pairs of the authors in the “club”; such scenarios are absent for mathematicians.

As a comparison, we applied some simple and reasonable methods to compare the coauthorship networks considered in this section. Motifs have been shown effective in distinguishing coauthorship networks from different scientific fields [16]. To compare high order coauthorship networks by motifs, we restrict attention to pairwise relationships. The dissimilarities between coauthorship networks are assigned as the differences between the summations of the weighted motifs in their corresponding pairwise networks. Analysis based on triangle motifs (weighted) results in a clear cluster between networks from CG, DCG, JAP and another cluster between networks from TSP, TWC, but cannot distinguish other networks very well.

Tetrahedron motif analysis (weighted) results in three clear clusters: networks from CG, DCG, JAP, networks from JMAA, SJNA, SPM, TPAMI, and networks from TSP, TWC. Other simple and common methods to compare pairwise networks yield similar results. Methods to compare networks via features give us similar observations as those based on the persistent methods; feature methods would generate 6 to 8 errors in classifications.

C. Engineering Communities with Different Research Interests

The network metric lower bounds succeed in distinguishing the different collaboration patterns in engineering and mathematics communities. We now illustrate that the lower bounds are also able to identify distinctive features of engineering communities with different research interests. To see this we consider the networks constructed from annual publications of TAC, TSP, and TWC.

Figure 13 shows the two dimensional Euclidean embeddings of the networks with respect to the summation of the lower bounds b^0 , b^1 , and b^2 . We expect more variations in annual networks because the time for averaging behavior is reduced. Besides, it is hard to argue that intrinsic and obvious differences exist in the collaboration patterns in automatic control, signal processing, and wireless communication communities. Still, networks constructed from the same journal but different annuals tend to be close to each other and form clustering structures. An unsupervised classification with one linear boundary in the embeddings run across the summation of lower bounds would generate 4 (20%) errors out of 20 networks in all three classification problems considered. The less obvious clustering structure formed by networks from different journals in Figure 13(c) compared to (a) and (b) also suggests that the

collaboration patterns in research communities of signal processing and wireless communication are more similar compared to that of automatic control.

VIII. CONCLUSION

We establish connections between high order networks and simplicial complexes and use the differences between the induced homological features to evaluate the differences between networks. We justify that this is a lower bound to two families of valid metrics in the space of high order networks modulo permutation isomorphisms. These lower bounds succeed in classifying weighted pairwise networks constructed from different processes, in distinguishing the collaboration patterns of engineering communities from mathematics communities, and in discriminating engineering communities with different research interests. Since the upper bounds of network distances are easy to determine using specific correspondences (see Definition 2), this together with persistence lower bounds would give us distance intervals on the network distances, which can be used in clustering of networks [42], [43].

APPENDIX A

PROOFS OF PROPOSITION 2

We refer authors to [32] for more formal notions on computational homology. Given any tuple $x_{0:k}$ with non-duplicating nodes, (10) indicates that the k -simplex ϕ^k defined by the convex hull $\text{conv}\{x_{0:k}\}$ appears strictly after any of its faces $\text{conv}\{x_{0:\widehat{s}:k}\}$ in the filtration. Suppose ϕ^k appears at time α and denote $\partial_k \phi^k = \sum_i \beta_i \psi_i^{k-1}$ with β_i the coefficients, then each ψ_i^{k-1} appears strictly before time α .

Now suppose that the appearance of ϕ^k trivializes a $(k-1)$ -th dimensional homological feature. This means that ϕ^k is the boundary to trivialize the $(k-1)$ -th dimensional cycle $\partial_k \phi^k$. Since each face ψ_i^{k-1} of ϕ^k appears strictly before time α , the cycle $\partial_k \phi^k$ results in a homological feature. The death time of this homological feature is α , or equivalently, the time represented by the relationship $r_X^k(x_{0:k})$.

On the other hand, if the appearance of ϕ^k does not trivialize a $(k-1)$ -th dimensional homological feature, then the $(k-1)$ -cycle $\partial_k \phi^k$ is in the collection of simplices appearing before or on time α . This means that $\partial_k \phi^k$ can be represented by a sum of the boundaries of some k -chains Φ_i^k ,

$$\partial_k \phi^k = \sum_i \beta_i \partial_k \Phi_i^k \quad (26)$$

with coefficients β_i and k -chains Φ_i^k appearing before or on time α . By the definition of k -chains, $\Phi_i = \sum_j \beta'_j \psi_j^k$ with coefficients β'_j and k -simplices ψ_j^k appears before or on time α . Therefore, (26) can be written as $\partial_k \phi^k = \sum_j \beta''_j \partial_k \psi_j^k$. Rearranging terms,

$$\partial_k \left(\sum_j \beta''_j \psi_j^k - \phi^k \right) = 0. \quad (27)$$

This implies that $\sum_i \beta''_i \psi_i^k - \phi^k$ is a k -cycle. There must be a new cycle formed since ϕ^k just appears. The cycle cannot

be trivialized immediately since any $(k+1)$ -chain Ψ^{k+1} with $\partial_{k+1} \Psi^{k+1} = \sum_i \beta''_i \psi_i^k - \phi^k$ would involve a simplex $[x_{0:k,l}]$ for some node x_l with tuple $x_{0:k,l}$ consisted of non-repeating elements where this simplex $[x_{0:k,l}]$ appears strictly after α . Therefore we have a k -th dimensional homological feature with birth time α , or equivalently, the time denoted by the relationship $r_X^k(x_{0:k})$. This concludes the proof.

APPENDIX B

PROOFS OF FACT 1 AND FACT 2

Proof of Fact 1: We need to prove the (i) symmetry property and (ii) identity property of high order networks, (iii) $\mathcal{L}(A_X^K)$ being a valid filtration, and (iv) $\mathcal{L}(A_X^K)$ and $\mathcal{L}(D_X^K)$ having identical persistent homology.

Proof of the symmetry property: For any tuples $a_{0:k}$, the symmetry property of dissimilarity network D_X^K and the definition of augmented networks imply $r_{A_X}^k(a_{[0:k]}) = r_X^k(x_{[C_0:C_k]}) = r_X^k(x_{C_0:C_k}) = r_{A_X}^k(a_{[0:k]})$ for any reordering $a_{[0:k]}$. This shows the symmetry property of A_X^K . ■

Proof of the identity property: Given a tuple $a_{0:k}$, if its subtuple $a_{l_0:l_{\bar{k}}}$ have same set of unique elements as that of $a_{0:k}$, according to the Definition of augmented networks, we would have that $x_{C_0:C_k}$ and $x_{C_{l_0}:C_{l_{\bar{k}}}}$ also possess identical set of unique elements. The identity property of dissimilarity network D_X^K and the definition of augmented networks yield $r_{A_X}^k(a_{[0:k]}) = r_X^k(x_{C_0:C_k}) = r_X^k(x_{C_{l_0}:C_{l_{\bar{k}}}}) = r_{A_X}^k(a_{[l_0:l_{\bar{k}]})$. This shows the identity property of A_X^K . ■

Proof of $\mathcal{L}(A_X^K)$ being a valid filtration: The order increasing property of relationship function in augmented networks holds true due to the fact that

$$r_A^k(a_{[0:k]}) = r_X^k(x_{C_0:C_k}) \geq r_X^{k-1}(x_{C_0:C_{k-1}}) = r_A^{k-1}(a_{[0:k-1]}). \quad (28)$$

The remaining proof is identical to the proof of Proposition 1. ■

Proof of $\mathcal{L}(A_X^K)$ and $\mathcal{L}(D_X^K)$ having identical persistence diagrams: First, for each point $x \in X$, pick one pair (x, y) from the correspondence C to construct C_0 that is a subset of C . If we define a map that maps each pair (x, y) to its first element x , this gives a bijective projection from C_0 to X . Construct $A_{C_0,X}^K$ and $A_{C,X}^K$ as the augmented networks using the respective correspondence. It then follows naturally that $A_{C_0,X}^K$ is isomorphic to the original dissimilarity network D_X^K and so do the corresponding filtrations. Denote $\mathcal{L} := \mathcal{L}(A_{C,X}^K)$ and $\mathcal{L}_0 := \mathcal{L}(A_{C_0,X}^K) = \mathcal{L}(D_X^K)$.

Next, consider the projection of the filtration \mathcal{L} onto \mathcal{L}_0 , where each vertex (x, y) in C is mapped to (x, y_0) in C_0 who share the first element in the pair. Denote L^α and L_0^α as the simplicial complexes that collect simplices appearing before or prior to α in the respective filtration. For each value of α , the projection defines a retraction of simplicial complexes $L^\alpha \rightarrow L_0^\alpha$. Since the relationship functions are the same in the augmented network $A_{C,X}^K$ as in the original network $A_{C_0,X}^K$, whenever a simplex σ appears in L^α , not only its projection appears in L_0^α at the same time, but all the simplices connecting vertices of σ

and vertices of its projection are also already in L_0^α . Hence, as a simplicial map, the projection $L^\alpha \rightarrow L_0^\alpha$ is contiguous to the identity of L^α . As a result, the section $L_0^\alpha \rightarrow L^\alpha$ is a homotopy equivalence. Its induced homomorphism at the homology level is therefore an isomorphism, and consequently the two persistence diagrams are isomorphic. ■

Having demonstrated all statements, the proof completes. ■

Proof of Fact 2: We need to prove the two statements in (21) and (22). Given any $0 \leq k' \leq k \leq K$, it follows from (2) that for any correspondence C between X and Y ,

$$\Gamma_{X,Y}^k(C) = \max_{(x_{0:k}, y_{0:k}) \in C} |r_X^k(x_{0:k}) - r_Y^k(y_{0:k})|, \quad (29)$$

$$\Gamma_{X,Y}^{k'}(C) = \max_{(x_{0:k'}, y_{0:k'}) \in C} |r_X^{k'}(x_{0:k'}) - r_Y^{k'}(y_{0:k'})|. \quad (30)$$

For the $(x_{0:k'}, y_{0:k'}) \in C$ achieving the maximum difference $|r_X^{k'}(x_{0:k'}) - r_Y^{k'}(y_{0:k'})|$ in $\Gamma_{X,Y}^{k'}(C)$, we can construct another correspondent pair $(x_{0:k}^*, y_{0:k}^*)$ such that $x_{0:k}^*$ is a sub-tuple of $x_{0:k}^*$ with identical set of unique elements and $y_{0:k}^*$ is a sub-tuple of $y_{0:k}^*$ with identical set of unique elements. It follows from the identity property of high order networks that $r_X^{k'}(x_{0:k}^*) = r_X^k(x_{0:k}^*)$ and $r_Y^{k'}(y_{0:k}^*) = r_Y^k(y_{0:k}^*)$. This implies that taking the maximum $|r_X^k(x_{0:k}) - r_Y^k(y_{0:k})|$ over $(x_{0:k}, y_{0:k}) \in C$ cannot yield a lower difference than $|r_X^{k'}(x_{0:k}^*) - r_Y^{k'}(y_{0:k}^*)|$, i.e.

$$\Gamma_{X,Y}^k(C) \geq |r_X^{k'}(x_{0:k}^*) - r_Y^{k'}(y_{0:k}^*)| = \Gamma_{X,Y}^{k'}(C). \quad (31)$$

Since (31) holds true for any correspondence $C \in \mathcal{C}(X, Y)$, the inequality must hold true when we take the minimum over all correspondences $\mathcal{C}(X, Y)$ between X and Y ,

$$\min_{C \in \mathcal{C}(X, Y)} \Gamma_{X,Y}^k(C) \geq \min_{C' \in \mathcal{C}(X, Y)} \Gamma_{X,Y}^{k'}(C'). \quad (32)$$

Substituting the definition of k -order and k' -order network distances into (32) yields $d_D^k \geq d_D^{k'}$, concluding the proof of (21).

Also, it follows from (4) that for any correspondence C between the node sets X and Y , the network difference between N_X^K and N_Y^K measured by C is

$$\|\Gamma_{X,Y}^K(C)\|_\infty = \max_{k=0,1,\dots,K} \{\Gamma_{X,Y}^k(C)\}. \quad (33)$$

From (31) we know that for any k' , $\Gamma_{X,Y}^k(C) \geq \Gamma_{X,Y}^{k'}(C)$, and therefore for any correspondence C , $\max_{k=0}^K \{\Gamma_{X,Y}^k(C)\} = \Gamma_{X,Y}^K(C)$. Substituting this into (33) and taking a minimum over all correspondences concludes the proof of (22).

Having proven the two statements, the proof completes. ■

REFERENCES

- [1] W. Huang and A. Ribeiro, "Persistent homology approximations to network distances," in *Proc. Global Conf. Signal Inf. Process.*, Orlando, FL, USA, 2015, pp. 1002–1006.
- [2] W. Huang and A. Ribeiro, "Persistent homology lower bounds on network distances," in *Proc. 2016 IEEE Int. Conf. Acoust., Speech Signal Process.*, Shanghai, China, 2016, pp. 4845–4849.
- [3] W. Huang and A. Ribeiro, "Persistent homology lower bounds on distances 1271 in the space of networks," in *Proc. Asilomar Conf. Signals Syst. Comput.*, Pacific Grove CA, USA, Nov. 6–9, 2016.
- [4] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993, pp. 31–38.
- [5] S. Wasserman and K. Faust, *Social Network Analysis: Methods And Applications* (ser. Structural Analysis in the Social Sciences). Cambridge, U.K.: Cambridge Univ. Press, 1994.
- [6] R. Ghrist and A. Muhammad, "Coverage and hole-detection in sensor networks via homology," in *Proc. Int. Symp. Inf. Process. Sensor Netw.*, 2005, pp. 254–260.
- [7] H. Chintakunta and H. Krim, "Divide and Conquer: Localizing Coverage Holes in Sensor Networks," in *Proc. 2010 7th Annu. IEEE Commun. Soc. Conf. Sensor, Mesh Ad Hoc Commun. Netw.*, 2010, pp. 1–8.
- [8] B. Zhang, "Hypernetworks: A molecular evolutionary architecture for cognitive learning and memory," *IEEE Comput. Intell. Mag.*, vol. 3, no. 3, pp. 49–63, Aug. 2008.
- [9] W. Ren *et al.*, "Broadcasting in multi-radio multi-channel wireless networks using simplicial complexes," *Wireless Netw.*, vol. 19, no. 6, pp. 1121–1133, 2012.
- [10] J. Xu and V. Singh, "Unified hypergraph for image ranking in a multimodal context," in *Proc. 2012 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2012, pp. 2333–2336.
- [11] Y. Gao, M. Wang, D. Tao, R. Ji, and Q. Dai, "3-D object retrieval and recognition with hypergraph analysis," *IEEE Trans. Image Process.*, vol. 21, no. 9, pp. 4290–303, Sep. 2012.
- [12] A. Wilkerson, T. Moore, A. Swami, and H. Krim, "Simplifying the homology Of networks via strong collapses," in *Proc. 2013 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2013, pp. 5258–5262.
- [13] T. Wang and H. Krim, "Statistical classification of social networks," in *Proc. 2012 IEEE Int. Conf. Acoust., Speech Signal Process.*, 2012, pp. 3977–3980.
- [14] R. Singh, J. Xu, and B. Berger, "Global alignment of multiple protein interaction networks with application to functional orthology detection," *Proc. Nat. Acad. Sci.*, vol. 105, no. 35, pp. 12 763–12 768, 2008.
- [15] L. Peng, L. Liu, S. Chen, and Q. Sheng, "A network comparison algorithm for predicting the conservative interaction regions in protein-protein interaction network," in *Proc. 2010 IEEE 5th Int. Conf. Bio-Inspired Comput., Theories Appl.*, 2010, pp. 34–39.
- [16] S. Choobdar, P. Ribeiro, S. Bugla, and F. Silva, "Comparison of co-authorship networks across scientific fields using motifs," in *Proc. 2012 IEEE/ACM Int. Conf. Adv. Social Netw. Anal. Mining*, 2012, pp. 147–152.
- [17] L. Yong, Z. Yan, and C. Lei, "Protein-protein interaction network comparison based on wavelet and principal component analysis," in *Proc. 2010 IEEE Int. Conf. Bioinform. Biomed. Workshops*, 2010, pp. 430–437.
- [18] T. Milenković and N. Pržulj, "Uncovering biological network function via graphlet degree signatures," *Cancer Informat.*, vol. 6, pp. 257–273, Jan. 2008.
- [19] F. Memoli, "Gromov-Hausdorff distances in Euclidean spaces," in *Proc. 2008 IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Workshops*, 2008, pp. 1–8.
- [20] G. Carlsson, F. Memoli, A. Ribeiro, and S. Segarra, "Axiomatic construction of hierarchical clustering in asymmetric networks," 2014. [Online]. Available: <http://arxiv.org/abs/1301.7724>
- [21] W. Huang and A. Ribeiro, "Metrics in the space of high order networks," *IEEE Trans. Signal Process.*, vol. 64, no. 3, pp. 615–629, Feb. 2016.
- [22] W. Huang and A. Ribeiro, "Metrics in the space of high order proximity networks," in *Proc. 2015 IEEE Int. Conf. Acoust., Speech Signal Process.*, Brisbane, Australia, 2015, pp. 4135–4139.
- [23] S. Biasotti *et al.*, "Describing shapes by geometrical-topological properties of real functions," *ACM Comput. Surveys*, vol. 40, no. 4, 2008, Art. no. 12.
- [24] A. Cerri, B. D. Fabio, M. Ferri, P. Frosini, and C. Landi, "Betti numbers in multidimensional persistent homology are stable functions," *Math. Methods Appl. Sci.*, vol. 36, no. 12, pp. 1543–1557, 2013.
- [25] B. Di Fabio and P. Frosini, "Filtrations induced by continuous functions," *Topology Appl.*, vol. 160, no. 12, pp. 1413–1422, 2013.
- [26] H. Edelsbrunner, D. Letscher, and A. Zomorodian, "Topological persistence and simplification," in *Proc. 41st Annu. Symp. Found. Comput. Sci.*, 2000, pp. 454–463.
- [27] M. Lesnick, "The theory of the interleaving distance on multidimensional persistence modules," *Found. Comput. Math.*, vol. 15, no. 3, pp. 613–650, 2015.
- [28] K. Mischaikow and V. Nanda, "Morse theory for filtrations and efficient computation of persistent homology," *Discrete Comput. Geom.*, vol. 50, no. 2, pp. 330–353, 2013.
- [29] C. Berge, *Graphs And Hypergraphs*. Amsterdam, The Netherlands: North-Holland, 1976.
- [30] D. Burago, Y. Burago, and S. Ivanov, *A Course In Metric Geometry*, vol. 33. Providence, RI, USA: Amer. Math. Soc., 2001.

- [31] A. Zomorodian and G. Carlsson, "Computing persistent homology," *Discrete Comput. Geom.*, vol. 33, pp. 249–274, 2005.
- [32] H. Edelsbrunner and J. L. Harer, *Computational Topology: An Introduction*. Providence, RI, USA: Amer. Math. Soc., 2010.
- [33] S. Y. Oudot, *Persistence Theory: From Quiver Representations to Data Analysis*, vol. 209. Providence, RI, USA: Amer. Math. Soc., 2015.
- [34] R. Burkard, M. Dell'Amico, and S. Martello, *Assignment Problems*. Philadelphia, PA, USA: SIAM, 2009.
- [35] F. Chazal, D. Cohen-Steiner, L. J. Guibas, F. Mémoli, and S. Y. Oudot, "Gromov-Hausdorff stable signatures for shapes using persistence," *Proc. Eurograph. Symp. Geom. Process.*, vol. 28, no. 5, pp. 1393–1403, 2009.
- [36] F. Chazal, V. De Silva, and S. Oudot, "Persistence stability for geometric complexes," *Geometriae Dedicata*, vol. 173, no. 1, pp. 193–214, 2014.
- [37] A. M. Davie and A. J. Stothers, "Improved bound for complexity of matrix multiplication," *Proc. Roy. Soc. Edinburgh A, Math.*, vol. 143, no. 2, pp. 351–369, 2013.
- [38] P. Erdős and A. Rényi, "On the evolution of random graphs," *Publ. Math. Inst. Hungary Acad. Sci.*, vol. 5, pp. 17–61, 1960.
- [39] A. Tausz, M. Vejdemo-Johansson, and H. Adams, "JavaPlex: A research software package for persistent (co)homology," in *Proc. Int. Conf. Math. Sci.*, 2014, pp. 129–136.
- [40] M. A. A. Cox and T. F. Cox, "Multidimensional scaling," in *Handbook of Data Visualization* (ser. Springer Handbooks Computational Statistics). Berlin, Germany: Springer, 2008, pp. 315–347.
- [41] "Engineering Village: The place to find answers to engineering questions." (2014). [Online]. Available: <http://www.engineeringvillage.com/search/quick.url>
- [42] W. Huang and A. Ribeiro, "Hierarchical clustering given confidence intervals of metric distances," *IEEE Trans. Signal Process.*, Oct. 2016, submitted. [Online]. Available: <https://arxiv.org/abs/1610.04274>
- [43] W. Huang and A. Ribeiro, "Axiomatic hierarchical clustering for intervals of metric distances," in *Proc. Global Conf. Signal Inf. Process.*, Washington DC, USA, 2016, to be published.



Alejandro Ribeiro (S'02–M'07) received the B.Sc. degree in electrical engineering from the Universidad de la Republica Oriental del Uruguay, Montevideo, Uruguay, in 1998 and the M.Sc. and Ph.D. degrees in electrical engineering from the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, in 2005 and 2007, respectively. From 1998 to 2003, he was a member of the technical staff at Bellsouth Montevideo. After his M.Sc. and Ph.D. studies, in 2008 he joined the University of Pennsylvania (Penn), Philadelphia, where he is currently the Rosenbluth Associate Professor in the Department of Electrical and Systems Engineering. His research interests include the applications of statistical signal processing to the study of networks and networked phenomena. His focus is on structured representations of networked data structures, graph signal processing, network optimization, robot teams, and networked control. He received the 2014 O. Hugo Schuck Best Paper Award, the 2012 S. Reid Warren, Jr., Award presented by Penn's undergraduate student body for outstanding teaching, the NSF CAREER Award in 2010, and paper awards at the 2016 SSP Workshop, 2016 SAM Workshop, 2015 Asilomar SSC Conference, ACC 2013, ICASSP 2006, and ICASSP 2005. He is a Fulbright Scholar and a Penn Fellow.



Weiyu Huang (S'15) received the B.Eng. (Hons.) degree in electronics and telecommunication from Australian National University (ANU), Canberra, ACT, Australia, in 2012. He is currently working toward the Ph.D. degree in electrical and systems engineering at the University of Pennsylvania (Penn), Philadelphia. From 2011 to 2013, he was a Broadcasting Engineer and Policy Officer with the Australian Communication and Media Authority. His research interests include network theory, pattern recognition, graph signal processing, and the study of networked

data arising in human, social, and technological networks. He received the National Scholarship offered by the Ministry of Education, China, University Medal (best graduates), and H. A. Johns Award (best academic achievement) offered by the ANU, as well as the Department Best Ph.D. Colloquium Award by Penn for academic year 2015–2016.