



Distances and isomorphism between networks: stability and convergence of network invariants

Samir Chowdhury¹ · Facundo Mémoli²

Received: 7 August 2020 / Revised: 24 September 2022 / Accepted: 10 October 2022

© The Author(s), under exclusive licence to Springer Nature Switzerland AG 2022

Abstract

We develop the theoretical foundations of a generalized Gromov–Hausdorff distance between functions on networks that has recently been applied to various subfields of topological data analysis and optimal transport. These functional representations of networks, or networks for short, specialize in the finite setting to (possibly asymmetric) adjacency matrices and derived representations such as distance or kernel matrices. Existing literature utilizing these constructions cannot, however, benefit from continuous formulations because the continuum limits of finite networks under this distance are not well-understood. For example, while there are currently numerous persistent homology methods on finite networks, it is unclear if these methods produce well-defined persistence diagrams in the infinite setting. We resolve this situation by introducing the collection of compact networks that arises by taking continuum limits of finite networks and developing sampling results showing that this collection admits well-defined persistence diagrams. The difference between the network setting and metric setting arises as follows. For metric spaces, the isomorphism class of the Gromov–Hausdorff distance consists of isometric spaces and is thus very simple. For networks, the isomorphism class is rather complex, and contains representatives having different cardinalities and different topologies. We provide an exact characterization of a suitable notion of isomorphism for compact networks as well as alternative, stronger characterizations under additional topological regularity assumptions. Toward data applications, we describe a unified framework for developing quantitatively stable network invariants, provide basic examples, and cast existing results on the stability

✉ Samir Chowdhury
samirc@stanford.edu

Facundo Mémoli
memoli@math.osu.edu

¹ Department of Psychiatry and Behavioral Sciences, Stanford University, Stanford, CA 94305, USA

² Departments of Mathematics and Computer Science and Engineering, The Ohio State University, Columbus, OH 43210, USA

of persistent homology methods in this extended framework. To illustrate our theoretical results, we introduce a model of directed circles with finite reversibility and characterize their Dowker persistence diagrams.

Keywords Network data analysis · Directed networks · Generalized metric spaces · Gromov–Hausdorff distance · Gromov–Wasserstein distance

Mathematics Subject Classification 68T09 · 55N31 · 54E35

Contents

1	Introduction	1
1.1	Background	1
1.2	Setup and use cases	1
1.3	Main contributions and results	1
1.4	Related literature	1
1.5	Organization	1
2	Networks, isomorphism, and network distances	1
2.1	Definitions	1
2.2	The network distance	1
2.3	Weak isomorphism	1
2.4	Interpolating between strong and weak isomorphism	1
2.5	Coherence and Kuratowski-like embeddings	1
2.6	Skeletons	1
2.7	The second network distance	1
2.8	Special families: dissimilarity networks and directed metric spaces	1
2.9	Network models: the directed circles	1
2.10	The complexity of computing $d_{\mathcal{N}}$	1
2.11	Comparison of $d_{\mathcal{N}}$ with the cut metric on graphs	1
2.11.1	The cut distance between finite graphs	1
2.11.2	The cut distance and the Gromov–Hausdorff distance	1
2.12	Proofs from Section 2	1
2.12.1	Proofs related to Theorem 2.3.7	1
2.12.2	Proofs related to Theorem 2.6.3	1
2.12.3	Proof of Example 2.7.2	1
2.12.4	Proof of Proposition 2.7.5	1
2.12.5	Proof of Theorem 2.8.4	1
3	Metric structure of networks	1
3.1	Compact networks and finite sampling	1
3.2	Weak isomorphism in compact networks and characterization via motifs	1
3.3	The completeness of \mathcal{CN}/\cong^w	1
3.4	Precompact families in \mathcal{CN}/\cong^w	1
3.5	Geodesic structure on \mathcal{CN}/\cong^w	1
3.6	Proofs from Section 3	1
3.6.1	Proof of Theorem 3.1.3	1
3.6.2	Proof of Theorem 3.2.1	1
3.6.3	Proofs related to Theorem 3.2.2	1
3.6.4	Proofs related to Theorem 3.3.1	1
3.6.5	Proofs related to Theorem 3.5.4	1
4	Invariants, stability, and convergence	1
4.1	Preliminary setup	1
4.2	Network invariants and basic examples	1
4.3	Quantitative stability of basic invariants	1
4.4	Persistent homology invariants and their convergence	1

4.5	Vietoris–Rips, Dowker, path, and ordered tuple persistent homology methods on networks
4.6	Hierarchical clustering methods on asymmetric networks
4.6.1	Chain cost and path-connectedness
4.6.2	The nonreciprocal clustering method: definition and convergence
4.6.3	The reciprocal clustering method: definition and convergence
4.7	Clustering and persistence on the directed circle with finite reversibility
4.8	Proofs from Section 4
4.8.1	Proofs related to stability of basic invariants
4.8.2	Proofs related to Theorem 4.6.10
4.8.3	Proofs related to Theorem 4.6.13
5	Discussion
A	Experiments
A.1	Computational aspects and an algorithm for computing minimum matchings
A.2	Computational example: randomly generated networks
A.3	Computational example: simulated hippocampal networks
	References

1 Introduction

1.1 Background

Networks which show the relationships within and between complex systems are key tools in a variety of current research areas. Network analysis techniques such as modularity, core-periphery structure, clustering coefficients, and centrality (Newman 2010) have enjoyed enormous success in being used by researchers to derive insights from complex datasets, and this early success has ushered in a variety of approaches from different disciplines toward network data analysis. One such perspective that has appeared recently is to view a network as a *generalized metric space* (gms). In this viewpoint, a network is a set X along with a real-valued *edge weight function* $\omega_X : X \times X \rightarrow \mathbb{R}$. When X is finite, the pair (X, ω_X) can be represented as a real-valued square matrix. Such a matrix may encode the adjacency matrix of a (possibly directed) combinatorial graph or any real-valued function on pairs of graph vertices derived from the adjacency matrix. Examples of the latter include shortest path length matrices or positive semidefinite matrices such as the heat kernel obtained by exponentiating the graph Laplacian.

The gms perspective enables the import of numerous tools from the setting of metric spaces to the setting of networks. This program was initiated by the authors of Carlsson et al. (2013, 2014), who showed that single linkage hierarchical clustering could be defined for *directed* networks, and that a certain reformulation of the well-established *Gromov–Hausdorff* (GH) distance (d_{GH}) (Gromov 1981, 1999)—a *network distance*, denoted $d_{\mathcal{N}}$ —could be used to show that such a clustering method would be stable with respect to perturbations of the input data (Carlsson and Mémoli 2010). This network distance was defined between pairs of finite networks $(X, \omega_X), (Y, \omega_Y)$ via a combinatorial optimization problem that would search over correspondences between the points of X and Y as follows (see Definition 2.2.7 for a formal definition as used

in this work):

$$d_{\mathcal{N}}(X, Y) := \frac{1}{2} \min \left\{ \max_{(x,y),(x',y') \in R} |\omega_X(x, x') - \omega_Y(y, y')| : R \in \{0, 1\}^{|X| \times |Y|} \right. \\ \left. \text{with all row, column sums nonzero} \right\}.$$

Beyond its origins in metric geometry (Burago et al. 2001), the Gromov–Hausdorff distance between metric spaces has found applications in the context of shape and data analysis (Mémoli and Sapiro 2004; Mémoli 2011, 2012; Carlsson and Mémoli 2010; Chazal et al. 2009b). Thus one expects that an appropriate modification of the Gromov–Hausdorff distance would be a valuable tool in network analysis. Following this thread, the gms viewpoint enabled the development of *persistent homology* methods for directed networks (Chowdhury and Mémoli 2016b, 2018b; Turner 2019; Chowdhury and Mémoli 2018c; Dey et al. 2020) that came with stability guarantees obtained via $d_{\mathcal{N}}$. However, the theoretical guarantees of these methods were limited by the theoretical study of $d_{\mathcal{N}}$, which has to date been limited to the setting of finite networks. In this work, we remove this bottleneck and provide a comprehensive study of $d_{\mathcal{N}}$ and the theoretical framework surrounding its applications to hierarchical clustering, persistent homology, and machine learning using (not necessarily metric) dissimilarities.

1.2 Setup and use cases

Henceforth we will gradually stop referring to generalized metric spaces in favor of the “network” terminology, but we stress that this gms perspective will always be implicit. For the present discussion, the collection of networks is defined as (we will shortly add extra conditions; see Definition 2.1.2 for the formal definition):

$$\mathcal{N} := \{(X, \omega_X) : X \text{ a set}, \omega_X : X \times X \rightarrow \mathbb{R} \text{ any function}\}.$$

A (possibly directed, weighted) graph $G = (V, E)$ admits a variety of principled embeddings into \mathcal{N} , including:

- *Adjacency*: ω_V is defined to be the adjacency function of the graph.
- *Distance*: (for connected graphs/strongly connected digraphs) ω_V is defined to be the geodesic distance function on the graph.
- *Kernel*: ω_V is defined to be any of a number of graph kernels, many of which encode multiscale structure.

In this work, we study the metric structure of $(\mathcal{N}, d_{\mathcal{N}})$ with the goal of producing a common framework for different embeddings into \mathcal{N} . This naturally opens up a direction of future work on developing new embedding techniques—possibly parametrized and supervised, following Litman and Bronstein (2013)—that benefit from sharing a common ambient space.

Toward obtaining such characterizations, we first observe that the $d_{\mathcal{N}}$ function defined above can be immediately extended to infinite networks $(X, \omega_X), (Y, \omega_Y)$ by

replacing the min, max with inf, sup. We also observe that $d_{\mathcal{N}}$ only sees ω_X, ω_Y , and thus it is sensitive to structures on X, Y only if these structures affect ω_X, ω_Y . One such structure, and the minimum needed to at least be able to talk about limits and continuity, is a topology. Because topologies on X, Y seem to be a very basic requirement, we redefine \mathcal{N} as follows (cf. Definition 2.1.2):

$$\mathcal{N} := \{(X, \omega_X) : X \text{ a first countable topological space,} \\ \omega_X \text{ continuous w.r.t. product topology}\}.$$

Here, first countability is a mild assumption needed to prove some of our technical results. An interesting point to note is that given $(X, \omega_X), (Y, \omega_Y)$, there may be numerous first countable topologies on X, Y against which ω_X, ω_Y remain continuous. Referring to these as admissible topologies, we note that $d_{\mathcal{N}}(X, Y)$ is non sensitive to such admissible topological changes, as they do not perturb ω_X, ω_Y . Later in this work, we study the effects of constraining these admissible topologies. For now we remark on particular subcollections of \mathcal{N} . The following definition of *compact networks* is crucial for subsequent constructions:

$$\mathcal{CN} := \{(X, \omega_X) \in \mathcal{N} : X \text{ a compact, first countable topological space,} \\ \omega_X \text{ continuous w.r.t. product topology}\}.$$

As an example of the difference between \mathcal{N} and \mathcal{CN} , we show that compact networks admit the following characterization via *tripods*: if $X, Y \in \mathcal{CN}$, then $d_{\mathcal{N}}(X, Y) = 0$ if and only if there exists a set Z with surjective maps $\varphi_X : Z \rightarrow X, \varphi_Y : Z \rightarrow Y$ such that $\omega_X(\varphi_X(z), \varphi_X(z')) = \omega_Y(\varphi_Y(z), \varphi_Y(z'))$ for all $z, z' \in Z$ (see Sturm 2012; Mémoli 2017; Blumberg and Lesnick 2017; Bauer et al. 2020 for related applications of this technique). In particular, this result is not true without some form of the compactness assumption.

From the practical perspective, a first step is to clarify the types of networks that motivate our constructions, and to which our theory can be applied. One primary motivation has come from the domain of neuroimaging (Sporns 2011, 2012), where the technique of building a network out of thresholded correlations between neural systems (cf. Fig. 1) has led to great successes in identifying the neural correlates of human behavior. The underlying correlation matrices are called *functional connectivity* matrices, and are square, non-metric matrices with values in the interval $[-1, 1]$ (Friston 1994) (cf. Fig. 2). It has been shown that embedding such matrices into spaces with non-Euclidean geometry is beneficial for downstream analysis (Venkatesh et al. 2020), suggesting that building new spaces with different geometries may yield further insights from the data. For search systems in computer vision, utilizing non-metric dissimilarity matrices has been shown to be effective (Jegou et al. 2008; Hörster et al. 2007), and our work characterizes such matrices as members of a “nice” space \mathcal{CN} . From these perspectives, our study of \mathcal{CN} provides a new geometry in which to embed data and also provides posthoc theoretical justification for empirical methods involving non-metric dissimilarity measures.




Fig. 1 Four unit vectors and their inner products, i.e. their correlations, as well as their representations in the form of a matrix and a directed, weighted graph (self-loops are omitted from the figure). Networks with positive and negative values and zeros in arbitrary positions are the object of study in this work

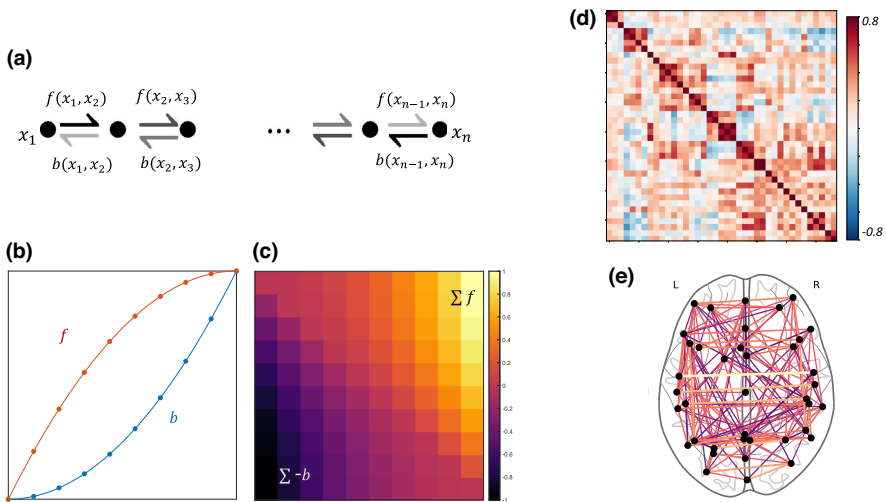


Fig. 2 Examples of networks derived from natural systems that violate metric axioms. **a** A simple (unnormalized) Markov chain $X = \{x_1, \dots, x_n\}$ with asymmetric forward and backward transition probabilities, as used in modeling biological/physical systems at broken detailed balance (Hill 2005; Perez et al. 2020). **b** Plot of $\sum_{i=1}^k f(x_i, x_{i+1})$ and $\sum_{i=1}^k b(x_i, x_{i+1})$ for $1 \leq k \leq n$. **c** Weight function ω_X defined for $i \leq j$ as $\omega_X(x_i, x_j) := \sum_{k=i}^j f(x_k, x_{k+1})$ and $\omega_X(x_j, x_i) := \sum_{k=i}^j -b(x_k, x_{k+1})$. **d, e** Functional connectivity (i.e. correlation) matrix and graph with edges corresponding to positive correlation values in the 80th percentile. Thresholding to obtain a graph is done in practice so that one may apply graph-connectivity tools for posthoc analysis, but the caveat is that there is no consensus on how to threshold weights, or how to deal with negative correlations. Our network framework provides posthoc analysis tools that accept arbitrary matrices without thresholding, thus avoiding potential information loss

Raw data in a variety of scientific applications may violate symmetry and/or triangle inequality, and this can be observed in natural phenomena such as human perception (Tversky and Gati 1982), broken detailed balance (Hill 2005), and discrepancies between time and distance measures (Matsumoto 1989). Exploratory data analysis techniques such as hierarchical clustering and persistent homology can be extended to accept nonmetric data as input. This was observed and successfully explored in Carlson et al. (2013, 2014, 2018), Chowdhury and Mémoli (2018b, c) and Turner (2019). These works focused on the setting of finite datasets due to their emphasis on appli-

cations, but an underlying question—that was made explicit in Turner (2019)—was that of characterizing the convergence of these techniques as the sample size grows to infinity. Already in these works and surrounding literature (Carlsson and Mémoli 2010; Chazal et al. 2014), it was implicit that one way to prove such a result would be to use a stability property of an appropriately reformulated GH distance along with a notion of dense sampling analogous to taking ε -nets in a metric space. Looking toward the needs of a framework which can work with matrices that are possibly asymmetric and can have both positive and negative entries, we study the structure of \mathcal{CN} and propose it as a convenient ambient space for the collection of finite networks. Studying \mathcal{CN} necessitates the departure from spectral techniques and well-grounded structures such as the cone of symmetric positive definite matrices (Pennec et al. 2006), but generates new possibilities via the lenses of applied topology and related machine learning approaches.

In this work we develop foundational tools for the theoretical integration of (possibly directed) networks into applied topology and broader machine learning frameworks. We study the notion of isomorphism under $d_{\mathcal{N}}$ and produce multiple characterizations of isomorphism that hold under a hierarchy of topological regularity assumptions. We explore the limits obtained via $d_{\mathcal{N}}$ -convergent sequences, and use these insights to give conditions on networks admitting well-defined *persistence diagrams*, i.e. networks on which persistent homology can be meaningfully applied. Toward grounding hierarchical clustering methods on directed networks, we develop the notion of sampling from networks equipped with probability measures and characterize the continuum limits of these methods. Finally, we derive network invariants that provide lower bounds on $d_{\mathcal{N}}$, and provide algorithms and implementations for comparing networks via these lower bounds.

1.3 Main contributions and results

The main contributions of this work are as follows:

- I We introduce $(\mathcal{CN}, d_{\mathcal{N}})$ and characterize its metric structure, including multiple perspectives on its isomorphism structure. Among these, we highlight the results that require the minimal amount of setup:
 - (a) $(\mathcal{CN}, d_{\mathcal{N}})$ is a complete, geodesic, pseudometric space that includes all square matrices, notably matrices violating metric axioms (Theorems 2.3.7, 3.3.1, 3.5.4).
 - (b) (Weak) Isomorphism in $(\mathcal{CN}, d_{\mathcal{N}})$ is exactly characterized via tripods (Theorem 3.2.1).
- II We extend known persistent homology methods on finite, possibly asymmetric networks to the infinite networks in \mathcal{CN} (Corollary 4.5.9). We also study convergence of hierarchical clustering methods for finite, possibly asymmetric networks on certain classes of infinite¹ networks equipped with Borel probability measures (Theorems 4.6.10, 4.6.13).
- III We define a directed network model—the family of compact directed circles with finite reversibility—and characterize its Dowker persistent homology by linking to

¹ Compact Polish spaces equipped with a bounded, measurable weight function ω , cf. Definition 4.1.1.

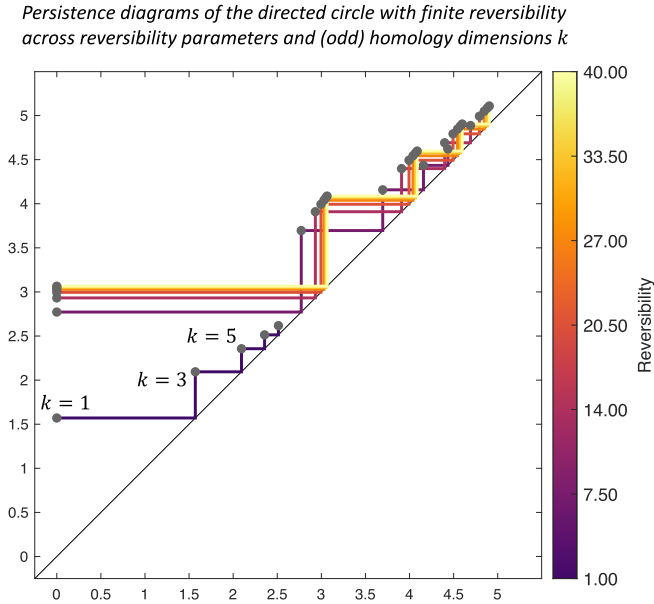


Fig. 3 Dowker persistence diagrams of $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$ across different reversibility parameters. Each broken line connects the persistence diagrams in homology dimensions $k = 1, 3, 5, 7, 9$ for a single reversibility parameter. This illustrates the exact characterization given by Theorem 4.7.3

existing results (Adamaszek et al. 2016; Adamaszek and Adams 2017) (cf. Fig. 3 and Theorem 4.7.3).

We now mention some additional results contained in this work. We show that the collection $(\mathcal{CN}, d_{\mathcal{N}})$ is complete and geodesic. Crucially, networks in \mathcal{CN} also admit ε -approximation by finite networks in the $d_{\mathcal{N}}$ sense. This result is used to show that compact networks have well-defined persistence diagrams, and to prove convergence results for both hierarchical clustering and persistent homology methods on networks. One may take a network (X, ω_X) and perturb the topology without incurring changes in $d_{\mathcal{N}}(X, \cdot)$, as long as ω_X remains unperturbed while still satisfying continuity. Combining this observation with the results stated above suggests that we would like descriptions of topologies on a pair (X, ω_X) that are controlled by ω_X . We provide such descriptions, which we call the *weak and strong coherent* topologies, by means of a Kuratowski-like embedding. Here we consider $\omega_X(x, \cdot)$, $\omega_X(\cdot, x)$ in the l^∞ function space, construct a metric in this space, and then study the topologies on X that are compatible with the metric topology in l^∞ . An interesting remark for practitioners is that an ε -ball in a coherent network may include points from disparate regions in the network. This is commonly seen in neuroimaging applications, where it is standard practice to study statistical correlations between disparate regions of the brain that coordinate to produce function.

When restricted to compact networks equipped with coherent topologies, we recover a result showing that a network can be recovered from the knowledge of its $n \times n$ subnetworks, for $n \in \mathbb{N}$. This generalizes a result about reconstructibility

of compact metric spaces from curvature sets (Gromov 1999). Notably in the metric setting, this result on curvature sets is the crucial ingredient in the work of Chowdhury et al. (2020) on producing families of polynomial-time lower bounds to the GH distance.

1.4 Related literature

A related thread which has received much attention from the machine learning community recently is that of the *Gromov–Wasserstein (GW) distance* between metric measure spaces,² which was introduced in Mémoli (2007, 2011) to simplify the difficult combinatorial optimization problem posed above by taking a relaxation based on *optimal transport (OT)*. This was later used in Hendrikson (2016) for network classification by taking graph geodesic distances to define the underlying metric. This perspective was further developed in Peyré et al. (2016), which discussed GW-averaging of distance and kernel matrix representations equipped with measures, and then in Chowdhury and Mémoli (2019), where the authors showed that all restrictions on the representations could be removed, and one would still obtain a pseudometric structure on such “measure networks” with well-defined notions of isomorphism. This theoretical framework was leveraged in Xu et al. (2019a, b) and Chowdhury and Needham (2020a), which used adjacency and heat kernel representations on graphs, respectively, for a variety of learning tasks. Further applications of GW distances were demonstrated in Chowdhury and Needham (2020b), where the authors developed the notion of GW-averaging and principal component analysis from a Riemannian perspective. Yet another direction in this line of work appeared in the formulation of Gromov–Wasserstein Factorization (Hongtengl 2020), which unrolled the loopy computations of GW distances into a deep learning framework for graph dictionary learning. The use of measures in the network-GW setting has important theoretical and practical implications, and our work in this paper contributes to the understanding of the “measure network” setting by studying the “network” setting in isolation and explaining which properties of the former are inherited from the latter.

For the generalized GH distance $d_{\mathcal{N}}$, the critical difference between the metric and network settings is in the zero set $\{d_{\mathcal{N}}(X, \cdot) = 0\}$, which can informally be thought of as the collection of symmetries. The collection of compact metric spaces equipped with d_{GH} is a metric space up to isometry. So two compact metric spaces (X, d_X) , (Y, d_Y) satisfying $d_{\text{GH}}(X, Y) = 0$ only differ by a bijection. In the setting of networks (X, ω_X) , (Y, ω_Y) with $d_{\mathcal{N}}$, however, one may have $d_{\mathcal{N}}(X, Y) = 0$ even when X and Y have different cardinalities. One far-reaching consequence of this phenomenon is the following seminal observation by Sturm (2012) in the related setting of GW distance: a certain subfamily of networks equipped with a GW distance admits Riemannian-like structure such as tangent spaces and exponential maps, but the symmetries force this collection to have the structure of a Riemannian *orbifold* rather than a manifold. The close relationship between GH and GW distances then makes it relevant and interesting to study the notion of *isomorphism* that is compatible with $d_{\mathcal{N}}$. This is one of the main directions of this work. Mildly relaxing the notion of isomorphism

² Metric spaces equipped with a Borel probability measure.

suggests considering limits and ε -approximation under $d_{\mathcal{N}}$, and these constructions are also developed in this work.

Our approaches for studying $(\mathcal{N}, d_{\mathcal{N}})$ are inspired by techniques used to study the structure $(\mathcal{M}, d_{\text{GH}})$ (Burago et al. 2001; Petersen 2006; Ivanov et al. 2016; Chowdhury and Mémoli 2018a) as well as methods used by Sturm (2012) to study the structure of metric measure (*mm*) spaces equipped with the GW distance. For example, the completion of the space of *mm* spaces turns out to be the space of relaxed *mm* spaces that satisfy the triangle inequality almost everywhere. This in turn is embedded in an ambient space which is isometric to the quotient space of symmetric L^2 functions on the unit square modulo measure-preserving transformations of the unit interval. In our setting, the idea of progressively relaxing the ambient space proceeds as follows: the collection of finite networks \mathcal{FN} is nested inside a collection of compact networks \mathcal{CN} that is complete. This space admits a natural notion of equivalence class. Working in a space where the elements need not satisfy metric requirements allows one to perform linear operations such as “addition” of networks as well as operations where the size of a network is changed while remaining in the equivalence class. These latter properties turn out to be fundamental in a parallel theory of Riemannian statistics on networks equipped with measures (Chowdhury and Needham 2020b).

For the related setting of quasimetric spaces, i.e. metric spaces where the symmetry condition is relaxed, there has been significant development of the topological consequences of quasimetry. Such spaces admit *forward-open* and *backward-open* ε -balls, and thus they may admit multiple topologies concurrently (Kelly 1963). These in turn lead to “forward” and “backward” notions of Cauchy sequences, completeness, and boundedness, which have been used in Shen and Zhao (2010) to extend results such as Gromov’s precompactness theorem from Riemannian to Finsler manifolds. Such results require some form of topological regularity. Shen and Zhao (2010) the regularity is provided by a condition called *finite reversibility* that places bounds on the asymmetry of the space. We study this condition in the context of a model of *directed circles with finite reversibility*, and in particular, we characterize the *Dowker persistent homology* (cf. Sect. 4.5) of such circles across all reversibility parameters to explain the difference caused by asymmetry. Furthermore, we study a weaker topological regularity condition called *coherence* that is implied by finite reversibility. This allows us to obtain a reconstruction result which, in the metric setting, states that compact metric spaces can be recovered from their curvature sets (Gromov 1999, Section 3.27 $\frac{1}{2}$). This in turn leads to efficient computational implementations for estimating d_{GH} -type distances (Oles et al. 2019). While our focus is on characterizing settings where asymmetric networks can be studied via theoretical and computational tools developed for metric spaces, we remark that numerous works dating back to Busemann have studied specialized versions of length structures, geodesics, and curvature for quasimetric spaces (Busemann 1950; Zaustinsky 1959; Cobzas 2012; Mennucci 2013). We also remark on an interesting connection to the work of Perrault-Joncas and Meila (2011), where directed graphs are modeled as samples from a manifold equipped with an asymmetric kernel: the directionality of the kernel is derived from a vector field on the manifold, which in turn provides the topological regularity suggested above.

The “generalized metric” approach to studying networks utilizes different techniques than combinatorial and spectral approaches (Newman 2010). A possible bridge

between metric and combinatorial approaches is given by observing the structural similarities between GH-type distances and the *cut metric* (Lovász 2012; Borgs et al. 2008) that has been deeply influential in statistical physics, theoretical computer science, and extremal graph theory. This connection is not fully fleshed out at present, but we report on a potential approach and leave an explicit connection open for future work. A different approach is provided by the notion of *structure space* (Jain and Obermayer 2009), wherein networks are modeled to have fixed size by appending null nodes (Calissano et al. 2020). This latter approach also gives rise to well-defined notions of averaging and principal component analysis, and suggests that future work could further elucidate these connections.

1.5 Organization

The presentation is divided into three parts. Section 2 introduces networks, notions of isomorphism, the network distance, topological considerations in defining families of networks, and a particular family of directed networks called the directed circles with finite reversibility. Section 3 develops metric properties of the space of networks equipped with the network distance and characterizes the isomorphism structure in this space. In particular, it sets up results on finite sampling and convergence of compact networks that we later use for proving convergence of network invariants. In Sect. 4 we set up the general framework of network invariants, which includes hierarchical clustering and persistent homology for infinite networks. We then establish the stability and convergence of these invariants.

Certain subsections fall into natural groups that may be read reasonably independently of other parts. The directed circles with finite reversibility are introduced in Sect. 2.9, and the full characterization of their (Dowker) persistent homology is presented in Sect. 4.7. The crucial structural aspects of compact networks, namely the sampling and isomorphism properties, are described in Sects. 3.1 and 3.2, and their application toward obtaining well-defined persistence diagrams is provided in Sect. 4.5.

Computational experiments using these invariants are provided in the Appendix. To clarify presentation, longer proofs are relegated to the end of the section in which they appear.

2 Networks, isomorphism, and network distances

In this section, we formulate definitions for networks, provide examples, construct some *model* networks, define distances between networks, and relate different notions of isomorphism. The notion of a unique limit for a d_N -convergent sequence of finite networks is somewhat complex, and the results in this section show why it is necessary to work in an appropriate subspace of \mathcal{N} .

2.1 Definitions

For real-world applications, the object of interest is often the collection of all finite networks, which we denote by \mathcal{FN} . Formally, one writes:

$$\mathcal{FN} := \{(X, \omega_X) : X \text{ a finite set, } \omega_X : X \times X \rightarrow \mathbb{R} \text{ any map}\}.$$

Example 2.1.1 (*Finite graphs and metric spaces*) This definition of a finite network is an immediate relaxation of the definition of a finite metric space. The name “network” is justified because graphs can be viewed as networks in the sense defined above. Given a connected, undirected graph $G = (V, E)$ having adjacency matrix A , degree matrix D , and geodesic distance matrix d_G , the pairs (V, d_G) and (V, A) , are both examples of networks. Consider also the graph Laplacian $L := D - A$, or its normalized form $\mathcal{L} := I - D^{-1}AD$. The pairs (V, L) , (V, \mathcal{L}) are also networks. In the case of (strongly connected), directed graphs, one could consider the *directed Laplacian* studied by Chung (2005) that is defined by deriving a Markov transition matrix and considering its associated Perron vector. In summary, there are a variety of quantities with different properties that can be derived from the combinatorial structure of a graph, and the definition of a finite network is general enough to comprise all of these. Each choice of a weight function that one may derive from a graph is a “lens” through which to study the data—such perspectives have been successfully harnessed in data analysis contexts (Singh et al. 2007).

In order to build a satisfactory theoretical foundation, one also needs to develop a formalism for infinite networks. Thus we proceed with the following definition.

Definition 2.1.2 (*Networks \mathcal{N}*) Let X be a first countable topological space, and let ω_X be a continuous function from $X \times X$ (endowed with the product topology) to \mathbb{R} . By a *network*, we will mean a pair (X, ω_X) . We will denote the collection of all networks by \mathcal{N} .

Notice in particular that \mathcal{N} includes metric spaces (they are first countable, and the distance function is continuous) as well as spaces that are quasi-metric or directed (no symmetry), pseudometric (no nondegeneracy), semimetric (no triangle inequality), or all of the above. Recall that a space is first countable if each point in the space has a countable local basis (see Steen and Seebach 1978, p. 7 for more details). First countability is a technical condition guaranteeing that when the underlying topological space of a network is compact, it is also sequentially compact.

Given a network (X, ω_X) , we will refer to the points of X as *nodes* and ω_X as the *weight function* of X . Pairs of nodes will be referred to as *edges*. Given a nonempty subset $A \subset X$, we will refer to $(A, \omega_X|_{A \times A})$ as the *sub-network* of X induced by A . For notational convenience, we will often write $X \in \mathcal{N}$ to mean $(X, \omega_X) \in \mathcal{N}$.

Recall that any finite set X can be equipped with the discrete topology, and any map $\omega_X : X \times X \rightarrow \mathbb{R}$ is continuous with respect to the discrete topology. Thus the elements of \mathcal{FN} trivially fit into the framework of \mathcal{N} . Throughout the paper, we will always understand finite networks to be equipped with the discrete topology.

While we are interested in \mathcal{FN} for practical applications, a key ingredient of our theoretical framework is the collection of *compact networks*. We define these to be the

networks (X, ω_X) satisfying the additional constraint that X is compact. The collection of compact networks is denoted \mathcal{CN} . Specifically, we write:

$$\mathcal{CN} := \{(X, \omega_X) : X \text{ compact, first countable topological space,} \\ \omega_X : X \times X \rightarrow \mathbb{R} \text{ continuous}\}.$$

Compact networks are of special practical interest because they can be finitely approximated in a manner that we will make precise in Sect. 3.1. Real world networks that are amenable to computational tasks are necessarily finite and may be viewed as samples drawn from an underlying compact network, so whenever possible, we will state our results for compact networks. Occasionally we will provide examples of noncompact networks to illustrate interesting theoretical points.

The aforementioned finite reversibility property is defined for the special subclass of *dissimilarity networks* that is described as follows.

Definition 2.1.3 (*Dissimilarity networks* \mathcal{N}^{dis}) A *dissimilarity network* is a network (X, A_X) where $A_X : X \times X \rightarrow \mathbb{R}_+$ and $A_X(x, x') = 0$ if and only if $x = x'$. Neither symmetry nor triangle inequality is assumed. The collection of all such networks is denoted \mathcal{N}^{dis} . The finite and compact settings are denoted as $\mathcal{FN}^{\text{dis}}$ and $\mathcal{CN}^{\text{dis}}$, respectively.

Definition 2.1.4 (*The reversibility parameter* ρ_X) The *reversibility* ρ_X of a dissimilarity network (X, A_X) is defined to be

$$\rho_X := \sup_{x \neq x' \in X} \frac{A_X(x, x')}{A_X(x', x)}.$$

(X, A_X) is said to have *finite reversibility* if $\rho_X < \infty$. Notice that $\rho_X \geq 1$ always, with equality iff ω_X is symmetric. Finitely reversible networks will feature heavily in Sects. 2.9 and 4.7. See also Figs. 3 and 12.

Dissimilarity networks satisfying the triangle inequality, but not symmetry, include the special class of objects called *directed metric spaces*, which we define below.

Definition 2.1.5 Let (X, A_X) be a dissimilarity network. Given any $x \in X$ and $r \in \mathbb{R}_+$, the *forward-open ball* of radius r centered at x is

$$B^+(x, r) := \{x' \in X : A_X(x, x') < r\}.$$

The *forward-open topology induced by* A_X is the topology on X generated by the collection $\{B^+(x, r) : x \in X, r > 0\}$. The idea of forward open balls is prevalent in the study of Finsler geometry; see Bao et al. (2012, p. 149) for details.

Definition 2.1.6 (*Directed metric spaces*) A *directed metric space* or *quasi-metric space* is a dissimilarity network (X, v_X) such that X is equipped with the forward-open topology induced by v_X and $v_X : X \times X \rightarrow \mathbb{R}_+$ satisfies:

$$v_X(x, x'') \leq v_X(x, x') + v_X(x', x'') \quad \text{for all } x, x', x'' \in X.$$

The function ν_X is called a directed metric or quasi-metric on X . Notice that compact directed metric spaces constitute a subfamily of $\mathcal{CN}^{\text{dis}}$.

Definition 2.1.7 (*Ultrametric/strong triangle inequality*) A network (X, ω_X) is said to satisfy the *ultrametric* or *strong triangle inequality* if

$$\omega_X(x, x') \leq \max \{ \omega_X(x, x''), \omega_X(x'', x') \} \quad \text{for all } x, x', x'' \in X.$$

Metric spaces satisfying this property are simply known as *ultrametric spaces*. The collections of networks satisfying this inequality are denoted \mathcal{N}^{ult} , $\mathcal{CN}^{\text{ult}}$, and $\mathcal{FN}^{\text{ult}}$ in the general, compact, and finite cases, respectively.

Example 2.1.8 Finite metric spaces and finite ultrametric spaces constitute basic examples of dissimilarity networks. Also note that finite dissimilarity networks are basic examples of networks that satisfy finite reversibility (cf. Definition 2.1.4).

Dissimilarity networks satisfying the symmetry condition, but not the triangle inequality, have a long history dating back to Fréchet (1906) and continuing with work by Pitcher et al. (1918), Niemytzki (1927), Galvin and Shore (1984, 1991), and many others, as summarized in Gruenhage (1984). One of the interesting directions in this line of work was the development of a “local triangle inequality” and related metrization theorems (Niemytzki 1927), which has been continued more recently in Waszkiewicz (2013).


Directed metric spaces with finite reversibility were studied in Shen and Zhao (2010), and constitute important examples of networks that are strictly non-metric. More specifically, the authors of Shen and Zhao (2010) extended notions of Hausdorff distance and Gromov–Hausdorff distance to the setting of directed metric spaces with finite reversibility, and our network distance $d_{\mathcal{N}}$ subsumes this theory while extending it to even more general settings.

Remark 2.1.9 (*Finsler metrics*) An interesting class of directed metric spaces arises from studying Finsler manifolds. A *Finsler manifold* (M, F) is a smooth, connected manifold M equipped with an asymmetric norm F (called a *Finsler function*) defined on each tangent space of M (Bao et al. 2012). A Finsler function induces a directed metric $d_F : M \times M \rightarrow \mathbb{R}_+$ as follows: for each $x, x' \in M$,

$$d_F(x, x') := \inf \left\{ \int_a^b F(\gamma(t), \dot{\gamma}(t)) dt : \gamma : [a, b] \rightarrow M \text{ a smooth curve joining } x \text{ and } x' \right\}.$$

Finsler metric spaces have received interest in the applied literature. In Sabau et al. (2013), the authors prove that Finsler metric spaces with *reversible geodesics* (i.e. the reverse curve $\gamma'(t) := \gamma(1 - t)$ of any geodesic $\gamma : [0, 1] \rightarrow M$ is also a geodesic) is a *weighted quasi-metric* (Sabau et al. 2013, p. 2). Such objects have been shown to be essential in biological sequence comparison (Stojmirović and Yi-Kuo 2009). In Sect. 2.9 we will study directed circles with finite reversibility, which have the property that ε -balls in such circles are not spherically symmetric. This is motivated by the more general setting of Randers manifolds (Bao et al. 2004), which are a special subfamily of Finsler manifolds.

Fig. 4 Networks over one and two nodes with their weight functions



Returning to the setting of general networks, a natural question in understanding the structure of \mathcal{N} would be: which elements of \mathcal{N} are equivalent? A suitable answer to this question requires us to develop notions of *isomorphism* that show various degrees of restrictiveness. These notions of isomorphism form a recurrent theme throughout this paper.

We first develop the notion of *strong isomorphism* of networks. The definition follows below.

Definition 2.1.10 (*Weight preserving maps*) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$. A map $\varphi : X \rightarrow Y$ is *weight preserving* if:

$$\omega_X(x, x') = \omega_Y(\varphi(x), \varphi(x')) \quad \text{for all } x, x' \in X.$$

Definition 2.1.11 (*Strong isomorphism*) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$. To say (X, ω_X) and (Y, ω_Y) are *strongly isomorphic* means that there exists a weight preserving bijection $\varphi : X \rightarrow Y$. We will denote a strong isomorphism between networks by $X \cong^s Y$. Note that this notion is exactly the usual notion of isomorphism between weighted graphs.

Strongly isomorphic networks formalize the idea that the information contained in a network should be preserved when we relabel the nodes in a compatible way.

Example 2.1.12 Networks with one or two nodes (cf. Fig. 4) will be very instructive in providing examples and counterexamples, so we introduce them now with some special terminology.

- A network with one node p can be specified by $\alpha \in \mathbb{R}$, and we denote this by $N_1(\alpha)$. We have $N_1(\alpha) \cong^s N_1(\alpha')$ if and only if $\alpha = \alpha'$.
- A network with two nodes will be denoted by $N_2(\Omega)$, where $\Omega = \begin{pmatrix} \alpha & \delta \\ \gamma & \beta \end{pmatrix} \in \mathbb{R}^{2 \times 2}$. Given $\Omega, \Omega' \in \mathbb{R}^{2 \times 2}$, $N_2(\Omega) \cong^s N_2(\Omega')$ if and only if there exists a permutation matrix P of size 2×2 such that $\Omega' = P \Omega P^T$.
- Any k -by- k matrix $\Sigma \in \mathbb{R}^{k \times k}$ induces a network on k nodes, which we refer to as $N_k(\Sigma)$. Notice that $N_k(\Sigma) \cong^s N_\ell(\Sigma')$ if and only if $k = \ell$ and there exists a permutation matrix P of size k such that $\Sigma' = P \Sigma P^T$.

Having defined a notion of isomorphism between networks, the next goal is to present the network distance $d_{\mathcal{N}}$ that is the central focus of this paper, and verify that $d_{\mathcal{N}}$ is compatible with strong isomorphism. We remind the reader that restricted formulations of this network distance have appeared in earlier applications of *hierarchical clustering* (Carlsson et al. 2014, 2013) and *persistent homology* (Chowdhury and Mémoli 2016b, 2018b,c) methods to network data, and our overarching goal in

this paper is to provide a theoretical foundation for this useful notion of network distance. In our presentation, we use a formulation of $d_{\mathcal{N}}$ that is more general than any other version available in the existing literature. As such, we proceed pedagogically and motivate the definition of $d_{\mathcal{N}}$ by tracing its roots in the metric space literature. Before ending this section, however, we will provide a final definition that is related to the matrices in Example 2.1.12.

For a sequence $(x_i)_{i=1}^n$ of nodes in a network X , we will denote the associated weight matrix by $((\omega_X(x_i, x_j)))_{i,j=1}^n$. Entry (i, j) of this matrix is simply $\omega_X(x_i, x_j)$.

Definition 2.1.13 (*Motif sets M_n*) For each $n \in \mathbb{N}$ and each $X \in \mathcal{CN}$, define $\Psi_X^n : X^n \rightarrow \mathbb{R}^{n \times n}$ to be the map $(x_1, \dots, x_n) \mapsto ((\omega_X(x_i, x_j)))_{i,j=1}^n$. Note that Ψ_X^n is simply a map that sends each sequence of length n to its corresponding weight matrix. Let $\mathcal{C}(\mathbb{R}^{n \times n})$ denote the closed subsets of $\mathbb{R}^{n \times n}$. Then let $M_n : \mathcal{CN} \rightarrow \mathcal{C}(\mathbb{R}^{n \times n})$ denote the map defined by

$$(X, \omega_X) \mapsto \left\{ \Psi_X^n(x_1, \dots, x_n) : x_1, \dots, x_n \in X \right\}.$$

We refer to $M_n(X)$ as the n -motif set of X . Notice that the image of M_n is closed in $\mathbb{R}^{n \times n}$ because each coordinate is the continuous image of the compact set $X \times X$ under ω_X , hence the image of M_n is compact in $\mathbb{R}^{n \times n}$ and hence closed.

Notice that for $X \in \mathcal{FN}$ and for a fixed $n \in \mathbb{N}$, the set $M_n(X)$ is a *finite* subset of $\mathbb{R}^{n \times n}$. The interpretation is that $M_n(X)$ is a bag containing all the motifs of X that one can form by looking at all subnetworks of size n (with repetitions).

Example 2.1.14 For the networks from Example 2.1.12, we have $M_1(N_2(\Omega)) = \{\alpha, \beta\}$ and

$$M_2(N_2(\Omega)) = \left\{ \begin{pmatrix} \alpha & \alpha \\ \alpha & \alpha \end{pmatrix}, \begin{pmatrix} \beta & \beta \\ \beta & \beta \end{pmatrix}, \begin{pmatrix} \alpha & \delta \\ \gamma & \beta \end{pmatrix}, \begin{pmatrix} \beta & \gamma \\ \delta & \alpha \end{pmatrix} \right\}, \quad M_2(N_1(\alpha)) = \left\{ \begin{pmatrix} \alpha & \alpha \\ \alpha & \alpha \end{pmatrix} \right\}.$$

Remark 2.1.15 Our definition of motif sets is inspired by a definition made by Gromov, termed “curvature classes,” in the context of compact metric spaces (Gromov 1999, §3.27).

2.2 The network distance

One strategy for defining a notion of distance between networks would be to take a well-understood notion of distance between metric spaces and extend it to all networks. The network distance $d_{\mathcal{N}}$ arises by following this strategy and extending the well-known Gromov–Hausdorff distance d_{GH} between compact metric spaces (Gromov 1981; Burago et al. 2001; Petersen 2006). The definition of d_{GH} is rooted in the *Hausdorff distance* d_{H} between closed subsets of a metric space. Given a metric space (Z, d_Z) and closed subsets $A, B \subseteq Z$, one defines:

$$d_{\text{H}}^Z(A, B) := \max \left(\sup_{a \in A} \inf_{b \in B} d_Z(a, b), \sup_{b \in B} \inf_{a \in A} d_Z(a, b) \right).$$

Definition 2.2.1 Given metric spaces (X, d_X) and (Y, d_Y) , the *Gromov–Hausdorff distance* d_{GH} between them is defined as:

$$d_{\text{GH}}((X, d_X), (Y, d_Y)) := \inf \{ d_{\text{H}}^Z(\varphi(X), \psi(Y)) : Z \text{ a metric space,} \\ \varphi : X \rightarrow Z, \psi : Y \rightarrow Z \text{ isometric embeddings} \}.$$

The Gromov–Hausdorff distance dates back to at least the early 1980s (Gromov 1981), and it satisfies numerous desirable properties. It is a valid metric on the collection of isometry classes of compact metric spaces, is complete, admits many precompact families, and has well-understood notions of convergence (Burago et al. 2001, Chapter 7). Moreover, it has found real-world applications in the *shape matching* (Mémoli and Sapiro 2004, 2005) and persistent homology literature (Chazal et al. 2009b), and its computational aspects have been studied as well (Mémoli 2012). As such, it is a strong candidate for use in defining a network distance.

Unfortunately, the formulation of d_{GH} above is heavily dependent on a metric space structure, and the notion of Hausdorff distance may not make sense in the setting of networks. So d_{GH} as defined above cannot be directly extended to a network distance. However, it turns out that there is a reformulation of d_{GH} that utilizes the language of *correspondences* (Kalton and Ostrovskii 1999; Burago et al. 2001). We present this construction next, and note that the resulting network distance $d_{\mathcal{N}}$ will agree with d_{GH} when restricted to metric spaces.

Definition 2.2.2 (*Correspondence*) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$. A *correspondence between X and Y* is a relation $R \subseteq X \times Y$ such that $\pi_X(R) = X$ and $\pi_Y(R) = Y$, where π_X and π_Y are the canonical projections of $X \times Y$ onto X and Y , respectively. The collection of all correspondences between X and Y will be denoted $\mathcal{R}(X, Y)$, abbreviated to \mathcal{R} when the context is clear.

Example 2.2.3 (*1-point correspondence*) Let X be a set, and let $\{p\}$ be the set with one point. Then there is a unique correspondence $R = \{(x, p) : x \in X\}$ between X and $\{p\}$.

Example 2.2.4 (*Diagonal correspondence*) Let $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_n\}$ be two enumerated sets with the same cardinality. A useful correspondence is the *diagonal correspondence*, defined as $\Delta := \{(x_i, y_i) : 1 \leq i \leq n\}$. When X and Y are infinite sets with the same cardinality, and $\varphi : X \rightarrow Y$ is a given bijection, then we write the diagonal correspondence as $\Delta := \{(x, \varphi(x)) : x \in X\}$.

Definition 2.2.5 (*Distortion of a correspondence dis*) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$ and let $R \in \mathcal{R}(X, Y)$. The *distortion* of R is given by:

$$\text{dis}(R) := \sup_{(x, y), (x', y') \in R} |\omega_X(x, x') - \omega_Y(y, y')|.$$

Remark 2.2.6 (*Composition of correspondences*) Let $(X, \omega_X), (Y, \omega_Y), (Z, \omega_Z) \in \mathcal{N}$, and let $R \in \mathcal{R}(X, Y)$, $S \in \mathcal{R}(Y, Z)$. Then we define:

$$R \circ S := \{(x, z) \in X \times Z \mid \exists y, (x, y) \in R, (y, z) \in S\}.$$

In the proof of Theorem 2.3.7, we verify that $R \circ S \in \mathcal{R}(X, Z)$, and that $\text{dis}(R \circ S) \leq \text{dis}(R) + \text{dis}(S)$.

Definition 2.2.7 (The first network distance $d_{\mathcal{N}}$) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$. We define the network distance between X and Y as follows:

$$d_{\mathcal{N}}((X, \omega_X), (Y, \omega_Y)) := \frac{1}{2} \inf_{R \in \mathcal{R}} \text{dis}(R).$$

When the context is clear, we will often write $d_{\mathcal{N}}(X, Y)$ to denote $d_{\mathcal{N}}((X, \omega_X), (Y, \omega_Y))$. We define the collection of *optimal correspondences* \mathcal{R}^{opt} between X and Y to be the collection $\{R \in \mathcal{R}(X, Y) : \text{dis}(R) = 2d_{\mathcal{N}}(X, Y)\}$. This set is always nonempty when $X, Y \in \mathcal{FN}$, but may be empty in general (see Example 2.3.4).

Remark 2.2.8 The intuition behind the preceding definition of network distance may be better understood by examining the case of a finite network. Given a finite set X and two edge weight functions ω_X, ω'_X defined on it, we can use the ℓ^∞ distance as a measure of network similarity between (X, ω_X) and (X, ω'_X) :

$$\|\omega_X - \omega'_X\|_{\ell^\infty(X \times X)} := \max_{x, x' \in X} |\omega_X(x, x') - \omega'_X(x, x')|.$$

A generalization of the ℓ^∞ distance is required when dealing with networks having different sizes: Given two sets X and Y , we need to decide how to match up points of X with points of Y . Any such matching will yield a subset $R \subseteq X \times Y$ such that $\pi_X(R) = X$ and $\pi_Y(R) = Y$, where π_X and π_Y are the projection maps from $X \times Y$ to X and Y , respectively. This is precisely a correspondence, as defined above. A valid notion of network similarity may then be obtained as the distortion incurred by choosing an optimal correspondence—this is precisely the idea behind the definition of the network distance above.

Remark 2.2.9 Some simple but important remarks are the following:

1. When restricted to metric spaces, $d_{\mathcal{N}}$ agrees with d_{GH} . This can be seen from the reformulation of d_{GH} in terms of correspondences (Burago et al. 2001, Theorem 7.3.25), Kalton and Ostrovskii (1999). Whereas d_{GH} vanishes only on pairs of isometric compact metric spaces, $d_{\mathcal{N}}$ vanishes on a broader family of networks that we will describe more fully in Sect. 2.3.
2. Given $X, Y \in \mathcal{FN}$, the network distance reduces to the following:

$$d_{\mathcal{N}}(X, Y) = \frac{1}{2} \min_{R \in \mathcal{R}} \max_{(x, y), (x', y') \in R} |\omega_X(x, x') - \omega_Y(y, y')|.$$

Moreover, there is always at least one optimal correspondence R^{opt} for which $d_{\mathcal{N}}(X, Y)$ is achieved; this is a consequence of considering finite networks.

3. For any $X, Y \in \mathcal{CN}$, we have $\mathcal{R}(X, Y) \neq \emptyset$, and $d_{\mathcal{N}}(X, Y)$ is always bounded. Indeed, $X \times Y$ is always a valid correspondence between X and Y . So we have:

$$d_{\mathcal{N}}(X, Y) \leq \frac{1}{2} \text{dis}(X \times Y) \leq \frac{1}{2} \left(\sup_{x, x'} |\omega_X(x, x')| + \sup_{y, y' \in Y} |\omega_Y(y, y')| \right) < \infty.$$




Fig. 5 Relaxing the requirements on the maps of this “tripod structure” is a natural way to weaken the notion of strong isomorphism

Example 2.2.10 Now we give some examples to illustrate the preceding definitions.

- For $\alpha, \alpha' \in \mathbb{R}$ consider two networks with one node each: $N_1(\alpha) = (\{p\}, \alpha)$ and $N_1(\alpha') = (\{p'\}, \alpha')$. By Example 2.2.3 there is a unique correspondence $R = \{(p, p')\}$ between these two networks, so that $\text{dis}(R) = |\alpha - \alpha'|$ and as a result $d_{\mathcal{N}}(N_1(\alpha), N_1(\alpha')) = \frac{1}{2}|\alpha - \alpha'|$.
- Let $(X, \omega_X) \in \mathcal{FN}$ be any network and for $\alpha \in \mathbb{R}$ let $N_1(\alpha) = (\{p\}, \alpha)$. Then $R = \{(x, p), x \in X\}$ is the unique correspondence between X and $\{p\}$, so that

$$d_{\mathcal{N}}(X, N_1(\alpha)) = \frac{1}{2} \max_{x, x' \in X} |\omega_X(x, x') - \alpha|.$$

We now test whether $d_{\mathcal{N}}$ is compatible with strong isomorphism. Given two strongly isomorphic networks, i.e. networks $(X, \omega_X), (Y, \omega_Y)$ and a weight preserving bijection $\varphi : X \rightarrow Y$, it is easy to use the diagonal correspondence (Example 2.2.4) to verify that $d_{\mathcal{N}}(X, Y) = 0$. However, it is easy to see that the reverse implication is not true in general. Using the one-point correspondence (Example 2.2.3), one can see that $d_{\mathcal{N}}(N_1(1), N_2(\mathbb{1}_{2 \times 2})) = 0$. Here $\mathbb{1}_{n \times n}$ denotes the all-ones matrix of size $n \times n$ for any $n \in \mathbb{N}$. However, these two networks are not strongly isomorphic, because they do not even have the same cardinality. Thus we need to search for a different, perhaps weaker notion of isomorphism.

2.3 Weak isomorphism

To proceed in this direction, first notice that a strong isomorphism between two networks (X, ω_X) and (Y, ω_Y) , given by a bijection $f : X \rightarrow Y$, is equivalent to the following “tripod” condition (cf. Fig. 5): there exists a set Z and bijective maps $\varphi_X : Z \rightarrow X, \varphi_Y : Z \rightarrow Y$ such that $\omega_X(\varphi_X(z), \varphi_X(z')) = \omega_Y(\varphi_Y(z), \varphi_Y(z'))$ for each $z, z' \in Z$. To see this, simply let $Z = \{(x, f(x)) : x \in X\}$ and let φ_X, φ_Y be the projection maps on the first and second coordinates, respectively. Based on this observation, we make the next definition.

Definition 2.3.1 Let (X, ω_X) and $(Y, \omega_Y) \in \mathcal{N}$. We define X and Y to be *Type I weakly isomorphic*, denoted $X \cong_I^w Y$, if there exists a set Z and surjective maps $\varphi_X : Z \rightarrow X$ and $\varphi_Y : Z \rightarrow Y$ such that

$$\omega_X(\varphi_X(z), \varphi_X(z')) = \omega_Y(\varphi_Y(z), \varphi_Y(z')) \quad \text{for each } z, z' \in Z. \quad (1)$$

Notice that Type I weak isomorphism is in fact a *relaxation* of the notion of strong isomorphism. Indeed, if in addition to being surjective, we require the maps ϕ_X and ϕ_Y




Fig. 6 Note that Remark 2.3.2 does not fully characterize weak isomorphism, even for finite networks: All three networks above, with the given weight matrices, are Type I weakly isomorphic since C maps surjectively onto A and B . But there are no surjective, weight preserving maps $A \rightarrow B$ or $B \rightarrow A$

to be injective, then the strong notion of isomorphism is recovered. In this case, the map $\phi_Y \circ \phi_X^{-1} : X \rightarrow Y$ would be a weight preserving bijection between the networks X and Y . The relaxation of strong isomorphism to a Type I weak isomorphism is illustrated in Fig. 5. Also observe that the relaxation is *strict*. For example, the networks $X = N_1(1)$ and $Y = N_2(\mathbb{1}_{2 \times 2})$, are weakly but not strongly isomorphic via the map that sends both nodes of Y to the single node of X .

Remark 2.3.2 (*Surjective maps induce Type I isomorphism*) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{CN}$ and suppose $\varphi : X \rightarrow Y$ is a surjective map such that $\omega_X(x, x') = \omega_Y(\varphi(x'), \varphi(x'))$ for all $x, x' \in X$. Then X and Y are Type I weakly isomorphic. This result follows from Definition 2.3.1 by: (1) choosing $Z = X$, (2) letting ϕ_X be the identity map, and (3) letting $\phi_Y = \varphi$. The converse implication, i.e. that Type I weak isomorphism implies the existence of a surjective map as above, is not true: an example is shown in Fig. 6.

When dealing with infinite networks, it will turn out that an even weaker notion of isomorphism is required. We define this weakening next.

Definition 2.3.3 Let (X, ω_X) and $(Y, \omega_Y) \in \mathcal{N}$. We define X and Y to be *Type II weakly isomorphic*, denoted $X \cong_{\Pi}^w Y$, if for each $\varepsilon > 0$, there exists a set Z_ε and surjective maps $\phi_X^\varepsilon : Z_\varepsilon \rightarrow X$ and $\phi_Y^\varepsilon : Z_\varepsilon \rightarrow Y$ such that

$$|\omega_X(\phi_X^\varepsilon(z), \phi_X^\varepsilon(z')) - \omega_Y(\phi_Y^\varepsilon(z), \phi_Y^\varepsilon(z'))| < \varepsilon \quad \text{for all } z, z' \in Z_\varepsilon. \quad (2)$$

Example 2.3.4 (*Infinite networks without optimal correspondences*) The following example illustrates the reason we had to develop multiple notions of weak isomorphism. The key idea is that the infimum in Definition 2.2.7 is not necessarily obtained when X and Y are infinite networks. To see this, let (X, ω_X) denote $[0, 1]$ equipped with the Euclidean distance, and let (Y, ω_Y) denote $\mathbb{Q} \cap [0, 1]$ with the restriction of the Euclidean distance. Since the closure of Y in $[0, 1]$ is just X , the Hausdorff distance between X and Y is zero (recall that given $A, B \subseteq \mathbb{R}$, we have $d_{\mathbb{H}}^{\mathbb{R}}(A, B) = 0$ if and only if $\overline{A} = \overline{B}$ Burago et al. 2001, Proposition 7.3.3). It follows from the definition of d_{GH} (Definition 2.2.1) and the equivalence of $d_{\mathcal{N}}$ and d_{GH} on metric spaces (Remark 2.2.9) that $d_{\mathcal{N}}(X, Y) = 0$.

However, one cannot define an optimal correspondence between X and Y . To see this, assume towards a contradiction that R^{opt} is such an optimal correspondence,

i.e. $\text{dis}(R^{\text{opt}}) = 0$. For each $x \in X$, there exists $y_x \in Y$ such that $(x, y_x) \in R^{\text{opt}}$. By making a choice of $y_x \in Y$ for each $x \in X$, define a map $f : X \rightarrow Y$ given by $x \mapsto y_x$. Then $d_X(x, x') = d_Y(f(x), f(x'))$ for each $x, x' \in X$. Thus f is an isometric embedding from X into itself (note that $Y \subseteq X$). But $X = [0, 1]$ is compact, and an isometric embedding from a compact metric space into itself must be surjective (Burago et al. 2001, Theorem 1.6.14). This is a contradiction, because $f(X) \subseteq Y \neq X$.

We observe that $d_{\mathcal{N}}(X, Y) = 0$ and so X and Y are weakly isomorphic of Type II, but not of Type I. To see this, assume towards a contradiction that X and Y are Type I weakly isomorphic. Let Z be a set with surjective maps $\varphi_X : Z \rightarrow X$ and $\varphi_Y : Z \rightarrow Y$ satisfying $\omega_X \circ (\varphi_X, \varphi_X) = \omega_Y \circ (\varphi_Y, \varphi_Y)$. Then $\{(\varphi_X(z), \varphi_Y(z)) : z \in Z\}$ is an optimal correspondence. This is a contradiction by the previous reasoning.

Recall that our motivation for introducing notions of isomorphism on \mathcal{N} was to determine which networks deserve to be considered equivalent. It is easy to see that strong isomorphism induces an equivalence class on \mathcal{N} . The same is true for both types of weak isomorphism, and we record this result in the following proposition.

Proposition 2.3.5 *Weak isomorphism of Types I and II both induce equivalence relations on \mathcal{N} .*

In the setting of \mathcal{FN} , it is not difficult to show that the two types of weak isomorphism coincide. This is the content of the next proposition. By virtue of this result, there is no ambiguity in dropping the “Type I/II” modifier when saying that two finite networks are weakly isomorphic.

Proposition 2.3.6 *Let $X, Y \in \mathcal{FN}$ be finite networks. Then X and Y are Type I weakly isomorphic if and only if they are Type II weakly isomorphic.*

Type I weak isomorphisms will play a vital role in the content of this paper, but for now, we focus on Type II weak isomorphism. The next theorem justifies calling $d_{\mathcal{N}}$ a network distance, and shows that $d_{\mathcal{N}}$ is compatible with Type II weak isomorphism.

Theorem 2.3.7 *$d_{\mathcal{N}}$ is a metric on \mathcal{N} modulo Type II weak isomorphism.*

Proofs of the preceding results are provided in Sect. 2.12.1. For finite networks, we immediately obtain:

The restriction of $d_{\mathcal{N}}$ to \mathcal{FN} yields a metric modulo Type I weak isomorphism.

The proof of Proposition 2.3.6 will follow from the proof of Theorem 2.3.7. In fact, an even stronger result is true: weak isomorphism of Types I and II coincide for compact networks as well. We present this statement in Sect. 3.2. By virtue of this latter result, we will use the notation \cong^w without I/II qualifiers to mean Type I weak isomorphism between compact networks.

Remark 2.3.8 (Making $d_{\mathcal{N}}$ more sensitive to topology) As will be clear from the proof of Theorem 2.3.7, $d_{\mathcal{N}}$ is actually a metric modulo Type II weak isomorphism on the collection $\{(X, \omega_X) : X \text{ a set}, \omega_X : X \times X \rightarrow \mathbb{R} \text{ any function}\}$. In other words, when applied in this general form, $d_{\mathcal{N}}$ is insensitive to topology. However, restricting to subcollections of set-function pairs may improve this sensitivity. When one starts with a

network in \mathcal{N} (comprising a continuous weight function over a first countable topological space) it is possible to refine the topology while maintaining first countability, even taking the discrete topology at the extreme, and remain in the same Type II weak isomorphism class. Alternatively, one may start with a network in \mathcal{CN} and refine the underlying topology by adding a finite number of open sets—the resulting network will still be in \mathcal{CN} , and will even be in the same Type I weak isomorphism class. In each of these cases, we progressively narrowed down the types of topological changes (possibly infinite refinements, and then finite refinements) against which $d_{\mathcal{N}}$ would be guaranteed to be insensitive by studying nested subcollections of set-function pairs. Even the insensitivity to refinements need not be problematic, as it will turn out that each Type I weak isomorphism class has representatives with “nice” topology. These ideas will be introduced in the next section.

We end the current subsection with the following definition, which is reminiscent of the notion of ε -nets in metric spaces.

Definition 2.3.9 (*ε -Approximations*) Let $\varepsilon > 0$. A network $(X, \omega_X) \in \mathcal{N}$ is said to be ε -approximable by $(Y, \omega_Y) \in \mathcal{N}$ if $d_{\mathcal{N}}(X, Y) < \varepsilon$. In this case, Y is said to be an ε -approximation of X . Typically, we will be interested in the case where X is infinite and Y is finite, i.e. in ε -approximating infinite networks by finite networks.

2.4 Interpolating between strong and weak isomorphism

As we saw in the simple examples discussed above, strong isomorphism implies weak isomorphism, and weak isomorphism does not in general imply strong isomorphism. However, we will show that weakly isomorphic networks live over a “base space” of strongly isomorphic networks. The following definitions enable us to formulate the appropriate statement.

Definition 2.4.1 (*Automorphisms*) Let $(X, \omega_X) \in \mathcal{CN}$. We define the *automorphisms* (*Aut*) of X to be the collection

$$\text{Aut}(X) := \{\varphi : X \rightarrow X : \varphi \text{ a weight preserving bijection}\}.$$

Definition 2.4.2 (*Poset of weak isomorphism*) Let $(X, \omega_X) \in \mathcal{CN}$. Define a set $\mathbf{p}(X)$ as follows:

$$\mathbf{p}(X) := \{(Y, \omega_Y) \in \mathcal{CN} : \text{there exists a surjective, weight preserving map } \varphi : X \rightarrow Y\}.$$

Next we define a partial order \leq on $\mathbf{p}(X)$ as follows: for any $(Y, \omega_Y), (Z, \omega_Z) \in \mathbf{p}(X)$,

$$(Y, \omega_Y) \leq (Z, \omega_Z) \iff \text{there exists a surjective, weight preserving map } \varphi : Z \rightarrow Y.$$

Then the set $\mathbf{p}(X)$ equipped with \leq is called the *poset of weak isomorphism of X* .

Definition 2.4.3 (*Terminal networks in \mathcal{CN}*) Let $(X, \omega_X) \in \mathcal{CN}$. A compact network $Z \in \mathbf{p}(X)$ is *terminal* if:




Fig. 7 *Left* Z represents a terminal object in $\mathbf{p}(X)$, and f, g are weight preserving surjections $X \rightarrow Z$. Here $\varphi \in \text{Aut}(Z)$ is such that $g = \varphi \circ f$. *Right* Here we show more of the poset structure of $\mathbf{p}(X)$. In this case we have $X \geq V \geq Y \dots \geq Z$

1. For each $Y \in \mathbf{p}(X)$, there exists a weight preserving surjection $\varphi : Y \rightarrow Z$.
2. Let $Y \in \mathbf{p}(X)$. If $f : Y \rightarrow Z$ and $g : Y \rightarrow Z$ are weight preserving surjections, then there exists $\varphi \in \text{Aut}(Z)$ such that $g = \varphi \circ f$ (Fig. 7).

In Sect. 2.6 we define a construction called the *skeleton* of a network and show that it is terminal. One of our main results (Theorem 3.2.2) shows that, under some mild topological regularity conditions,

weakly isomorphic networks have strongly isomorphic skeletons.

A terminal network captures the idea of a minimal substructure of a network. One may ask if anything interesting can be said about superstructures of a network. This motivates the following construction of a “blow-up” network. We provide an illustration in Fig. 8.

Definition 2.4.4 Let (X, ω_X) be any network. Let $\mathbf{k} = (k_x)_{x \in X}$ be a choice of an index set k_x for each node $x \in X$. Consider the network $X[\mathbf{k}]$ with node set $\bigcup_{x \in X} \{(x, i) : i \in k_x\}$ and weights ω given as follows: for $x, x' \in X$ and for $i \in k_x, i' \in k_{x'}$,

$$\omega((x, i), (x', i')) := \omega_X(x, x').$$

The topology on $X[\mathbf{k}]$ is given as follows: the open sets are of the form $\bigcup_{x \in U} \{(x, i) : i \in k_x\}$, where U is open in X . By construction, $X[\mathbf{k}]$ is first countable with respect to this topology. We will call any such $X[\mathbf{k}]$ a *blow-up network* of X .

In a blow-up network of X , each node $x \in X$ is replaced by another network, indexed by k_x . All internal weights of this network are constant and all outgoing weights are preserved from the original network. If X is compact, then so is $X[\mathbf{k}]$.

We also observe that X is weakly isomorphic to any of its blow-ups $Y = X[\mathbf{k}]$. To see this, let $Z = X[\mathbf{k}]$, let $\phi_Y : Z \rightarrow Y$ be the map sending each (x, i) to (x, i) , and let $\phi_X : Z \rightarrow X$ be the map sending each (x, i) to x . Then ϕ_X, ϕ_Y are surjective, weight preserving maps from Z onto X and Y respectively. By Remark 2.3.2, we obtain $X \cong_1^w Y$.

Later in Proposition 2.7.5, we will show that Type I weakly isomorphic networks may be blown-up to strongly isomorphic networks.




Fig. 8 Interpolating between the skeleton and blow-up constructions

2.5 Coherence and Kuratowski-like embeddings

Thus far we have maintained one-way control over the topology of a network, i.e. the topology of X restricts the collection of weight functions ω_X that are continuous and hence admissible. In certain settings it will be useful to impose control in the other direction, i.e. to restrict the collection of admissible topologies via ω_X . We produce this type of topological control via the embeddings $x \mapsto \omega_X(x, \cdot)$, $x \mapsto \omega_X(\cdot, x)$. These are variants of the well-known *Kuratowski embedding* (Heinonen 2012, Chapter 12) that is defined for compact metric spaces (X, d_X) as the map $x \mapsto d_X(x, \cdot)$. Such a map gives an isometric embedding of (X, d_X) into the Banach space $l^\infty(X)$.

Let $(X, \omega_X) \in \mathcal{N}$. For any $x, x' \in X$, define:

$$\begin{aligned} \delta_X^+(x, x') &:= \|\omega_X(x, \cdot) - \omega_X(x', \cdot)\|_\infty, & \delta_X^-(x, x') &:= \|\omega_X(\cdot, x) - \omega_X(\cdot, x')\|_\infty \\ \delta_X(x, x') &:= \max(\delta_X^+(x, x'), \delta_X^-(x, x')). \end{aligned} \quad (3)$$

Both δ_X^+ and δ_X^- are pseudometrics generated by the asymmetric Kuratowski-like embeddings, and δ_X , being the max of two pseudometrics, is also a pseudometric. A related generalization is later provided in Definition 3.6.7, and a figure demonstrating its use is provided in Fig. 9.

Definition 2.5.1 Let $(X, \omega_X) \in \mathcal{N}$. We say that X has a *weak coherent topology* if for any $x \in X$ and any sequence $(x_n)_n$ in X , one has

$$x_n \rightarrow x \quad \text{iff} \quad \delta_X(x, x_n) \rightarrow 0. \quad (4)$$

Definition 2.5.2 Let $(X, \omega_X) \in \mathcal{N}$. We say that X has a *strong coherent topology* if for any $x \in X$, any sequence $(x_n)_n$ in X , and any subspace $(S, \omega_X|_{S \times S})$ containing x and $\{x_n\}_n$, the following are equivalent:

- $x_n \rightarrow x$
- $\delta_X(x, x_n) \rightarrow 0$
- $\delta_S(x, x_n) \rightarrow 0$

Metric spaces satisfy strong coherence due to the triangle inequality—in particular, strong coherence encodes a generalized consequence of the linear constraints defining metric spaces. For clarification, let us verify that the statements in Definition 2.5.2 hold for metric spaces. In a metric space (X, d_X) , to say that a sequence $(x_n)_n$ converges to a point $x \in X$ means that $d_X(x_n, x) \rightarrow 0$, i.e. it suffices to check the value of d_X on the points $S := \{x\} \cup \{x_n\}_n$. It is well-known that the Kuratowski embedding is an isometric embedding for metric spaces, and hence $d_X(x, x_n) \rightarrow 0$ is equivalent to $\delta_X(x, x_n) \rightarrow 0$. Next, considering the space (S, d_S) where $d_S := d_X|_{S \times S}$, one applies the triangle inequality to obtain

$$\|d_S(x_n, \cdot) - d_S(x, \cdot)\|_\infty \rightarrow 0 \quad \text{and} \quad \|d_S(\cdot, x_n) - d_S(\cdot, x)\|_\infty \rightarrow 0.$$

Hence $\delta_S(x, x_n) \rightarrow 0$. Finally consider (S, d_S) by itself and suppose $\delta_S(x, x_n) \rightarrow 0$ for some $x, \{x_n\}_n$ in S . Consider the isometric embedding given by the inclusion $(S, d_S) \hookrightarrow (X, d_X)$. As a function, $d_X : X \times X \rightarrow \mathbb{R}_+$ extends d_S , and the linear constraints coming from the triangle inequality force the following:

$$\|d_X(x_n, \cdot) - d_X(x, \cdot)\|_\infty \rightarrow 0 \quad \text{and} \quad \|d_X(\cdot, x_n) - d_X(\cdot, x)\|_\infty \rightarrow 0.$$

The notion of strong coherence captures this property of isometric embedding without specifying, a priori, the form of these linear constraints.

We refer to such topologies as the *weak* or *strong coherent topology generated by* ω_X , and refer to networks satisfying the property of Eq. (4) as *weak* or *strong coherent networks*.

Remark 2.5.3 Directed metric spaces with finite reversibility comprise a collection of objects that are strictly more general than metric spaces and still satisfy strong coherence.

Definition 2.5.4 We refer to δ_X above as the *canonical pseudometric* of the network (X, ω_X) .

Remark 2.5.5 Kelly (1963) has studied separation axioms in spaces with multiple topologies, as is the case here with the topologies generated by δ_X^+, δ_X^- . Also note that in a first countable space, the topology is determined precisely by the convergent sequences. Generalizing this property leads to *sequential spaces*, which have been studied by Franklin (1965).

Remark 2.5.6 (*Weak coherence is strictly weaker than strong coherence*) The construction of δ_X above is *extrinsic* in the sense that to assign $\delta_X(x, x')$, we used knowledge




Fig. 9 To provide intuition for the Kuratowski-like embedding metric δ (cf. Definition 2.5.4), we plot $\delta(e^0, e^{i\theta})$ using ω_{S^1} (left) and $\omega_{S^1, \rho}$ (right) as θ varies along S^1 . We have rescaled S^1 to be parametrized by $[0, 1)$ for convenience. The figure on the left shows that for the strictly directed circle, $\delta(e^0, \cdot)$ is minimized at the antipode $e^{i\pi}$, and that for all ε values less than the diameter, the δ -ball of radius ε centered at e^0 contains only e^0 . The same holds for any other point $e^{i\theta}$ by rotational symmetry. Thus the metric topology induced by the Kuratowski-like embedding is actually discrete. Compare this with the figure on the right. The different colored lines correspond to different choices of ρ , taken on a logarithmic scale. We observe that δ takes on an M-shape that converges to the shape of the geodesic distance as $\rho \downarrow 1$. The shape of the plot also tells us that for any finite reversibility parameter, the δ -ball of radius ε centered at e^0 (and thus any other point by rotational symmetry) is centrally symmetric, as is the case for a ball of radius ε using the standard geodesic metric

from *all* of X . While δ_X formally gives us a metric, its construction suffers from consistency issues when embedding into a larger space. For example, suppose we have an injective, weight-preserving map $f : X \rightarrow Y$. A priori, we may have

$$\begin{aligned} \delta_X(x, x') &\neq \max(\|\omega_Y(f(x), \cdot) - \omega_Y(f(x'), \cdot)\|_\infty, \|\omega_Y(\cdot, f(x)) - \omega_Y(\cdot, f(x'))\|_\infty) \\ &= \delta_Y(f(x), f(x')) \end{aligned}$$

for $x, x' \in X$, which is an issue that one does not face with isometric embeddings of bona fide metric spaces. This is illustrated in Fig. 10. Consequently, given $x_n \rightarrow x$, one may have $\delta_X(x, x_n) \rightarrow 0$ and $\delta_Y(f(x), f(x_n)) \not\rightarrow 0$. An alternative viewpoint is the following. When *extending* a metric space (cf. Melleray 2007), say by a single point, one needs to choose distances from the new point to all other points that satisfy the linear constraints posed by the triangle inequality. When extending a network, however, there are no constraints in choosing weights. This flexibility comes at the cost of the inconsistent embedding described above. For strong coherence, however, one has that if $\delta_X(x, x_n) = \delta_{f(X)}(f(x), f(x_n)) \rightarrow 0$, then $\delta_Y(f(x), f(x_n)) \rightarrow 0$ as well. This property is used below to generalize the familiar result that isometric embeddings of metric spaces are continuous.

Proposition 2.5.7 *Let $(X, \omega_X), (Y, \omega_Y)$ be strongly coherent networks, and let $f : X \rightarrow Y$ be a weight-preserving map. Then f is continuous.*

Proof Let $V \subseteq Y$ be open. We need to show that $U := f^{-1}(V)$ is open. By first-countability, it suffices to show that any sequence converging to a point of U is




Fig. 10 “Isometrically” embedding into a larger space is tricky for networks. Let δ denote the canonical pseudometric obtained via the Kuratowski embedding as in Definition 2.5.4. Consider the embedding of Z into the metric space X . We have $\delta_X(a, b) = d_X(a, b) = 1 = \delta_Z(a, b)$. However, $\delta_Y(a, b) = 2 \neq 1 = \delta_Z(a, b)$. In Lemma 3.6.4 we show that this discrepancy can be resolved if the domain and codomain have the same motif sets

eventually inside U . Let $(x_n)_n$ be a sequence in X converging to $x \in U$. Then we have

$$\|\omega_Y(f(x_n), \cdot)|_{f(X)} - \omega_Y(f(x), \cdot)|_{f(X)}\|_\infty = \|\omega_X(x_n, \cdot) - \omega_X(x, \cdot)\|_\infty \rightarrow 0.$$

Here the first equality holds because f is weight-preserving, and the limit holds by coherence in X . Similarly we also have

$$\|\omega_Y(\cdot, f(x_n))|_{f(X)} - \omega_Y(\cdot, f(x))|_{f(X)}\|_\infty \rightarrow 0, \text{ and thus } \delta_{f(X)}(f(x), f(x_n)) \rightarrow 0.$$

By strong coherence, this property transfers to all of Y , i.e. $\delta_Y(f(x), f(x_n)) \rightarrow 0$. Thus we have $f(x_n) \rightarrow f(x)$. But then there must exist $N \in \mathbb{N}$ such that $f(x_n) \in V$ for all $n \geq N$. Then $x_n \in U$ for all $n \geq N$. Since $(x_n)_n$ was arbitrary, it follows that U is open. This concludes the proof. \square

Corollary 2.5.8 (Uniqueness) *Let (X, ω_X) be a strongly coherent network with a topology τ_X . Let τ' be another strongly coherent topology on X . Then $\tau' = \tau_X$.*

Proof The identity map is weight-preserving, and so we apply Proposition 2.5.7 to show that each topology is finer than the other. \square

Our ultimate application of coherence (Theorem 3.2.2) will be to improve a statement about weak isomorphism to a statement about strong isomorphism. In such a setting, we have equality of motif sets. It turns out that equality of motif sets allows us to prove a form of the preceding result using *only weak coherence*. We will develop this in Sect. 3.2.

2.6 Skeletons

The preceding construction naturally suggests the equivalence relation \sim defined on X as follows:

$$x \sim x' \text{ iff } \omega_X(x, z) = \omega_X(x', z) \text{ and } \omega_X(z, x) = \omega_X(z, x') \text{ for all } z \in X.$$

We refer to the process of taking equivalence classes as *passing to the skeleton*. Note that for a pseudometric d_X , one has $d_X(x, x') = 0$ iff $[x] = [x']$, i.e. $d_{X/\sim}$ is a bona fide metric. Next define $\sigma : X \rightarrow X/\sim$ to be the canonical map sending any $x \in X$ to its equivalence class $[x] \in X/\sim$. Also define $\omega_{X/\sim}([x], [x']) := \omega_X(x, x')$ for $[x], [x'] \in X/\sim$. To check that this map is well-defined, let $a, a' \in X$ be such that $a \sim x$ and $a' \sim x'$. Then,

$$\omega_X(a, a') = \omega_X(x, a') = \omega_X(x, x'),$$

where the first equality holds because $a \sim x$, and the second equality holds because $a' \sim x'$. We equip X/\sim with the quotient topology, i.e. a set is open in X/\sim if and only if its preimage under σ is open in X . Then σ is a surjective, continuous map.

Observe that when X is compact, X/\sim is the continuous image of a compact space and so is compact. In general, first countability of a topological space is *not* preserved under a surjective continuous map, but it is preserved when the surjective, continuous map is also open (Steen and Seebach 1978, p. 27). The following proposition gives a sufficient condition on X which will ensure that X/\sim is first countable.

Proposition 2.6.1 *Suppose $(X, \omega_X) \in \mathcal{N}$ has a weakly coherent topology. Then the map $\sigma : X \rightarrow X/\sim$ is an open map, i.e. it maps open sets to open sets.*

Proof of Proposition 2.6.1 Let $U \subseteq X$ be open. Because X/\sim has the quotient topology, we need to show $\sigma^{-1}(\sigma(U))$ is open. For convenience, define $V := \sigma^{-1}(\sigma(U))$. By first countability, it suffices to show that any sequence converging to a point in V is eventually inside V . Let $(v_n)_{n \in \mathbb{N}}$ be any sequence in X converging to a point $v \in V$. Note that we have $\sigma(v) = [v] = [x]$ for some $x \in U$. Because $x \sim v$, we have by weak coherence that $v_n \rightarrow x$ as well. But because $x \in U$ and U is open, there exists $N \in \mathbb{N}$ such that $v_n \in U \subseteq V$ for all $n \geq N$. Thus V is open. This concludes the proof. \square

Definition 2.6.2 (*The skeleton $\text{sk}(X)$ of a compact network*) Suppose $(X, \omega_X) \in \mathcal{CN}$ has a weakly coherent topology. The *skeleton* of X is defined to be $(\text{sk}(X), \omega_{\text{sk}(X)}) \in \mathcal{CN}$, where $\text{sk}(X) := X/\sim$, and

$$\omega_{\text{sk}(X)}([x], [x']) := \omega_X(x, x') \quad \text{for all } [x], [x'] \in \text{sk}(X).$$

Observe that $\text{sk}(X)$ is compact because X is compact, and first countable by Proposition 2.6.1 and the fact that the image of first countable space under an open, surjective, and continuous map is also first countable (Steen and Seebach 1978, p. 27). Furthermore, $\omega_{\text{sk}(X)}$ is well defined by the definition of \sim .

Theorem 2.6.3 (Skeletons are terminal) *Let $(X, \omega_X) \in \mathcal{CN}$ be a network with a weakly coherent topology. Then $(\text{sk}(X), \omega_{\text{sk}(X)}) \in \mathcal{CN}$ is terminal in $\mathbf{p}(X)$.*

The proof of this and related results are provided in Sect. 2.12.2.

2.7 The second network distance

Even though the definition of $d_{\mathcal{N}}$ is very general, in some restricted settings it may be convenient to consider a network distance that is easier to formulate. For example, in computational purposes it suffices to assume that we are computing distances between finite networks. Also, a potential reduction in computational cost is obtained if we restrict ourselves to computing distortions of bijections instead of general correspondences. The next definition (compare with Definition 2.2.7) arises from such considerations.

Definition 2.7.1 (*The second network distance $\widehat{d}_{\mathcal{N}}$*) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$ be such that $|X| = |Y|$. Then define:

$$\widehat{d}_{\mathcal{N}}(X, Y) := \frac{1}{2} \inf_{\varphi} \sup_{x, x' \in X} |\omega_X(x, x') - \omega_Y(\varphi(x), \varphi(x'))|,$$

where $\varphi : X \rightarrow Y$ ranges over all bijections from X to Y (at least one bijection exists because $|X| = |Y|$).

In analogy with the case for $d_{\mathcal{N}}$, we note that $X \cong^s Y$ implies $\widehat{d}_{\mathcal{N}}(X, Y) = 0$, and also that $\widehat{d}_{\mathcal{N}}$ satisfies symmetry and triangle inequality. In the setting of finite networks, and in contrast with $d_{\mathcal{N}}$, we have that $\widehat{d}_{\mathcal{N}}(X, Y) = 0$ also immediately implies $X \cong^s Y$. It turns out via Example 2.7.2 that $d_{\mathcal{N}}$ and $\widehat{d}_{\mathcal{N}}$ agree on networks over two nodes. However, the two notions do not agree in general. In particular, a minimal example where $d_{\mathcal{N}} \neq \widehat{d}_{\mathcal{N}}$ occurs for three node networks, as we show in Remark 2.7.3.

Example 2.7.2 (*Networks with two nodes*) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{FN}$ where $X = \{x_1, x_2\}$ and $Y = \{y_1, y_2\}$. Then we claim $d_{\mathcal{N}}(X, Y) = \widehat{d}_{\mathcal{N}}(X, Y)$. Furthermore, if $X = N_2\left(\begin{pmatrix} \alpha & \delta \\ \beta & \gamma \end{pmatrix}\right)$ and $Y = N_2\left(\begin{pmatrix} \alpha' & \delta' \\ \beta' & \gamma' \end{pmatrix}\right)$, then we have the explicit formula:

$$d_{\mathcal{N}}(X, Y) = \frac{1}{2} \min(\Gamma_1, \Gamma_2), \quad \text{where}$$

$$\Gamma_1 = \max(|\alpha - \alpha'|, |\beta - \beta'|, |\delta - \delta'|, |\gamma - \gamma'|),$$

$$\Gamma_2 = \max(|\alpha - \gamma'|, |\gamma - \alpha'|, |\delta - \beta'|, |\beta - \delta'|).$$

Details for this calculation are in Sect. 2.12.3.

Remark 2.7.3 (*A three-node example where $d_{\mathcal{N}} \neq \widehat{d}_{\mathcal{N}}$*) Assume (X, ω_X) and (Y, ω_Y) are two networks with the same cardinality. Then

$$d_{\mathcal{N}}(X, Y) \leq \widehat{d}_{\mathcal{N}}(X, Y).$$

The inequality holds because each bijection induces a correspondence, and we are minimizing over all correspondences to obtain $d_{\mathcal{N}}$. However, the inequality may be strict, as demonstrated by the following example. Let $X = \{x_1, \dots, x_3\}$ and let $Y = \{y_1, \dots, y_3\}$. Define $\omega_X(x_1, x_1) = \omega_X(x_3, x_3) = \omega_X(x_1, x_3) = 1, \omega_X = 0$

elsewhere, and define $\omega_Y(y_3, y_3) = 1, \omega_Y = 0$ elsewhere. In terms of matrices, $X = N_3(\Sigma_X)$ and $Y = N_3(\Sigma_Y)$, where

$$\Sigma_X = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \Sigma_Y = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Define $\Gamma(x, x', y, y') = |\omega_X(x, x') - \omega_Y(y, y')|$ for $x, x' \in X, y, y' \in Y$. Let φ be any bijection. Then we have:

$$\begin{aligned} \max_{x, x' \in X} \Gamma(x, x', \varphi(x), \varphi(x')) &= \max\{\Gamma(x_1, x_3, \varphi(x_1), \varphi(x_3)), \Gamma(x_1, x_1, \varphi(x_1), \varphi(x_1)), \\ &\quad \Gamma(x_3, x_3, \varphi(x_3), \varphi(x_3)), \Gamma(\varphi^{-1}(y_3), \varphi^{-1}(y_3), y_3, y_3)\} \\ &= 1. \end{aligned}$$

So $\widehat{d}_{\mathcal{N}}(X, Y) = \frac{1}{2}$. On the other hand, consider the correspondence

$$R = \{(x_1, y_3), (x_2, y_2), (x_3, y_3), (x_2, y_1)\}.$$

Then $\max_{(x, y), (x', y') \in R} |\omega_X(x, x') - \omega_Y(y, y')| = 0$. Thus $d_{\mathcal{N}}(X, Y) = 0 < \widehat{d}_{\mathcal{N}}(X, Y)$.

Example 2.7.4 (Networks with three nodes) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{FN}$, where we write $X = \{x_1, x_2, x_3\}$ and $Y = \{y_1, y_2, y_3\}$. Because we do not necessarily have $d_{\mathcal{N}} = \widehat{d}_{\mathcal{N}}$ on three node networks by Remark 2.7.3, the computation of $d_{\mathcal{N}}$ becomes more difficult than in the two node case presented in Example 2.7.2. A certain reduction is still possible, which we present next. Consider the following list \mathcal{L} of matrices representing correspondences, where a 1 in position (i, j) means that (x_i, y_j) belongs to the correspondence.

$$\begin{array}{ccccc} \begin{pmatrix} 1 & & \\ & 1 & \\ & & \end{pmatrix} & \begin{pmatrix} & 1 & \\ & & 1 \\ & & \end{pmatrix} & \begin{pmatrix} 1 & & \\ & 1 & \\ & & 1 \end{pmatrix} & \begin{pmatrix} 1 & 1 & \\ & & 1 \\ & & \end{pmatrix} & \begin{pmatrix} 1 & & \\ & 1 & 1 \\ & & \end{pmatrix} \\ \hline \begin{pmatrix} 1 & 1 & \\ & & 1 \\ & & \end{pmatrix} & \begin{pmatrix} & 1 & \\ & 1 & \\ & & 1 \end{pmatrix} & \begin{pmatrix} 1 & & \\ & 1 & 1 \\ & & \end{pmatrix} & \begin{pmatrix} 1 & 1 & \\ & 1 & \\ & & 1 \end{pmatrix} & \begin{pmatrix} 1 & 1 & \\ & 1 & 1 \\ & & \end{pmatrix} \\ \hline \begin{pmatrix} 1 & & 1 \\ & 1 & \\ & & 1 \end{pmatrix} & \begin{pmatrix} & 1 & \\ & 1 & \\ & & 1 \end{pmatrix} & \begin{pmatrix} 1 & & 1 \\ & 1 & 1 \\ & & \end{pmatrix} & \begin{pmatrix} 1 & 1 & \\ & 1 & \\ & & 1 \end{pmatrix} & \begin{pmatrix} 1 & 1 & \\ & 1 & 1 \\ & & \end{pmatrix} \end{array}$$

Now let $R \in \mathcal{R}(X, Y)$ be any correspondence. Then R contains a correspondence $S \in \mathcal{R}(X, Y)$ such that the matrix form of S is listed in \mathcal{L} . Thus $\text{dis}(R) \geq \text{dis}(S)$, since we are maximizing over a larger set. It follows that $d_{\mathcal{N}}(X, Y)$ is obtained by taking $\arg \min \frac{1}{2} \text{dis}(S)$ over all correspondences $S \in \mathcal{R}(X, Y)$ with matrix forms listed in \mathcal{L} .

For an example of this calculation, let S denote the correspondence $\{(x_1, y_1), (x_2, y_2), (x_3, y_3)\}$ represented by the matrix $\begin{pmatrix} 1 & & \\ & 1 & \\ & & 1 \end{pmatrix}$. Then $\text{dis}(S)$ is the




Fig. 11 The two networks on the left have different cardinalities, but computing correspondences shows that $d_{\mathcal{N}}(X, Y) = 1$. Similarly one computes $d_{\mathcal{N}}(X, Z) = 0$, and thus $d_{\mathcal{N}}(Y, Z) = 0$ by triangle inequality. On the other hand, the bijection given by the red arrows shows $\hat{d}_{\mathcal{N}}(Y, Z) = 1$. Applying Proposition 2.7.5 then recovers $d_{\mathcal{N}}(X, Y) = 1$

maximum among the following:

$$\begin{aligned} & |\omega_X(x_1, x_1) - \omega_Y(y_1, y_1)| \quad |\omega_X(x_1, x_2) - \omega_Y(y_1, y_2)| \quad |\omega_X(x_1, x_3) - \omega_Y(y_1, y_3)| \\ & |\omega_X(x_2, x_1) - \omega_Y(y_2, y_1)| \quad |\omega_X(x_2, x_2) - \omega_Y(y_2, y_2)| \quad |\omega_X(x_2, x_3) - \omega_Y(y_2, y_3)| \\ & |\omega_X(x_3, x_1) - \omega_Y(y_3, y_1)| \quad |\omega_X(x_3, x_2) - \omega_Y(y_3, y_2)| \quad |\omega_X(x_3, x_3) - \omega_Y(y_3, y_3)|. \end{aligned}$$

The following proposition provides an explicit connection between $d_{\mathcal{N}}$ and $\hat{d}_{\mathcal{N}}$. The proof is provided in Sect. 2.12.4, and an illustration is provided in Fig. 11.

Proposition 2.7.5 *Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$. Then,*

$$d_{\mathcal{N}}(X, Y) = \inf \left\{ \hat{d}_{\mathcal{N}}(X', Y') : X', Y' \in \mathcal{N}, X' \cong_{\Pi}^w X, Y' \cong_{\Pi}^w Y, \text{ and } |X'| = |Y'| \right\}.$$

The proof of this result is provided in Sect. 2.12.4. In words, the result states that the $d_{\mathcal{N}}$ distance between two networks X, Y may be achieved by blowing up to networks X', Y' of equal cardinality and then optimizing over bijections. Note that when restricted to $X, Y \in \mathcal{FN}$, the Type II weak isomorphism can be interchanged with Type I weak isomorphism.

Remark 2.7.6 (*Computational aspects of $d_{\mathcal{N}}$ and $\hat{d}_{\mathcal{N}}$*) The relation between $d_{\mathcal{N}}$ and $\hat{d}_{\mathcal{N}}$ in the setting of finite networks has been leveraged in Chowdhury and Needham (2020b) to compute network Fréchet means and geodesic principal components. From a computational complexity perspective, even though $\hat{d}_{\mathcal{N}}$ has a simpler formulation than $d_{\mathcal{N}}$, computing $\hat{d}_{\mathcal{N}}$ still turns out to be a hard problem (cf. Sect. 2.10).

Instead of trying to compute $d_{\mathcal{N}}$, we will focus on finding network invariants that can be computed easily. This is the content of Sect. 4. For each of these invariants, we will prove a stability result to demonstrate its validity as a proxy for $d_{\mathcal{N}}$.

2.8 Special families: dissimilarity networks and directed metric spaces

The second network distance $\hat{d}_{\mathcal{N}}$ that we introduced in the previous section turned out to be compatible with strong isomorphism. Interestingly, by narrowing down the domain of $d_{\mathcal{N}}$ to the setting of compact dissimilarity networks (cf. Definition 2.1.3), we obtain a subfamily of \mathcal{N} where $d_{\mathcal{N}}$ is compatible with strong isomorphism.

Theorem 2.8.1 (Carlsson et al. 2013) *The restriction of $d_{\mathcal{N}}$ to $\mathcal{FN}^{\text{dis}}$ is a metric modulo strong isomorphism.*

The expression for $d_{\mathcal{N}}$ was used in the context of $\mathcal{FN}^{\text{dis}}$ in Carlsson et al. (2013, 2014) to study the stability properties of hierarchical clustering methods on metric spaces and directed dissimilarity networks. That setting is considerably simpler than the situation in this paper, because in general we allow $d_{\mathcal{N}}(X, Y) = 0$ for $X, Y \in \mathcal{N}$ even when X and Y are not strongly isomorphic. In Theorem 2.8.4 below, we provide an extension of Theorem 2.8.1 to a class of compact dissimilarity networks that contains all finite dissimilarity networks.

The following definition gives a continuous relaxation of the triangle inequality that is motivated by applications to computer vision (Fagin and Stockmeyer 1998).

Definition 2.8.2 (Ψ -relaxed triangle inequality) Let $\Psi : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a continuous function such that $\Psi(0, 0) = 0$. A dissimilarity network (X, A_X) is said to satisfy a Ψ -relaxed triangle inequality if we have

$$A_X(x, x') \leq \Psi(A_X(x, x''), A_X(x', x'')) \quad \text{for all } x, x', x'' \in X.$$

This condition automatically encodes a notion of reversibility:

$$\begin{aligned} A_X(x, x') &\leq \Psi(A_X(x, x), A_X(x', x)) = \Psi(0, A_X(x', x)), \\ A_X(x', x) &\leq \Psi(A_X(x', x'), A_X(x, x')) = \Psi(0, A_X(x, x')). \end{aligned}$$

In the sequel, whenever we write “ $(X, A_X) \in \mathcal{N}^{\text{dis}}$ has a Ψ -relaxed triangle inequality” without explicit reference to a map Ψ , we mean that there exists a function $\Psi : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfying the conditions above.

Remark 2.8.3 Any finite dissimilarity network is finitely reversible and has a Ψ -relaxed triangle inequality. For example, Ψ can be taken to be a bump function that vanishes outside $\mathbb{R}_+^2 \setminus U$ —where U is some open set containing $\text{im}(A_X)$ and excluding $(0, 0)$ —and constant at $\max_{x, x' \in X} A_X(x, x')$ on U .

We end this section with a strengthening of Theorem 2.8.1 to the setting of compact networks. Recall that the interesting part of Theorem 2.8.1 was to show that $d_{\mathcal{N}}(X, Y) = 0 \implies X \cong^s Y$; we generalize this result to a certain class of *compact* dissimilarity networks. The proof is provided in Sect. 2.12.5.

Theorem 2.8.4 Let $(X, A_X), (Y, A_Y) \in \mathcal{CN}^{\text{dis}}$ be equipped with the forward-open topologies induced by A_X and A_Y , respectively. Suppose also that at least one of the two networks has a Ψ -relaxed triangle inequality. Then $d_{\mathcal{N}}(X, Y) = 0 \implies X \cong^s Y$.

Remark 2.8.5 (Generalizations of Theorem 2.8.1) For any finite dissimilarity network (X, A_X) , the discrete topology is precisely the topology induced by A_X . We have already stated before that finite dissimilarity networks trivially satisfy finite reversibility and Ψ -relaxed triangle inequality. It follows that Theorem 2.8.4 is a bona fide generalization of Theorem 2.8.1.




Fig. 12 The directed circle $(\mathbb{S}^1, \omega_{\mathbb{S}^1})$ and the directed circle $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$ with reversibility ρ , for some $\rho \in [1, \infty)$ (cf. Definition 2.1.4). The arrows show that traveling in a clockwise direction is possibly only in the directed circle with reversibility ρ . However, this incurs a penalty modulated by ρ , hence the shorter arrow in the clockwise direction

2.9 Network models: the directed circles

The collections \mathcal{N} , \mathcal{CN} , and \mathcal{FN} contain the collections of all metric spaces, compact metric spaces, and finite metric spaces, respectively. It is interesting to identify networks in these families that are not just metric spaces. Here we provide models of *directed circles*, which are infinite, asymmetric networks. See Fig. 12 for an illustration, and also Sect. 4.7 for applications of hierarchical clustering and persistent homology methods to these asymmetric network models.

Define $\mathbb{S}^1 := \{e^{i\theta} \in \mathbb{C} : \theta \in [0, 2\pi)\}$, i.e. the standard unit circle in the complex plane with the standard topology. For any $x, y \in \mathbb{S}^1$, define $\omega_{\mathbb{S}^1}(x, y)$ to be the counterclockwise geodesic distance (i.e. arc length) from x to y . Then $(\mathbb{S}^1, \omega_{\mathbb{S}^1})$ becomes a model of a directed circle.

An issue that arises now is that $\omega_{\mathbb{S}^1}$ is not continuous with respect to the standard topology on \mathbb{S}^1 . Thus $(\mathbb{S}^1, \omega_{\mathbb{S}^1})$ is a set-function pair that can be accepted as input into $d_{\mathcal{N}}$ -computations, but does not enjoy the nicer properties of \mathcal{N} . One recourse is to define $(\mathbb{S}^1, \omega_{\mathbb{S}^1})$ to be equipped with the discrete topology. This topology is first countable and makes $\omega_{\mathbb{S}^1}$ continuous, and so $(\mathbb{S}^1, \omega_{\mathbb{S}^1}) \in \mathcal{N}$.

However, the resulting network is not compact. Moreover, a coarser topology does not work to make $(\mathbb{S}^1, \omega_{\mathbb{S}^1})$ fit the framework of \mathcal{CN} . To see why, let $x \in \mathbb{S}^1$. Suppose $\omega_{\mathbb{S}^1}$ is continuous with respect to some topology on \mathbb{S}^1 . Fix $0 < \varepsilon \ll 2\pi$, and define $V := \omega_{\mathbb{S}^1}^{-1}[(-\varepsilon, \varepsilon)]$. Then V is open in the product topology, and in particular contains (x, x) . Since V is a union of open rectangles, there exists an open set $U \subseteq \mathbb{S}^1$ such that $(x, x) \in U \times U \subseteq V$. Suppose towards a contradiction that $U \neq \{x\}$. Then there exists $y \in U$, for some $y \neq x$. Because $(x, y) \in U \times U$, one has $\omega_{\mathbb{S}^1}(x, y) \in (0, \varepsilon)$. But then $\omega_{\mathbb{S}^1}(y, x) \in (2\pi - \varepsilon, 2\pi)$, which is a contradiction because $(y, x) \in U \times U \subseteq \omega_{\mathbb{S}^1}^{-1}[(-\varepsilon, \varepsilon)]$. Thus $U = \{x\}$, and hence the topology on $(\mathbb{S}^1, \omega_{\mathbb{S}^1})$ necessarily contains all singletons as open sets. Such a topology cannot be compact.

We would still like a model of a directed circle with sufficient topological regularity to belong to \mathcal{CN} . We therefore produce a family of *directed circles with finite reversibility* as follows. Fix a reversibility parameter $\rho \geq 1$ (cf. Definition 2.1.4). Then

for each $x, y \in \mathbb{S}^1$, define $\omega_{\mathbb{S}^1, \rho}$ as:

$$\omega_{\mathbb{S}^1, \rho}(x, y) := \min(\omega_{\mathbb{S}^1}(x, y), \rho \omega_{\mathbb{S}^1}(y, x)).$$

In particular, $\omega_{\mathbb{S}^1, \rho}$ has reversibility ρ . Finally, we equip \mathbb{S}^1 with the standard subspace topology generated by the open balls in \mathbb{C} . In this case, \mathbb{S}^1 is compact and first countable. It remains to check that $\omega_{\mathbb{S}^1, \rho}$ is continuous. Before proceeding to the next result, we set the notation $d_{\mathbb{S}^1}$ to denote the standard (not necessarily counterclockwise) geodesic distance on \mathbb{S}^1 . Note that $d_{\mathbb{S}^1}(x, y) = \min(\omega_{\mathbb{S}^1}(x, y), \omega_{\mathbb{S}^1}(y, x))$, and so we always have $\omega_{\mathbb{S}^1, \rho}(x, y) \leq \rho d_{\mathbb{S}^1}(x, y)$.

Proposition 2.9.1 $\omega_{\mathbb{S}^1, \rho} : \mathbb{S}^1 \times \mathbb{S}^1 \rightarrow \mathbb{R}$ is continuous.

Proof of Proposition 2.9.1 It suffices to show that the preimages of basic open sets under $\omega_{\mathbb{S}^1, \rho}$ are open. Let (a, b) be an open interval in \mathbb{R} , and let $(x, y) \in \omega_{\mathbb{S}^1, \rho}^{-1}[(a, b)]$. Set $\varepsilon := \frac{1}{2\rho} \min\{\omega_{\mathbb{S}^1, \rho}(x, y) - a, b - \omega_{\mathbb{S}^1, \rho}(x, y)\}$. Let $x' \in B(x, \varepsilon)$ and $y' \in B(y, \varepsilon)$, i.e. $d_{\mathbb{S}^1}(x, x') < \varepsilon$ and $d_{\mathbb{S}^1}(y, y') < \varepsilon$. It suffices to show $\omega_{\mathbb{S}^1, \rho}(x', y') \in \omega_{\mathbb{S}^1, \rho}^{-1}[(a, b)]$, for which, in turn, it suffices to show:

$$|\omega_{\mathbb{S}^1, \rho}(x, y) - \omega_{\mathbb{S}^1, \rho}(x', y')| < 2\rho\varepsilon = \min\{\omega_{\mathbb{S}^1, \rho}(x, y) - a, b - \omega_{\mathbb{S}^1, \rho}(x, y)\}.$$

Claim 1 Let $x, y, z \in \mathbb{S}^1$. Then $\omega_{\mathbb{S}^1, \rho}(x, z) \leq \omega_{\mathbb{S}^1, \rho}(x, y) + \omega_{\mathbb{S}^1, \rho}(y, z)$.

Assuming the claim, we have $\omega_{\mathbb{S}^1, \rho}(x, y) \leq \omega_{\mathbb{S}^1, \rho}(x, x') + \omega_{\mathbb{S}^1, \rho}(x', y') + \omega_{\mathbb{S}^1, \rho}(y', y)$ and so

$$\begin{aligned} \omega_{\mathbb{S}^1, \rho}(x, y) - \omega_{\mathbb{S}^1, \rho}(x', y') &\leq \omega_{\mathbb{S}^1, \rho}(x, x') + \omega_{\mathbb{S}^1, \rho}(y', y) \leq \rho d_{\mathbb{S}^1}(x, x') + \rho d_{\mathbb{S}^1}(y', y) \\ &= \rho(d_{\mathbb{S}^1}(x, x') + d_{\mathbb{S}^1}(y, y')). \end{aligned}$$

Similarly we have $\omega_{\mathbb{S}^1, \rho}(x', y') - \omega_{\mathbb{S}^1, \rho}(x, y) \leq \rho d_{\mathbb{S}^1}(x', x) + \rho d_{\mathbb{S}^1}(y, y') = \rho(d_{\mathbb{S}^1}(x, x') + d_{\mathbb{S}^1}(y, y'))$. We thus have $|\omega_{\mathbb{S}^1, \rho}(x, y) - \omega_{\mathbb{S}^1, \rho}(x', y')| < 2\rho\varepsilon$. Since $x' \in B(x, \varepsilon)$, $y' \in B(y, \varepsilon)$ were arbitrary, it follows that $\omega_{\mathbb{S}^1, \rho}^{-1}[(a, b)]$ is open.

To conclude the proof, we need to verify the triangle inequality in the claim. First note that for $x, y, z \in \mathbb{S}^1$, we have the simpler triangle inequality $\omega_{\mathbb{S}^1}(x, z) \leq \omega_{\mathbb{S}^1}(x, y) + \omega_{\mathbb{S}^1}(y, z)$. To see this, note that the counterclockwise arc from x to z is either equal to, or strictly included in, the union of the counterclockwise arcs from x to y and from y to z . Next we treat four cases corresponding to the ways in which $\omega_{\mathbb{S}^1, \rho}(x, y)$, $\omega_{\mathbb{S}^1, \rho}(y, z)$ may be realized.

First suppose $\omega_{\mathbb{S}^1, \rho}(x, y) = \omega_{\mathbb{S}^1}(x, y)$ and $\omega_{\mathbb{S}^1, \rho}(y, z) = \omega_{\mathbb{S}^1}(y, z)$. Then $\omega_{\mathbb{S}^1, \rho}(x, z) \leq \omega_{\mathbb{S}^1}(x, z) \leq \omega_{\mathbb{S}^1}(x, y) + \omega_{\mathbb{S}^1}(y, z) = \omega_{\mathbb{S}^1, \rho}(x, y) + \omega_{\mathbb{S}^1, \rho}(y, z)$. Here the first inequality follows by definition and the second follows by the triangle inequality for $\omega_{\mathbb{S}^1}$. Similarly, for the case $\omega_{\mathbb{S}^1, \rho}(x, y) = \rho \omega_{\mathbb{S}^1}(y, x)$ and $\omega_{\mathbb{S}^1, \rho}(y, z) = \rho \omega_{\mathbb{S}^1}(z, y)$, we have $\omega_{\mathbb{S}^1, \rho}(x, z) \leq \rho \omega_{\mathbb{S}^1}(z, x) \leq \rho(\omega_{\mathbb{S}^1}(z, y) + \omega_{\mathbb{S}^1}(y, x)) = \omega_{\mathbb{S}^1, \rho}(y, z) + \omega_{\mathbb{S}^1, \rho}(x, y)$.

Next suppose $\omega_{\mathbb{S}^1, \rho}(x, y) = \omega_{\mathbb{S}^1}(x, y)$ and $\omega_{\mathbb{S}^1, \rho}(y, z) = \rho \omega_{\mathbb{S}^1}(z, y)$. Note that either z belongs to the counterclockwise arc from x to y , or x belongs to the counterclockwise arc from z to y . In the first subcase, we have $\omega_{\mathbb{S}^1, \rho}(x, z) = \omega_{\mathbb{S}^1}(x, z) \leq \omega_{\mathbb{S}^1}(x, y) \leq \omega_{\mathbb{S}^1, \rho}(x, y) + \omega_{\mathbb{S}^1, \rho}(y, z)$. Here we use the assumption $\omega_{\mathbb{S}^1, \rho}(x, y) = \omega_{\mathbb{S}^1}(x, y)$. In the second subcase, we have $\omega_{\mathbb{S}^1, \rho}(x, z) \leq \rho \omega_{\mathbb{S}^1}(z, x) \leq \rho \omega_{\mathbb{S}^1}(z, y) \leq \omega_{\mathbb{S}^1, \rho}(y, z) + \omega_{\mathbb{S}^1, \rho}(x, y)$. Here we use the assumption $\omega_{\mathbb{S}^1, \rho}(y, z) = \rho \omega_{\mathbb{S}^1}(z, y)$.

Finally we have the case $\omega_{\mathbb{S}^1, \rho}(x, y) = \rho \omega_{\mathbb{S}^1}(y, x)$ and $\omega_{\mathbb{S}^1, \rho}(y, z) = \omega_{\mathbb{S}^1}(y, z)$. After splitting into subcases, this case proceeds analogously to the previous case. We thus prove the claim, and conclude the proof. \square

Definition 2.9.2 Let $\rho \in [1, \infty)$. We define the *directed unit circle with reversibility* ρ to be $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$. This is a compact, asymmetric network in \mathcal{CN} , specifically in $\mathcal{CN}^{\text{dis}}$. Figure 12 provides an illustration of such a network alongside $(\mathbb{S}^1, \omega_{\mathbb{S}^1})$ for comparison.

Remark 2.9.3 (*Directed circle with finite reversibility—forward-open topology version*) Instead of using the subspace topology generated by the standard topology on \mathbb{C} , we can also endow $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$ with the forward-open topology generated by $\omega_{\mathbb{S}^1, \rho}$. The open balls in this topology are precisely the open balls in the subspace topology induced by the standard topology, the only adjustment being the “center” of each ball. The directed metric space $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$ equipped with the forward-open topology is another example of a compact, asymmetric network in $\mathcal{CN}^{\text{dis}}$.

2.10 The complexity of computing $d_{\mathcal{N}}$

Computing $d_{\mathcal{N}}$ is at least as hard as computing the Gromov–Hausdorff distance between metric spaces, which has been shown by Schmiedl to not be computable in polynomial time unless $P = NP$ (Schmiedl 2017, 2015). Moreover, Schmiedl showed that even obtaining a 3-approximation of the Gromov–Hausdorff distance is not possible in polynomial time, unless $P = NP$. This result was obtained by starting with an instance of the 3-partition problem and constructing an instance of the Gromov–Hausdorff distance such that a 3-approximation of its optimal solution would solve the starting 3-partition problem.

Remark 2.10.1 For quite some time prior to Schmiedl’s proof, it was remarked (Mémoli 2012) that computing the Gromov–Hausdorff distance is reminiscent of the *quadratic bottleneck assignment problem* (QBAP) (Pardalos and Wolkowicz 1994), which is NP-hard (Burkard et al. 2009). We reproduce this argument below in the case of $d_{\mathcal{N}}$.

By Remark 2.7.3 and Proposition 2.7.5 we know that it is possible to obtain an upper bound on $d_{\mathcal{N}}$, in the case $|X| = |Y|$, by using $\widehat{d}_{\mathcal{N}}$. This problem turns out to be a case of the *quadratic bottleneck assignment problem* (QBAP) (Pardalos and Wolkowicz 1994). Let $X = \{x_1, \dots, x_n\}$ and let $Y = \{y_1, \dots, y_n\}$. Let Π denote the set of all $n \times n$ permutation matrices. Note that any $\pi \in \Pi$ can be written as $\pi = ((\pi_{ij}))_{i,j=1}^n$, where each $\pi_{ij} \in \{0, 1\}$. Then $\sum_j \pi_{ij} = 1$ for any i , and $\sum_i \pi_{ij} = 1$ for any j .

Computing $\widehat{d}_{\mathcal{N}}$ now becomes:

$$\widehat{d}_{\mathcal{N}}(X, Y) = \frac{1}{2} \min_{\pi \in \Pi} \max_{1 \leq i, k, j, l \leq n} \Gamma_{ijkl} \pi_{ij} \pi_{kl}, \quad \text{where} \\ \Gamma_{ijkl} = |\omega_X(x_i, x_k) - \omega_Y(y_j, y_l)|.$$

This is just the QBAP, which is known to be NP-hard (Burkard et al. 2009).

2.11 Comparison of $d_{\mathcal{N}}$ with the cut metric on graphs

In our work throughout this paper, we have developed the theoretical framework of a certain notion of network distance that has proven to be useful for applying methods from the topological data analysis literature to network data. In each of these applications, networks were modeled as generalizations of metric spaces, and so the appropriate notion of network distance turned out to be a generalization of the Gromov–Hausdorff distance between compact metric spaces. However, an alternative viewpoint would be to model networks as weighted, directed graphs. From this perspective, a well-known metric on the space of all graphs is the *cut metric* (Lovász 2012; Borgs et al. 2008). In particular, it is known that the completion of the space of all graphs with respect to the cut metric is compact (Lovász 2012, p. 149). An analogous result is known for the network distance, and we will establish it in a forthcoming publication. It turns out that there are other structural similarities between the cut metric and the network distance. In this section, we will develop an interpretation of an ℓ^∞ version of the cut metric in the setting of compact metric spaces, and show that it agrees with the Gromov–Hausdorff distance in this setting.

2.11.1 The cut distance between finite graphs

Let $G = (V, E)$ denote a vertex-weighted and edge-weighted graph on a vertex set $V = \{1, 2, \dots, n\}$. Let α_i denote the weight of node i , with the assumption that each $\alpha_i \geq 0$, and $\sum_i \alpha_i = 1$. Let $\beta_{ij} \in \mathbb{R}$ denote the weight of edge ij . For any $S, T \subseteq V$, define:

$$e_G(S, T) := \sum_{s \in S, t \in T} \alpha_s \alpha_t \beta_{st}.$$

Note for future reference that one may regard e_G as a function from $\text{pow}(V) \times \text{pow}(V)$ into \mathbb{R} .

Let A be an $n \times n$ matrix of real numbers. Some classical norms include the ℓ^1 norm $\|A\|_1 = n^{-2} \sum_{i,j=1}^n |A_{ij}|$, the ℓ_2 norm $\|A\|_2 = (n^{-2} \sum_{i,j=1}^n |A_{ij}|^2)^{1/2}$, and the ℓ^∞ norm $\|A\|_\infty = \max_{i,j} |A_{ij}|$. Note that the n^{-2} term is included for normalization.

The *cut norm* of A is defined as

$$\|A\|_{\square} := \frac{1}{n^2} \max_{S, T \subseteq \{1, \dots, n\}} \left| \sum_{i \in S, j \in T} A_{ij} \right|.$$

The *cut metric* or *cut distance* δ_{\square} between weighted graphs on the same node set $V = \{1, 2, \dots, n\}$ as

$$\delta_{\square}(G, G') := \max_{S, T \subseteq V} |e_G(S, T) - e_{G'}(S, T)|.$$

Next we consider weighted graphs with different numbers of nodes. Let G, G' be graphs on n and m nodes respectively, with node weights $(\alpha_i)_{i=1}^n, (\alpha'_k)_{k=1}^m$ and edge weights $(\beta_{ij})_{i,j=1}^n, (\beta'_{kl})_{k,l=1}^m$, respectively. A *fractional overlay* is a non-negative $n \times m$ matrix W such that

$$\sum_{i=1}^n W_{ik} = \alpha'_k \text{ for } 1 \leq k \leq m, \quad \text{and} \quad \sum_{k=1}^m W_{ik} = \alpha_i \text{ for } 1 \leq i \leq n.$$

Define $\mathcal{W}(G, G')$ to be the set of all fractional overlays between G and G' . Let $W \in \mathcal{W}(G, G')$. Consider the graphs $G(W), G'(W)$ on the node set $\{(i, k) : i \leq n, k \leq m, i, k \in \mathbb{N}\}$ defined in the following way: node (i, k) carries weight W_{ik} in both $G(W), G'(W)$, edge $((i, k), (j, l))$ carries weight β_{ij} in $G(W)$ and β'_{kl} in $G'(W)$. Then the cut distance d_{\square} between graphs of different sizes becomes

$$d_{\square}(G, G') := \min_{W \in \mathcal{W}(G, G')} \delta_{\square}(G(W), G'(W)). \quad (\star)$$

2.11.2 The cut distance and the Gromov–Hausdorff distance

In our interpretation, a fractional overlay is analogous to a *correspondence*, as defined in Sect. 2. We define correspondences between networks, but a similar definition can be made for metric spaces, and in the case of finite metric spaces, a correspondence can be regarded as a binary matrix. Since correspondences are used to define the Gromov–Hausdorff distance between compact metric spaces, and our definition of network distance is motivated by GH distance, we would like to reinterpret the cut distance in the setting of compact metric spaces. Our goal is to show that in this setting, a certain analogue of the cut distance agrees with the GH distance.

For any compact metric space (X, d_X) , let $\text{pow}(X)$ denote the nonempty elements of the power set of X , and let e_X be any \mathbb{R}_+ -valued function defined on $\text{pow}(X) \times \text{pow}(X)$. In analogy with the definition of e_G for graphs, one would like e_X to absorb information about the metric d_X on X .

Given a subset $R \subseteq X \times X$, let π_1 and π_2 denote the canonical projections to the X and X coordinates, respectively. Let Ξ denote the map that takes a compact metric space (X, d_X) and returns a function $e_X : \text{pow}(X) \times \text{pow}(X) \rightarrow \mathbb{R}$. We impose the following two conditions on the assignment $d_X \mapsto e_X$ induced by the map Ξ :

1. For all $x, x' \in X$,

$$\Xi(d_X)(\{x\}, \{x'\}) = d_X(x, x'). \quad (\#_1)$$

Thus $e_X = \Xi(d_X)$ and d_X agree on singleton sets.

2. For all $T, S \subseteq X \times Y$,

$$\begin{aligned} & |\Xi(d_X)(\pi_1(T), \pi_1(S)) - \Xi(d_Y)(\pi_2(T), \pi_2(S))| \\ & \leq \max_{t \in T, s \in S} |d_X(\pi_1(t), \pi_1(s)) - d_Y(\pi_2(t), \pi_2(s))|. \end{aligned} \quad (\#_2)$$

The latter of the two conditions above can be viewed as a continuity condition.

Example 2.11.1 Some natural candidates for the assignment $d_X \mapsto e_X$ which satisfy the two conditions above are:

- Ξ^H such that $e_X(A, B) = d_H^X(A, B)$, the Hausdorff distance,
- Ξ^{\max} such that $e_X(A, B) = \sup_{a \in A, b \in B} d_X(a, b)$,
- Ξ^{\min} such that $e_X(A, B) = \inf_{a \in A, b \in B} d_X(a, b)$.

It is clear that all these satisfy $(\#_1)$. In Proposition 2.11.4 we prove that they satisfy condition $(\#_2)$.

An analogue of the cut distance (\star) in the setting of compact metric spaces is the following.

Definition 2.11.2 (*An analogue of the cut distance for compact metric spaces*) Let Ξ be any map satisfying $(\#_1)$ and $(\#_2)$. Then for compact metric spaces X and Y , define

$$d_{\Xi}^{\Xi}(X, Y) := \frac{1}{2} \inf_R \sup_{S, T \subseteq R} |\Xi(d_X)(\pi_1(T), \pi_1(S)) - \Xi(d_Y)(\pi_2(T), \pi_2(S))|,$$

where R ranges over correspondences between X and Y , and π_1, π_2 are the canonical projections from $X \times Y$ onto X and Y respectively. We include the coefficient 2^{-1} to make comparison with d_{GH} simpler. We claim that this interpretation of d_{Ξ}^{Ξ} is identical to d_{GH} .

Note that one definition of the Gromov–Hausdorff distance for compact metric spaces (Burago et al. 2001, Theorem 7.3.25) is the following:

$$d_{GH}(X, Y) = \frac{1}{2} \inf_{R \in \mathcal{R}} \sup_{(x, y), (x', y') \in R} |d_X(x, x') - d_Y(y, y')|.$$

We also make the following definitions for *distortion*:

$$\begin{aligned} \text{dis}_{GH}(R) &= \sup_{(x, y), (x', y') \in R} |d_X(x, x') - d_Y(y, y')|, \\ \text{dis}_{\Xi}^{\Xi}(R) &= \sup_{S, T \subseteq R} |\Xi(d_X)(\pi_1(T), \pi_1(S)) - \Xi(d_Y)(\pi_2(T), \pi_2(S))|. \end{aligned}$$

Proposition 2.11.3 *For all compact metric spaces X, Y and any assignment Ξ satisfying $(\#_1)$ and $(\#_2)$ above, one has $d_{GH}(X, Y) = d_{\Xi}^{\Xi}(X, Y)$.*

Proof of Proposition 2.11.3 Write $e_X = \Xi(d_X)$ and $e_Y = \Xi(d_Y)$. Let $R \in \mathcal{R}(X, Y)$. In computing $\text{dis}_{\text{GH}}^{\Xi}(R)$, we take the supremum over all subsets of R , including singletons. Since e_X (resp. e_Y) agrees with d_X (resp. d_Y) on singletons, it follows that $\text{dis}_{\text{GH}}(R) \leq \text{dis}_{\Xi}^{\Xi}(R)$. Thus $d_{\text{GH}} \leq d_{\Xi}^{\Xi}$.

We now need to show $d_{\Xi}^{\Xi} \leq d_{\text{GH}}$.

Let $\eta > 0$ such that $d_{\text{GH}}(X, Y) < \eta$. Then, by one of the characterizations of d_{GH} (Burago et al. 2001, Chapter 7), there exists a joint-metric δ defined on $X \sqcup Y$ and a correspondence R such that $\delta(x, y) < \eta$ for all $(x, y) \in R$. In particular, δ agrees with d_X and d_Y when restricted to the appropriate spaces. Now we have

$$\begin{aligned} \text{dis}_{\Xi}^{\Xi}(R) &= \sup_{T, S \subseteq R} |e_X(\pi_1(T), \pi_1(S)) - e_Y(\pi_2(T), \pi_2(S))| \\ &\leq \sup_{T, S \subseteq R} \sup_{t \in T, s \in S} |d_X(\pi_1(t), \pi_1(s)) - d_Y(\pi_2(t), \pi_2(s))| \\ &= \sup_{T, S \subseteq R} \sup_{t \in T, s \in S} |\delta(\pi_1(t), \pi_1(s)) - \delta(\pi_2(t), \pi_2(s))| \end{aligned}$$

By the triangle inequality, we have the following for any $t \in T$ and $s \in S$:

$$\begin{aligned} |\delta(\pi_1(t), \pi_1(s)) - \delta(\pi_2(t), \pi_2(s))| &\leq |\delta(\pi_1(t), \pi_1(s)) - \delta(\pi_1(s), \pi_2(t))| \\ &\quad + |\delta(\pi_1(s), \pi_2(t)) - \delta(\pi_2(t), \pi_2(s))| \\ &\leq |\delta(\pi_1(t), \pi_2(t))| + |\delta(\pi_1(s), \pi_2(s))| \\ &< \eta + \eta = 2\eta. \end{aligned}$$

$$\text{dis}_{\Xi}^{\Xi}(R) \leq 2\eta.$$

This shows $d_{\Xi}^{\Xi}(X, Y) \leq \eta$, and so $d_{\Xi}^{\Xi}(X, Y) \leq d_{\text{GH}}(X, Y)$. So we conclude $d_{\text{GH}}(X, Y) = d_{\Xi}^{\Xi}(X, Y)$, under the assumptions made in this discussion. \square

Proposition 2.11.4 Each of the maps Ξ^{\min} , Ξ^{\max} , and Ξ^H satisfies condition (#2).

Proof We only give details for Ξ^H . The argument for the other cases is similar. Let X, Y be compact metric spaces, and let $T, S \subseteq X \times Y$. First we recall the Hausdorff distance between two closed subsets $E, F \subseteq X$:

$$d_H^X(E, F) = \max \left\{ \sup_{e \in E} \inf_{f \in F} d_X(e, f), \sup_{f \in F} \inf_{e \in E} d_X(e, f) \right\}.$$

So d_H^X between two sets in X is written as a max of two numbers a and b , and we have the general result

$$|\max(a, b) - \max(a', b')| \leq \max(|a - a'|, |b - b'|).$$

Another general result about “calculating” with sup is that $|\sup f - \sup g| \leq \sup |f - g|$ for real valued functions f and g . Both these properties are consequences of the

triangle inequality, and we use them here:

$$\begin{aligned}
 |d_H^X(\pi_1(T), \pi_1(S)) - d_H^Y(\pi_2(T), \pi_2(S))| &= |\max(a, b) - \max(a', b')| \\
 &\leq \max(|a - a'|, |b - b'|), \text{ where} \\
 a &= \sup_{t \in T} \inf_{s \in S} d_X(\pi_1(t), \pi_1(s)) \\
 a' &= \sup_{t \in T} \inf_{s \in S} d_Y(\pi_2(t), \pi_2(s)) \\
 b &= \sup_{s \in S} \inf_{t \in T} d_X(\pi_1(t), \pi_1(s)) \\
 b' &= \sup_{s \in S} \inf_{t \in T} d_Y(\pi_2(t), \pi_2(s))
 \end{aligned}$$

We consider only one of the terms $|a - a'|$; the other term can be treated similarly.

$$\begin{aligned}
 |a - a'| &= \left| \sup_{t \in T} \inf_{s \in S} d_X(\pi_1(t), \pi_1(s)) - \sup_{t \in T} \inf_{s \in S} d_Y(\pi_2(t), \pi_2(s)) \right| \\
 &\leq \sup_{t \in T} \left| \inf_{s \in S} d_X(\pi_1(t), \pi_1(s)) - \inf_{s \in S} d_Y(\pi_2(t), \pi_2(s)) \right| \\
 &\leq \sup_{t \in T, s \in S} |d_X(\pi_1(t), \pi_1(s)) - d_Y(\pi_2(t), \pi_2(s))|.
 \end{aligned}$$

The same bound holds for $|b - b'|$. Thus d_H satisfies condition (#₂), as claimed. \square

2.12 Proofs from Section 2

2.12.1 Proofs related to Theorem 2.3.7

Proposition 2.3.5 *Weak isomorphism of Types I and II both induce equivalence relations on \mathcal{N} .*

Proof of Proposition 2.3.5 The case for Type I weak isomorphism is similar to that of Type II, so we omit it. For Type II weak isomorphism, the reflexive and symmetric properties are easy to see, so we only provide details for verifying transitivity. Let $A, B, C \in \mathcal{N}$ be such that $A \cong_{\Pi}^w B$ and $B \cong_{\Pi}^w C$. Let $\varepsilon > 0$, and let P, S be sets with surjective maps $\varphi_A : P \rightarrow A$, $\varphi_B : P \rightarrow B$, $\psi_B : S \rightarrow B$, $\psi_C : S \rightarrow C$ such that:

$$\begin{aligned}
 |\omega_A(\varphi_A(p), \varphi_A(p')) - \omega_B(\varphi_B(p), \varphi_B(p'))| &< \varepsilon/2 \quad \text{for each } p, p' \in P, \text{ and} \\
 |\omega_B(\psi_B(s), \psi_B(s')) - \omega_C(\psi_C(s), \psi_C(s'))| &< \varepsilon/2 \quad \text{for each } s, s' \in S.
 \end{aligned}$$

Next define $T := \{(p, s) \in P \times S : \varphi_B(p) = \psi_B(s)\}$.

Claim 2 The projection maps $\pi_P : T \rightarrow P$ and $\pi_S : T \rightarrow S$ are surjective.

Proof Let $p \in P$. Then $\varphi_B(p) \in B$, and since $\psi_B : S \rightarrow B$ is surjective, there exists $s \in S$ such that $\psi_B(s) = \varphi_B(p)$. Thus $(p, s) \in T$, and $\pi_P(p, s) = p$. This suffices to show that $\pi_P : T \rightarrow P$ is a surjection. The case for $\pi_S : T \rightarrow S$ is similar. \square

It follows from the preceding claim that $\varphi_A \circ \pi_P : T \rightarrow A$ and $\psi_C \circ \pi_S : T \rightarrow C$ are surjective. Next let $(p, s), (p', s') \in T$. Then,

$$\begin{aligned} & |\omega_A(\varphi_A(\pi_P(p, s)), \varphi_A(\pi_P(p', s'))) - \omega_C(\psi_C(\pi_S(p, s)), \psi_C(\pi_S(p', s')))| \\ &= |\omega_A(\varphi_A(p), \varphi_A(p')) - \omega_C(\psi_C(s), \psi_C(s'))| \\ &= |\omega_A(\varphi_A(p), \varphi_A(p')) - \omega_B(\varphi_B(p), \varphi_B(p')) + \omega_B(\varphi_B(p), \varphi_B(p')) \\ &\quad - \omega_C(\psi_C(s), \psi_C(s'))| \\ &= |\omega_A(\varphi_A(p), \varphi_A(p')) - \omega_B(\varphi_B(p), \varphi_B(p')) + \omega_B(\psi_B(s), \psi_B(s')) \\ &\quad - \omega_C(\psi_C(s), \psi_C(s'))| \\ &< \varepsilon/2 + \varepsilon/2 = \varepsilon. \end{aligned}$$

Since $\varepsilon > 0$ was arbitrary, it follows that $A \cong_{\Pi}^w C$. \square

Theorem 2.3.7 $d_{\mathcal{N}}$ is a metric on \mathcal{N} modulo Type II weak isomorphism.

Proof of Theorem 2.3.7 It is clear that $d_{\mathcal{N}}(X, Y) \geq 0$. To show $d_{\mathcal{N}}(X, X) = 0$, consider the correspondence $R = \{(x, x) : x \in X\}$. Then for any $(x, x), (x', x') \in R$, we have $|\omega_X(x, x') - \omega_X(x, x')| = 0$. Thus $\text{dis}(R) = 0$ and $d_{\mathcal{N}}(X, X) = 0$.

Next we show symmetry, i.e. $d_{\mathcal{N}}(X, Y) \leq d_{\mathcal{N}}(Y, X)$ and $d_{\mathcal{N}}(Y, X) \leq d_{\mathcal{N}}(X, Y)$. The two cases are similar, so we just show the second inequality. Let $\eta > d_{\mathcal{N}}(X, Y)$. Let $R \in \mathcal{R}(X, Y)$ be such that $\text{dis}(R) < 2\eta$. Then define $\tilde{R} = \{(y, x) : (x, y) \in R\}$. Note that $\tilde{R} \in \mathcal{R}(Y, X)$. We have:

$$\begin{aligned} \text{dis}(\tilde{R}) &= \sup_{(y, x), (y', x') \in \tilde{R}} |\omega_Y(y, y') - \omega_X(x, x')| \\ &= \sup_{(x, y), (x', y') \in R} |\omega_Y(y, y') - \omega_X(x, x')| \\ &= \sup_{(x, y), (x', y') \in R} |\omega_X(x, x') - \omega_Y(y, y')| = \text{dis}(R). \end{aligned}$$

So $\text{dis}(R) = \text{dis}(\tilde{R})$. Then $d_{\mathcal{N}}(Y, X) = \frac{1}{2} \inf_{S \in \mathcal{R}(Y, X)} \text{dis}(S) \leq \frac{1}{2} \text{dis}(\tilde{R}) < \eta$. This shows $d_{\mathcal{N}}(Y, X) \leq d_{\mathcal{N}}(X, Y)$. The reverse inequality follows by a similar argument.

Next we prove the triangle inequality. Let $R \in \mathcal{R}(X, Y)$, $S \in \mathcal{R}(Y, Z)$, and let

$$R \circ S = \{(x, z) \in X \times Z \mid \exists y, (x, y) \in R, (y, z) \in S\}$$

First we claim that $R \circ S \in \mathcal{R}(X, Z)$. This is equivalent to checking that for each $x \in X$, there exists z such that $(x, z) \in R \circ S$, and for each $z \in Z$, there exists x such that $(x, z) \in R \circ S$. The proofs of these two conditions are similar, so we just prove the former. Let $x \in X$. Let $y \in Y$ be such that $(x, y) \in R$. Then there exists $z \in Z$ such that $(y, z) \in S$. Then $(x, z) \in R \circ S$.

Next we claim that $\text{dis}(R \circ S) \leq \text{dis}(R) + \text{dis}(S)$. Let $(x, z), (x', z') \in R \circ S$. Let $y \in Y$ be such that $(x, y) \in R$ and $(y, z) \in S$. Let $y' \in Y$ be such that $(x', y') \in R$, $(y', z') \in S$. Then we have:

$$|\omega_X(x, x') - \omega_Z(z, z')| = |\omega_X(x, x') - \omega_Y(y, y') + \omega_Y(y, y') - \omega_Z(z, z')|$$

$$\begin{aligned} & \leq |\omega_X(x, x') - \omega_Y(y, y')| + |\omega_Y(y, y') - \omega_Z(z, z')| \\ & \leq \text{dis}(R) + \text{dis}(S). \end{aligned}$$

This holds for any $(x, z), (x', z') \in R \circ S$, and proves the claim.

Now let $\eta_1 > d_{\mathcal{N}}(X, Y)$, let $\eta_2 > d_{\mathcal{N}}(Y, Z)$, and let $R \in \mathcal{R}(X, Y)$, $S \in \mathcal{R}(Y, Z)$ be such that $\text{dis}(R) < 2\eta_1$ and $\text{dis}(S) < 2\eta_2$. Then we have:

$$d_{\mathcal{N}}(X, Z) \leq \frac{1}{2} \text{dis}(R \circ S) \leq \frac{1}{2} \text{dis}(R) + \frac{1}{2} \text{dis}(S) < 2\eta_1 + 2\eta_2.$$

This shows that $d_{\mathcal{N}}(X, Z) \leq d_{\mathcal{N}}(X, Y) + d_{\mathcal{N}}(Y, Z)$, and proves the triangle inequality.

Finally, we claim that $X \cong_{\Pi}^w Y$ if and only if $d_{\mathcal{N}}(X, Y) = 0$. Suppose $d_{\mathcal{N}}(X, Y) = 0$. Let $\varepsilon > 0$, and let $R(\varepsilon) \in \mathcal{R}(X, Y)$ be such that $\text{dis}(R(\varepsilon)) < \varepsilon$. Then for any $z = (x, y), z' = (x', y') \in R(\varepsilon)$, we have $|\omega_X(x, x') - \omega_Y(y, y')| < \varepsilon$. But this is equivalent to writing $|\omega_X(\pi_X(z), \pi_X(z')) - \omega_Y(\pi_Y(z), \pi_Y(z'))| < \varepsilon$, where $\pi_X : R(\varepsilon) \rightarrow X$ and $\pi_Y : R(\varepsilon) \rightarrow Y$ are the canonical projection maps. This holds for each $\varepsilon > 0$. Thus $X \cong_{\Pi}^w Y$.

Conversely, suppose $X \cong_{\Pi}^w Y$, and for each $\varepsilon > 0$ let $Z(\varepsilon)$ be a set with surjective maps $\phi_X^\varepsilon : Z(\varepsilon) \rightarrow X$, $\phi_Y^\varepsilon : Z \rightarrow Y$ such that $|\omega_X(\phi_X(z), \phi_X(z')) - \omega_Y(\phi_Y(z), \phi_Y(z'))| < \varepsilon$ for all $z, z' \in Z(\varepsilon)$. For each $\varepsilon > 0$, let $R(\varepsilon) = \{(\phi_X^\varepsilon(z), \phi_Y^\varepsilon(z)) : z \in Z(\varepsilon)\}$. Then $R(\varepsilon) \in \mathcal{R}(X, Y)$ for each $\varepsilon > 0$, and $\text{dis}(R(\varepsilon)) = \sup_{z, z' \in Z} |\omega_X(\phi_X(z), \phi_X(z')) - \omega_Y(\phi_Y(z), \phi_Y(z'))| < \varepsilon$.

We conclude that $d_{\mathcal{N}}(X, Y) = 0$. Thus $d_{\mathcal{N}}$ is a metric modulo Type II weak isomorphism. \square

2.12.2 Proofs related to Theorem 2.6.3

The following lemma summarizes useful facts about weight preserving maps and the relation \sim .

Lemma 2.12.1 *Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$, and let $f : X \rightarrow Y$ be a weight preserving surjection. Then,*

1. *f preserves equivalence classes of \sim , i.e. $x \sim x'$ for $x, x' \in X$ iff $f(x) \sim f(x')$.*
2. *f preserves weights between equivalence classes, i.e. $\omega_{X/\sim}([x], [x']) = \omega_{Y/\sim}([f(x)], [f(x')])$ for any $[x], [x'] \in X/\sim$.*

Proof of Lemma 2.12.1 For the first assertion, let $x \sim x'$ for some $x, x' \in X$. We wish to show $f(x) \sim f(x')$. Let $y \in Y$, and write $y = f(z)$ for some $z \in X$. Then,

$$\begin{aligned} \omega_Y(f(x), y) &= \omega_Y(f(x), f(z)) = \omega_X(x, z) = \omega_X(x', z) \\ &= \omega_Y(f(x'), f(z)) = \omega_Y(f(x'), y). \end{aligned}$$

Similarly we have $\omega_Y(y, f(x)) = \omega_Y(y, f(x'))$ for any $y \in Y$. Thus $f(x) \sim f(x')$.

Conversely suppose $f(x) \sim f(x')$. Let $z \in X$. Then,

$$\omega_X(x, z) = \omega_Y(f(x), f(z)) = \omega_Y(f(x'), f(z)) = \omega_X(x', z),$$

and similarly we get $\omega_X(z, x) = \omega_X(z, x')$. Thus $x \sim x'$. This proves the first assertion.

The second assertion holds by definition:

$$\omega_{Y/\sim}([f(x)], [f(x')]) = \omega_Y(f(x), f(x')) = \omega_X(x, x') = \omega_{X/\sim}([x], [x']).$$

□

The following proposition shows that skeletons inherit the property of weak coherence (cf. Definition 2.5.1).

Proposition 2.12.2 *Let (X, ω_X) be a compact network with a weakly coherent topology. The quotient topology on $(\text{sk}(X), \omega_{\text{sk}(X)})$ is also weakly coherent.*

Proof of Proposition 2.12.2 Let $([x_n])_n$ be a sequence in $\text{sk}(X)$ converging to some $[x] \in \text{sk}(X)$, and let U be an open set in X containing x . By Proposition 2.6.1 we know $V := \sigma(U)$ is open. Thus all but finitely many of the terms of $([x_n])_n$ are inside V , and so all but finitely many of the terms of $(x_n)_n$ are inside U . Thus $x_n \rightarrow x$. By weak coherence, this means $\delta_X(x_n, x) \rightarrow 0$. Because σ preserves weights by construction, we then have $\delta_{\text{sk}(X)}([x_n], [x]) \rightarrow 0$.

This completes one direction of the proof. For the other direction, suppose we have $\delta_{\text{sk}(X)}([x_n], [x]) \rightarrow 0$. We need to show $[x_n] \rightarrow [x]$ in $\text{sk}(X)$.

Let V be an open set in $\text{sk}(X)$ containing $[x]$, and denote $U := \sigma^{-1}(V)$. Because σ preserves weights, we know $\delta_X(x_n, x) \rightarrow 0$. Then by coherence of X , we have $x_n \rightarrow x$. Thus $(x_n)_n$ is eventually inside U , and so $([x_n])_n$ is eventually inside V . Thus $[x_n] \rightarrow [x]$. This concludes the proof. □

In addition to weak coherence, the skeleton has the following useful property.

Proposition 2.12.3 *Let (X, ω_X) be a compact network with a weakly coherent topology. Then $(\text{sk}(X), \omega_{\text{sk}(X)})$ is Hausdorff.*

Proof of Proposition 2.12.3 Let $[x] \neq [x'] \in \text{sk}(X)$. By first countability, we take a countable open neighborhood base $\{U_n : n \in \mathbb{N}\}$ of $[x]$ such that $U_1 \supseteq U_2 \supseteq U_3 \dots$ (if necessary, we replace U_n by $\bigcap_{i=1}^n U_i$). Similarly, we take a countable open neighborhood base $\{V_n : n \in \mathbb{N}\}$ of $[x']$ such that $V_1 \supseteq V_2 \supseteq V_3 \dots$. To show that $\text{sk}(X)$ is Hausdorff, it suffices to show that there exists $n \in \mathbb{N}$ such that $U_n \cap V_n = \emptyset$.

Towards a contradiction, suppose $U_n \cap V_n \neq \emptyset$ for each $n \in \mathbb{N}$. For each $n \in \mathbb{N}$, let $[y_n] \in U_n \cap V_n$. Any open set containing $[x]$ contains U_N for some $N \in \mathbb{N}$, and thus contains $[y_n]$ for all $n \geq N$. Thus $[y_n] \rightarrow [x]$. Similarly, $[y_n] \rightarrow [x']$. Because $\text{sk}(X)$ has a weakly coherent topology (Proposition 2.12.2), we then have:

$$\begin{aligned} \|\omega_{\text{sk}(X)}([x], \cdot) - \omega_{\text{sk}(X)}([x'], \cdot)\|_\infty &\leq \|\omega_{\text{sk}(X)}([x], \cdot) - \omega_{\text{sk}(X)}([y_n], \cdot)\|_\infty \\ &\quad + \|\omega_{\text{sk}(X)}([y_n], \cdot) - \omega_{\text{sk}(X)}([x'], \cdot)\|_\infty \rightarrow 0. \end{aligned}$$

Thus $\omega_{\text{sk}(X)}([x], \cdot) = \omega_{\text{sk}(X)}([x'], \cdot)$ and similarly $\omega_{\text{sk}(X)}(\cdot, [x]) = \omega_{\text{sk}(X)}(\cdot, [x'])$. But then $x \sim x'$ and so $[x] = [x']$, a contradiction. □

We are now ready to prove that skeletons are terminal, in the sense of Definition 2.4.3 (also recall Definitions 2.4.1 and 2.4.2).

Theorem 2.6.3 (Skeletons are terminal) *Let $(X, \omega_X) \in \mathcal{CN}$ be a network with a weakly coherent topology. Then $(\text{sk}(X), \omega_{\text{sk}(X)}) \in \mathcal{CN}$ is terminal in $\mathbf{p}(X)$.*

Proof of Theorem 2.6.3 Let $Y \in \mathbf{p}(X)$. Let $f : X \rightarrow Y$ be a weight preserving surjection. We first prove that there exists a weight preserving surjection $g : Y \rightarrow \text{sk}(X)$.

Since f is surjective, for each $y \in Y$ we can write $y = f(x_y)$ for some $x_y \in X$. Then define $g : Y \rightarrow \text{sk}(X)$ by $g(y) := [x_y]$.

To see that g is surjective, let $[x] \in \text{sk}(X)$. Write $y = f(x)$. Then there exists $x_y \in X$ such that $f(x_y) = y$ and $g(y) = [x_y]$. Since f preserves equivalence classes (Lemma 2.12.1) and $f(x_y) = f(x)$, we have $x \sim x_y$. Thus $[x_y] = [x]$, and so $g(y) = [x]$.

To see that g preserves weights, let $y, y' \in Y$. Then,

$$\begin{aligned}\omega_Y(y, y') &= \omega_Y(f(x_y), f(x_{y'})) = \omega_X(x_y, x_{y'}) = \omega_{\text{sk}(X)}([x_y], [x_{y'}]) \\ &= \omega_{\text{sk}(X)}(g(y), g(y')).\end{aligned}$$

This proves that the skeleton satisfies the first condition for being terminal.

Next suppose $g : Y \rightarrow \text{sk}(X)$ and $h : Y \rightarrow \text{sk}(X)$ are two weight preserving surjections. We wish to show $h = \psi \circ g$ for some $\psi \in \text{Aut}(\text{sk}(X))$.

For each $[x] \in \text{sk}(X)$, we use the surjectivity of g to pick $y_x \in Y$ such that $g(y_x) = [x]$. Then we define $\psi : \text{sk}(X) \rightarrow \text{sk}(X)$ by $\psi([x]) = \psi(g(y_x)) := h(y_x)$.

To see that ψ is surjective, let $[x] \in \text{sk}(X)$. Since h is surjective, there exists $y'_x \in Y$ such that $h(y'_x) = [x]$. Write $[u] = g(y'_x)$. We have already chosen y_u such that $g(y_u) = [u]$. Since g preserves equivalence classes (Lemma 2.12.1), it follows that $y'_x \sim y_u$. Then,

$$\psi([u]) = \psi(g(y_u)) = h(y_u) = h(y'_x) = [x],$$

where the second-to-last equality holds because h preserves equivalence classes (Lemma 2.12.1).

To see that ψ is injective, let $[x], [x'] \in \text{sk}(X)$ be such that $\psi([x]) = h(y_x) = h(y_{x'}) = \psi([x'])$. Since h preserves equivalence classes (Lemma 2.12.1), we have $y_x \sim y_{x'}$. Next, $g(y_x) = [x]$ and $g(y_{x'}) = [x']$ by the choices we made earlier. Since $y_x \sim y_{x'}$ and g preserves clusters, we have $g(y_x) \sim g(y_{x'})$. Thus $[x] = [x']$.

Next we wish to show that ψ preserves weights. Let $[x], [x'] \in \text{sk}(X)$. Then,

$$\begin{aligned}\omega_{\text{sk}(X)}(\psi([x]), \psi([x'])) &= \omega_{\text{sk}(X)}(h(y_x), h(y_{x'})) = \omega_Y(y_x, y_{x'}) \\ &= \omega_{\text{sk}(X)}(g(y_x), g(y_{x'})) = \omega_{\text{sk}(X)}([x], [x']).\end{aligned}$$

Thus ψ is a bijective, weight preserving automorphism of $\text{sk}(X)$. Finally we wish to show that $h = \psi \circ g$. Let $y \in Y$, and write $g(y) = [x]$ for some $x \in X$. Since g preserves equivalence classes (Lemma 2.12.1), we have $y \sim y_x$, where $g(y_x) = [x]$. Then,

$$\psi(g(y)) = \psi([x]) = \psi(g(y_x)) = h(y_x) = h(y),$$

where the last equality holds because h preserves equivalence classes (Lemma 2.12.1). Thus for each $y \in Y$, we have $h(y) = \psi(g(y))$. This shows that the skeleton satisfies the second condition for being terminal. We conclude the proof. \square

2.12.3 Proof of Example 2.7.2

Example 2.7.2 (*Networks with two nodes*) Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{FN}$ where $X = \{x_1, x_2\}$ and $Y = \{y_1, y_2\}$. Then we claim $d_{\mathcal{N}}(X, Y) = \widehat{d}_{\mathcal{N}}(X, Y)$. Furthermore, if $X = N_2\left(\begin{pmatrix} \alpha & \delta \\ \beta & \gamma \end{pmatrix}\right)$ and $Y = N_2\left(\begin{pmatrix} \alpha' & \delta' \\ \beta' & \gamma' \end{pmatrix}\right)$, then we have the explicit formula:

$$\begin{aligned} d_{\mathcal{N}}(X, Y) &= \frac{1}{2} \min(\Gamma_1, \Gamma_2), \quad \text{where} \\ \Gamma_1 &= \max(|\alpha - \alpha'|, |\beta - \beta'|, |\delta - \delta'|, |\gamma - \gamma'|), \\ \Gamma_2 &= \max(|\alpha - \gamma'|, |\gamma - \alpha'|, |\delta - \beta'|, |\beta - \delta'|). \end{aligned}$$

Proof of Example 2.7.2 We start with some notation: for $x, x' \in X, y, y' \in Y$, let

$$\Gamma(x, x', y, y') = |\omega_X(x, x') - \omega_Y(y, y')|.$$

Let $\varphi : X \rightarrow Y$ be a bijection. Note that $R_\varphi := \{(x, \varphi(x)) : x \in X\}$ is a correspondence, and this holds for any bijection (actually any surjection) φ . Since we minimize over all correspondences for $d_{\mathcal{N}}$, we conclude $d_{\mathcal{N}}(X, Y) \leq \widehat{d}_{\mathcal{N}}(X, Y)$.

For the reverse inequality, we represent all the elements of $\mathcal{R}(X, Y)$ as 2-by-2 binary matrices R , where a 1 in position ij means $(x_i, y_j) \in R$. Denote the matrix representation of each $R \in \mathcal{R}(X, Y)$ by $\text{mat}(R)$, and the collection of such matrices as $\text{mat}(\mathcal{R})$. Then we have:

$$\text{mat}(\mathcal{R}) = \left\{ \begin{pmatrix} 1 & a \\ b & 1 \end{pmatrix} : a, b \in \{0, 1\} \right\} \cup \left\{ \begin{pmatrix} a & 1 \\ 1 & b \end{pmatrix} : a, b \in \{0, 1\} \right\}$$

Let $A = \{(x_1, y_1), (x_2, y_2)\}$ (in matrix notation, this is $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$) and let $B = \{(x_1, y_2), (x_2, y_1)\}$ (in matrix notation, this is $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$). Let $R \in \mathcal{R}(X, Y)$. Note that either $A \subseteq R$ or $B \subseteq R$. Suppose that $A \subseteq R$. Then we have:

$$\max_{(x, y), (x', y') \in A} \Gamma(x, x', y, y') \leq \max_{(x, y), (x', y') \in R} \Gamma(x, x', y, y')$$

Let $\Omega(A)$ denote the quantity on the left hand side. A similar result holds in the case $B \subseteq R$:

$$\max_{(x, y), (x', y') \in B} \Gamma(x, x', y, y') \leq \max_{(x, y), (x', y') \in R} \Gamma(x, x', y, y')$$

Let $\Omega(B)$ denote the quantity on the left hand side. Since either $A \subseteq R$ or $B \subseteq R$, we have

$$\min\{\Omega(A), \Omega(B)\} \leq \min_{R \in \mathcal{R}} \max_{(x, y), (x', y') \in R} \Gamma(x, x', y, y')$$

We may identify A with the bijection given by $x_1 \mapsto y_1$ and $x_2 \mapsto y_2$. Similarly we may identify B with the bijection sending $x_1 \mapsto y_2$, $x_2 \mapsto y_1$. Thus we have

$$\min_{\varphi} \max_{x, x' \in X} \Gamma(x, x', \varphi(x), \varphi(x')) \leq \min_{R \in \mathcal{R}} \max_{(x, y), (x', y') \in R} \Gamma(x, x', y, y').$$

So we have $\widehat{d}_{\mathcal{N}}(X, Y) \leq d_{\mathcal{N}}(X, Y)$. Thus $\widehat{d}_{\mathcal{N}} = d_{\mathcal{N}}$.

Next, let $\{p, q\}$ and $\{p', q'\}$ denote the vertex sets of X and Y . Consider the bijection φ given by $p \mapsto p'$, $q \mapsto q'$ and the bijection ψ given by $p \mapsto q'$, $q \mapsto p'$. Note that the weight matrix is determined by setting $\omega_X(p, p) = \alpha$, $\omega_X(p, q) = \delta$, $\omega_X(q, p) = \beta$, and $\omega_X(q, q) = \gamma$, and similarly for Y . Then we get $\text{dis}(\varphi) = \max(|\alpha - \alpha'|, |\beta - \beta'|, |\gamma - \gamma'|, |\delta - \delta'|)$ and $\text{dis}(\psi) = \max(|\alpha - \gamma'|, |\gamma - \alpha'|, |\delta - \beta'|, |\beta - \delta'|)$. The formula follows immediately. \square

2.12.4 Proof of Proposition 2.7.5

Proposition 2.7.5 *Let $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$. Then,*

$$d_{\mathcal{N}}(X, Y) = \inf \{ \widehat{d}_{\mathcal{N}}(X', Y') : X', Y' \in \mathcal{N}, X' \cong_{\Pi}^w X, Y' \cong_{\Pi}^w Y, \text{ and } |X'| = |Y'| \}.$$

Proof of Proposition 2.7.5 We begin with an observation. Given $X, Y \in \mathcal{N}$, let $X', Y' \in \mathcal{N}$ be such that $X \cong_{\Pi}^w X'$, $Y \cong_{\Pi}^w Y'$, and $|X'| = |Y'|$. Because $d_{\mathcal{N}}$ is a metric on \mathcal{N} modulo Type II weak isomorphism (Theorem 2.3.7), we have:

$$d_{\mathcal{N}}(X, Y) \leq d_{\mathcal{N}}(X, X') + d_{\mathcal{N}}(X', Y') + d_{\mathcal{N}}(Y', Y) = d_{\mathcal{N}}(X', Y') \leq \widehat{d}_{\mathcal{N}}(X', Y'),$$

where the last inequality follows from Remark 2.7.3.

Next let $\eta > d_{\mathcal{N}}(X, Y)$, and let $R \in \mathcal{R}(X, Y)$ be such that $\text{dis}(R) < 2\eta$. We wish to find networks X' and Y' such that $\widehat{d}_{\mathcal{N}}(X', Y') < \eta$. Write $Z = X \times Y$, and write $f : Z \rightarrow X$ and $g : Z \rightarrow Y$ to denote the (surjective) projection maps $(x, y) \mapsto x$ and $(x, y) \mapsto y$. Notice that we may write $R = \{(f(z), g(z)) : z \in R \subseteq Z\}$. In particular, by the definition of a correspondence, the restrictions of f, g to R are still surjective. Also notice that Z is a finite product of first countable spaces and is hence first countable. We equip R with the subspace topology, which is also first countable.

Define two weight functions $f^*\omega, g^*\omega : Z \times Z \rightarrow \mathbb{R}$ by $f^*\omega(z, z') = \omega_X(f(z), f(z'))$ and $g^*\omega(z, z') = \omega_Y(g(z), g(z'))$. Note that $f^*\omega$ is the composition of ω_X with the continuous projection $(x, y, x', y') \mapsto (x, x')$, and is thus continuous. The restriction of a continuous function to a subspace is continuous, and so $f^*\omega$ restricted to $R \times R$ is continuous. We note also that the subspace topology on $R \times R$ induced by $Z \times Z$ is just the product topology on $R \times R$. Similarly, $g^*\omega$ restricted to $R \times R$ is continuous. We now abuse notation slightly to write $f^*\omega, g^*\omega$ to mean their restrictions to $R \times R$. Let $(U, \omega_U) = (R, f^*\omega)$ and let $(V, \omega_V) = (R, g^*\omega)$. Then $U, V \in \mathcal{N}$. Note that $d_{\mathcal{N}}(X, U) = 0$ by Remark 2.3.2, because $|U| \geq |X|$ and for all $z, z' \in U$, we have $\omega_U(z, z') = f^*\omega(z, z') = \omega_X(f(z), f(z'))$ for the surjective map f . Similarly $d_{\mathcal{N}}(Y, V) = 0$.

Next let $\varphi : U \rightarrow V$ be the bijection $z \mapsto z$. Then we have:

$$\begin{aligned} \sup_{z, z' \in U} |\omega_U(z, z') - \omega_V(\varphi(z), \varphi(z'))| &= \sup_{z, z' \in U} |\omega_U(z, z') - \omega_V(z, z')| \\ &= \sup_{z, z' \in R} |\omega_X(f(z), f(z')) - \omega_Y(g(z), g(z'))| \\ &= \sup_{(x, y), (x', y') \in R} |\omega_X(x, x') - \omega_Y(y, y')| \\ &= \text{dis}(R). \text{ In particular,} \\ \inf_{\varphi: U \rightarrow V \text{ bijection}} \text{dis}(\varphi) &\leq \text{dis}(R). \end{aligned}$$

So there exist networks U, V with the same node set (and thus the same cardinality) such that $\widehat{d}_{\mathcal{N}}(U, V) \leq \frac{1}{2} \text{dis}(R) < \eta$. We have already shown that $d_{\mathcal{N}}(X, Y) \leq \widehat{d}_{\mathcal{N}}(U, V)$. Since $\eta > d_{\mathcal{N}}(X, Y)$ was arbitrary, it follows that we have:

$$d_{\mathcal{N}}(X, Y) = \inf \{ \widehat{d}_{\mathcal{N}}(X', Y') : X' \cong_{\Pi}^w X, Y' \cong_{\Pi}^w Y, \text{ and } |X'| = |Y'| \}.$$

□

2.12.5 Proof of Theorem 2.8.4

Theorem 2.8.4 *Let $(X, A_X), (Y, A_Y) \in \mathcal{CN}^{\text{dis}}$ be equipped with the forward-open topologies induced by A_X and A_Y , respectively. Suppose also that at least one of the two networks has a Ψ -relaxed triangle inequality. Then $d_{\mathcal{N}}(X, Y) = 0 \implies X \cong^s Y$.*

Proof of Theorem 2.8.4 By Theorem 3.2.1, we know that X and Y are Type I weakly isomorphic. So there exists a set V with surjections $\varphi_X : V \rightarrow X, \varphi_Y : V \rightarrow Y$ such that $A_X(\varphi_X(v), \varphi_X(v')) = A_Y(\varphi_Y(v), \varphi_Y(v'))$ for all $v, v' \in V$. Thus we obtain (not necessarily unique) maps $f : X \rightarrow Y$ and $g : Y \rightarrow X$ that are weight-preserving. Hence the composition $g \circ f : X \rightarrow X$ is a weight-preserving map. Without loss of generality, assume that X has a Ψ -relaxed triangle inequality. Recall that this means that there exists a continuous function $\Psi : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $\Psi(0, 0) = 0$ and $A_X(x, x') \leq \Psi(A_X(x, x''), A_X(x', x''))$ for all $x, x', x'' \in X$.

It is known that an isometric embedding from a compact metric space into itself must be bijective (Burago et al. 2001, Theorem 1.6.14). We now prove a similar result using the assumptions of our theorem. Let $h : X \rightarrow X$ be a weight-preserving map. By the assumption of a dissimilarity network, we know that f, g and h are injective.

We check that h is continuous, using the assumptions about the topology on X . Let $V \subseteq h(X)$ be open. Define $U := h^{-1}[V]$. We claim that U is open. Let $x \in U$, and consider $h(x) \in V$. Since V is open and the forward balls form a base for the topology, we pick $\varepsilon > 0$ such that $B^+(h(x), \varepsilon) \subseteq V$. Now let $x' \in B^+(x, \varepsilon)$. Then $A_X(h(x), h(x')) = A_X(x, x') < \varepsilon$, so $h(x') \in B^+(h(x), \varepsilon) \subseteq V$. Hence $x' \in U$. It follows that $B^+(x, \varepsilon) \subseteq U$. Hence U is open, and h is continuous.

Next we check that X is Hausdorff, using the Ψ_X -relaxed triangle inequality assumption. Let $x, x' \in X$, where $x \neq x'$. Using continuity of Ψ_X , let $\varepsilon > 0$ be such that $\Psi([0, \varepsilon), [0, \varepsilon)) \subseteq [0, A_X(x, x'))$. We wish to show that $B^+(x, \varepsilon) \cap B^+(x', \varepsilon) = \emptyset$. Towards a contradiction, suppose this is not the case and let $z \in B^+(x, \varepsilon) \cap B^+(x', \varepsilon)$. But then $A_X(x, x') \leq \Psi(A_X(x, z), A_X(x', z)) < A_X(x, x')$, a contradiction. It follows that X is Hausdorff.

Now $h(X)$ is compact, being the continuous image of a compact space, and it is closed in X because it is a compact subset of a Hausdorff space.

Finally we show that h is surjective. Towards a contradiction, suppose that the open set $X \setminus h(X)$ is nonempty, and let $x \in X \setminus h(X)$. Using the topology assumption on X , pick $\varepsilon > 0$ such that $B^+(x, \varepsilon) \subseteq X \setminus h(X)$. Define $x_0 := x$, and $x_n := h(x_{n-1})$ for each $n \in \mathbb{N}$. Then for each $n \in \mathbb{N}$, we have $A_X(x_0, x_n) \geq \varepsilon$. Since h is weight-preserving, we also have $A_X(x_k, x_{k+n}) \geq \varepsilon$ for all $k, n \in \mathbb{N}$. Since X is sequentially compact, the sequence $(x_k)_{k \geq 0}$ has a convergent subsequence $(x_{k_j})_{j \in \mathbb{N}}$ that limits to some $z \in X$. Thus $B^+(z, r)$ contains all but finitely many terms of this sequence, for any $r > 0$. Now for any $m, n \in \mathbb{N}$ we observe:

$$\begin{aligned} A_X(x_m, x_n) &\leq \Psi(A_X(x_m, z), A_X(x_n, z)), \quad \text{where} \\ A_X(x_m, z) &\leq \Psi(A_X(x_m, x_m), A_X(z, x_m)) \\ &= \Psi(0, A_X(z, x_m)), \end{aligned}$$

$$\text{and similarly } A_X(x_n, z) \leq \Psi(0, A_X(z, x_n)).$$

Since Ψ is continuous and vanishes at $(0, 0)$, we choose $\delta > 0$ such that $\Psi([0, \delta), [0, \delta)) \subseteq [0, \varepsilon)$. We also choose $\eta > 0$ such that $\Psi(0, [0, \eta)) \subseteq [0, \delta)$. Since $B^+(z, \eta)$ contains all but finitely many terms of the sequence $(x_{k_j})_{j \geq \mathbb{N}}$, we pick $N \in \mathbb{N}$ so that $x_{k_m} \in B^+(z, \eta)$, for all $m \geq N$. Let $m, n \geq N$. Then $A_X(z, x_{k_n}) < \eta$ and $A_X(z, x_{k_m}) < \eta$. Thus $A_X(x_{k_n}, z) \leq \Psi(0, A_X(z, x_{k_n})) < \delta$ and $A_X(x_{k_m}, z) \leq \Psi(0, A_X(z, x_{k_m})) < \delta$. It follows that $A_X(x_{k_m}, x_{k_n}) < \varepsilon$.

But this is a contradiction to what we have shown before. Thus h is surjective, hence bijective. Since h was an arbitrary weight-preserving map from X into itself, the same result holds for $g \circ f : X \rightarrow X$. This shows that g is surjective. It follows that $X \cong^s Y$. \square

3 Metric structure of networks

In this section we first characterize the weak isomorphism class of \mathcal{CN} via a notion of tripods. Viewing $(\mathcal{CN}/\sim, d_{\mathcal{N}})$ as a bona fide metric space, we then recover results on completeness, geodesics, and compact families as can be described for the collection of compact metric spaces equipped with d_{GH} .

With a slight abuse of notation, $d_{\mathcal{N}} : \mathcal{N}/\cong_{\text{II}}^w \times \mathcal{N}/\cong_{\text{II}}^w \rightarrow \mathbb{R}_+$ is defined as follows:

$$d_{\mathcal{N}}([X], [Y]) := d_{\mathcal{N}}(X, Y), \quad \text{for each } [X], [Y] \in \mathcal{N}/\cong_{\text{II}}^w.$$

To check that $d_{\mathcal{N}}$ is well-defined on $[X], [Y] \in \mathcal{N} / \cong_{\Pi}^w$, let $X' \in [X], Y' \in [Y]$. Then:

$$d_{\mathcal{N}}([X'], [Y']) = d_{\mathcal{N}}(X', Y') = d_{\mathcal{N}}(X, Y) = d_{\mathcal{N}}([X], [Y]),$$

where the second-to-last equality follows from the triangle inequality and the observation that $d_{\mathcal{N}}(X, X') = d_{\mathcal{N}}(Y, Y') = 0$.

Our main result is that the two types of weak isomorphism coincide in the setting of compact networks. As a stepping stone towards proving this result, we explore the notion of “sampling” finite networks from compact networks.

3.1 Compact networks and finite sampling

In this section, we prove that any compact network admits an approximation by a finite network up to arbitrary precision, in the sense of $d_{\mathcal{N}}$.

Example 3.1.1 (Some compact and noncompact networks) The nonreversible and finitely reversible directed circles in Sect. 2.9 are examples of noncompact and compact asymmetric networks, respectively. The finitely reversible circle $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$ can be “approximated” up to arbitrary precision by picking n equidistant points on \mathbb{S}^1 and equipping this collection with the restriction of $\omega_{\mathbb{S}^1, \rho}$. We view this process as “sampling” finite networks from a compact network. In the next result, we present this sampling process as a theorem that applies to *any* compact network.

Definition 3.1.2 (ε -systems) Let $\varepsilon > 0$. For any network (X, ω_X) , an ε -system on X is a finite open cover $\mathcal{U} = \{U_1, \dots, U_n\}$, $n \in \mathbb{N}$, of X such that for any $1 \leq i, j \leq n$, we have $\omega_X(U_i, U_j) \subseteq B(r_{ij}, \varepsilon)$ for some $r_{ij} \in \mathbb{R}$.

In some cases, we will be interested in the situation where X is a finite union of connected components $\{X_1, \dots, X_n\}$, $n \in \mathbb{N}$. By a *refined* ε -system, we will mean an ε -system such that each element of the ε -system is contained in precisely one connected component of X .

The next result shows that by sampling points from all the elements of an ε -system, one obtains a finite, quantitatively good approximation to the underlying network.

Theorem 3.1.3 (Sampling from a compact network) *Let (X, ω_X) be a compact network. Then for any $\varepsilon > 0$,*

- *there exists a refined $\varepsilon/4$ -system \mathcal{G} on X*
- *for X' any finite subset of X having nonempty intersection with each element in a refined $\varepsilon/4$ system \mathcal{G} , one has*

$$d_{\mathcal{N}}((X, \omega_X), (X', \omega_X|_{X' \times X'})) < \varepsilon.$$

In particular, one can always choose a finite subset $X' \subseteq X$ such that $d_{\mathcal{N}}(X, X') < \varepsilon$.

The proof of this result is in Sect. 3.6.1.

Remark 3.1.4 (Compact metric spaces admit ε -systems) When considering a compact metric space (X, d_X) , the preceding theorem relates to the well-known notion of taking finite ε -nets in a metric space. Recall that for $\varepsilon > 0$, a subset $S \subseteq X$ is an ε -net if for any point $x \in X$, we have $B(x, \varepsilon) \cap S \neq \emptyset$. Such an ε -net satisfies the nice property that $d_{\text{GH}}(X, S) < \varepsilon$ (Burago et al. 2001, 7.3.11). In particular, one can find a finite ε -net of (X, d_X) for any $\varepsilon > 0$ by compactness.

Moreover, one may always find an ε -system for a compact metric space, for any $\varepsilon > 0$. To see this, let (X, d_X) be a compact metric space and let $\varepsilon > 0$. Consider the open cover $\{B(x, \varepsilon/4) : x \in X\}$. By compactness, we can take a finite subcover $\mathcal{U} = \{B(x_i, \varepsilon/4) : 1 \leq i \leq n\}$ for some $n \in \mathbb{N}$. Now fix $1 \leq i, j \leq n$. Let $x \in B(x_i, \varepsilon/4), x' \in B(x_j, \varepsilon/4)$. Then by triangle inequality, one has:

$$|d_X(x, x') - d_X(x_i, x_j)| \leq d_X(x, x_i) + d_X(x_j, x') < \varepsilon/2.$$

Consequently $d_X(x, x') \in B(d_X(x_i, x_j), \varepsilon)$. Since x, x' were arbitrary, we have $d_X(B(x_i, \varepsilon/4), B(x_j, \varepsilon/4)) \subseteq B(d_X(x_i, x_j), \varepsilon)$. Setting $r_{ij} = d_X(x_i, x_j)$ now recovers the property of being an ε -system. Theorem 3.1.3 can be viewed as showing that this property of compact metric spaces can be obtained (albeit with more difficulty) as a consequence of compactness and continuity, rather than metric axioms.

Observe that we do not make quantitative estimates on the cardinality of the ε -approximation produced in Theorem 3.1.3. In the setting of compact metric spaces, the size of an ε -net relates to the rich theory of metric entropy developed by Kolmogorov and Tihomirov (Edgar 1993, Chapter 17).

By virtue of Theorem 3.1.3, one can always approximate a compact network up to any given precision. The next result implies that a sampled network limits to the underlying compact network as the sample gets more and more dense.

Corollary 3.1.5 (Limit of dense sampling) *Let (X, ω_X) be a compact network, and let $S = \{s_1, s_2, \dots\}$ be a countable dense subset of X with a fixed enumeration. For each $n \in \mathbb{N}$, let X_n be the finite network with node set $\{s_1, \dots, s_n\}$ and weight function $\omega_X|_{X_n \times X_n}$. Then we have:*

$$d_{\mathcal{N}}(X, X_n) \downarrow 0 \text{ as } n \rightarrow \infty.$$

Proof of Corollary 3.1.5 Using the first statement of Theorem 3.1.3, let $\mathcal{G} = \{G_1, \dots, G_q\}$ be a refined $\varepsilon/4$ -system on X . By density of S , choose $p(i) \in \mathbb{N}$ for $1 \leq i \leq q$ such that $s_{p(i)} \in G_i$ for each i . Then define

$$n := \max \{p(1), p(2), \dots, p(q)\}.$$

Now define X_n to be the network with node set $\{s_1, s_2, \dots, s_n\}$ and weight function given by the appropriate restriction of ω_X . Then by the second assertion of Theorem 3.1.3, one has $d_{\mathcal{N}}(X, X_n) \leq \varepsilon$. It follows from the construction that $d_{\mathcal{N}}(X, X_n) \downarrow 0$ as $n \rightarrow \infty$. \square

3.2 Weak isomorphism in compact networks and characterization via motifs

By Theorem 2.3.7, $d_{\mathcal{N}}$ is a proper metric on \mathcal{N} modulo Type II weak isomorphism, which is equivalent to Type I weak isomorphism when restricted to \mathcal{FN} . The comparison between $\mathbb{Q} \cap [0, 1]$ and $[0, 1]$ in Example 2.3.4 shows that in general, these two notions of weak isomorphism are not equivalent. This leads to the following natural question: when restricted to \mathcal{CN} , are we still able to recover equivalence between Type I and Type II weak isomorphism?

In the following theorem, we provide a positive answer to this question.

Theorem 3.2.1 (Weak isomorphism in \mathcal{CN}) *Let $X, Y \in \mathcal{CN}$. Then X and Y are Type II weakly isomorphic if and only if X and Y are Type I weakly isomorphic, i.e. there exists a set V and surjections $\varphi_X : V \rightarrow X$, $\varphi_Y : V \rightarrow Y$ such that:*

$$\omega_X(\varphi_X(v), \varphi_X(v')) = \omega_Y(\varphi_Y(v), \varphi_Y(v')) \text{ for all } v, v' \in V.$$

The proof of this result is in Sect. 3.6.2. As a consequence of Theorem 3.2.1, we see that weak isomorphisms of Types I and II coincide in the setting of \mathcal{CN} . Thus we recover a desirable notion of equivalence in the setting of compact networks. By imposing slightly more control over topology using the notion of weak coherence, we obtain the next result, which roughly states that weakly isomorphic compact networks have strongly isomorphic skeleta.

Theorem 3.2.2 *Suppose $(X, \omega_X), (Y, \omega_Y)$ are separable, compact networks with weakly coherent topologies. Then the following are equivalent:*

1. $X \cong^w Y$.
2. $M_n(X) = M_n(Y)$ for all $n \in \mathbb{N}$.
3. $\text{sk}(X) \cong^s \text{sk}(Y)$.

The proof of this result is in Sect. 3.6.3. Assertion (3) implies (1) immediately, but motif sets form the crucial tool required to pass from (1) to (3). When proving this result, we will use a property of motif sets called *stability* to verify that (1) implies (2). This property of stability is later contextualized and proved in Theorem 4.3.6.

To contrast this result with the metric setting, note that compact metric spaces are automatically separable and weakly coherent, and that weak and strong notions of isomorphism coincide for compact metric spaces. The directed circles with finite reversibility (Sect. 2.9) form a family of non-metric spaces that satisfy the assumptions of Theorem 3.2.2.

3.3 The completeness of \mathcal{CN}/\cong^w

A very natural question regarding \mathcal{CN}/\cong^w is if it is complete. This indeed turns out to be the case. The proof of the following result is in Sect. 3.6.4.

Theorem 3.3.1 *$(\mathcal{CN}/\cong^w, d_{\mathcal{N}})$ is a complete metric space containing $(\mathcal{FN}/\cong^w, d_{\mathcal{N}})$.*

The result of Theorem 3.3.1 can be summarized as follows:

A convergent sequence of finite networks limits to a compact topological space with a continuous weight function.

Remark 3.3.2 The technique of composed correspondences used in the proof of Theorem 3.3.1 can also be used to show that the collection of isometry classes of compact metric spaces endowed with the Gromov–Hausdorff distance is a complete metric space. Standard proofs of this fact (Petersen 2006, §10) do not use correspondences, relying instead on a method of endowing metrics on disjoint unions of spaces and then computing Hausdorff distances.

Completeness of \mathcal{CN}/\cong^w gives us a first useful criterion for convergence of networks. Ideally, we would also want a criterion for convergence along the lines of sequential compactness. In the setting of compact metric spaces, Gromov’s Precompactness Theorem implies that the topology induced by the Gromov–Hausdorff distance admits many *precompact families* of compact metric spaces (i.e. collections whose closure is compact) (Gromov 1981; Burago et al. 2001; Petersen 2006). Any sequence in such a precompact family has a subsequence converging to some limit point of the family. In the next section, we extend these results to the setting of networks. Namely, we show that there are many families of compact networks that are precompact under the metric topology induced by d_N .

3.4 Precompact families in \mathcal{CN}/\cong^w

We begin this section with some definitions.

Definition 3.4.1 (*Diameter for networks*, Chowdhury and Mémoli 2017) For any network (X, ω_X) , define $\text{diam}(X) := \sup_{x, x' \in X} |\omega_X(x, x')|$. For compact networks, the sup is replaced by max.

Definition 3.4.2 A family \mathcal{F} of weak isomorphism classes of compact networks is *uniformly approximable* if: (1) there exists $D \geq 0$ such that for every $[X] \in \mathcal{F}$, we have $\text{diam}(X) \leq D$, and (2) for every $\varepsilon > 0$, there exists $N(\varepsilon) \in \mathbb{N}$ such that for each $[X] \in \mathcal{F}$, there exists a finite network Y satisfying $|Y| \leq N(\varepsilon)$ and $d_N(X, Y) < \varepsilon$.

Remark 3.4.3 The preceding definition is an analogue of the definition of *uniformly totally bounded* families of compact metric spaces (Burago et al. 2001, Definition 7.4.13), which is used in formulating the precompactness result in the metric space setting. A family of compact metric spaces is said to be uniformly totally bounded if there exists $D \in \mathbb{R}_+$ such that each space has diameter bounded above by D , and for any $\varepsilon > 0$ there exists $N_\varepsilon \in \mathbb{N}$ such that each space in the family has an ε -net with cardinality bounded above by N_ε . Recall that given a metric space (X, d_X) and $\varepsilon > 0$, a subset $S \subseteq X$ is an ε -net if for any point $x \in X$, we have $B(x, \varepsilon) \cap S \neq \emptyset$. Such an ε -net satisfies the nice property that $d_{\text{GH}}(X, S) < \varepsilon$ (Burago et al. 2001, 7.3.11). Thus an ε -net is an ε -approximation of the underlying metric space in the Gromov–Hausdorff distance.

Theorem 3.4.4 *Let \mathfrak{F} be a uniformly approximable family in \mathcal{CN}/\cong^w . Then \mathfrak{F} is precompact, i.e. any sequence in \mathfrak{F} contains a subsequence that converges in \mathcal{CN}/\cong^w .*

Our proof is modeled on the proof of an analogous result for compact metric spaces proposed by Gromov (1981). We use one external fact (Chowdhury and Mémoli 2017, stability of diam): for compact networks X, Y such that $d_{\mathcal{N}}(X, Y) < \varepsilon$, we have $\text{diam}(X) \leq \text{diam}(Y) + 2\varepsilon$.

Proof of Theorem 3.4.4 Let $D \geq 0$ be such that $\text{diam}(X) \leq D$ for each $[X] \in \mathfrak{F}$. It suffices to prove that \mathfrak{F} is totally bounded, because Theorem 3.3.1 gives completeness, and these two properties together imply precompactness. Let $\varepsilon > 0$. We need to find a finite family $\mathcal{G} \subseteq \mathcal{CN}/\cong^w$ such that for every $[F] \in \mathfrak{F}$, there exists $[G] \in \mathcal{G}$ with $d_{\mathcal{N}}(F, G) < \varepsilon$. Define:

$$\mathcal{A} := \{A \in \mathcal{FN} : |A| \leq N(\varepsilon/2), d_{\mathcal{N}}(A, F) < \varepsilon/2 \text{ for some } [F] \in \mathfrak{F}\}.$$

Each element of \mathcal{A} is an $n \times n$ matrix, where $1 \leq n \leq N(\varepsilon/2)$. For each $A \in \mathcal{A}$, there exists $[F] \in \mathfrak{F}$ with $d_{\mathcal{N}}(A, F) < \varepsilon/2$, and by the fact stated above, we have $\text{diam}(A) \leq \text{diam}(F) + 2(\varepsilon/2) \leq D + \varepsilon$. Thus the matrices in \mathcal{A} have entries in $[-D - \varepsilon, D + \varepsilon]$. Let $N \gg 1$ be such that:

$$\frac{2D + 2\varepsilon}{N} < \frac{\varepsilon}{4},$$

and write the refinement of $[-D - \varepsilon, D + \varepsilon]$ into N pieces as:

$$W := \left\{ -D - \varepsilon + k \left(\frac{2D + 2\varepsilon}{N} \right) : 0 \leq k \leq N \right\}.$$

Write $\mathcal{A} = \bigsqcup_{i=1}^{N(\varepsilon/2)} \mathcal{A}_i$, where each \mathcal{A}_i consists of the $i \times i$ matrices of \mathcal{A} . For each i define:

$$\mathcal{G}_i := \left\{ (G_{pq})_{1 \leq p, q \leq i} : G_{pq} \in W \right\}, \text{ the } i \times i \text{ matrices with entries in } W.$$

Let $\mathcal{G} = \bigsqcup_{i=1}^{N(\varepsilon/2)} \mathcal{G}_i$ and note that this is a finite collection. Furthermore, for each $A_i \in \mathcal{A}_i$, there exists $G_i \in \mathcal{G}_i$ such that

$$\|A_i - G_i\|_{\infty} < \frac{\varepsilon}{4}.$$

Taking the diagonal correspondence between A_i and G_i , it follows that $d_{\mathcal{N}}(A_i, G_i) < \varepsilon/2$. Hence for any $[F] \in \mathfrak{F}$, there exists $A \in \mathcal{A}$ and $G \in \mathcal{G}$ such that

$$d_{\mathcal{N}}(F, G) \leq d_{\mathcal{N}}(F, A) + d_{\mathcal{N}}(A, G) < \varepsilon/2 + \varepsilon/2 = \varepsilon.$$

This shows that \mathfrak{F} is totally bounded, and concludes the proof. \square

3.5 Geodesic structure on \mathcal{CN}/\cong^w

Thus far, we have motivated our discussion of compact networks by viewing them as limiting objects of finite networks. By the results of the preceding section, we know that $(\mathcal{CN}/\cong^w, d_{\mathcal{N}})$ is complete and obeys a well-behaved compactness criterion. In this section, we prove that this metric space is also *geodesic*, i.e. any two compact networks can be joined by a rectifiable curve with length equal to the distance between the two networks.

Geodesic spaces can have a variety of practical implications. For example, geodesic spaces that are also complete and locally compact are *proper* (i.e. any closed, bounded subset is compact), by virtue of the Hopf–Rinow theorem (Burago et al. 2001, §2.5.3). Any probability measure with finite second moment supported on such a space has a *barycenter* (Ohta 2012, Lemma 3.2), i.e. a “center of mass”. Conceivably, such a result can be applied to a compact, geodesically convex region of $(\mathcal{CN}/\cong^w, d_{\mathcal{N}})$ to compute an “average” network from a collection of networks. Such a result is of interest in statistical inference, e.g. when one wishes to represent a noisy collection of networks by a single network. Similar results on barycenters of geodesic spaces can be found in Le Gouic and Loubes (2017) and Leustean et al. (2016). We leave a treatment of this topic from a probabilistic framework as future work, and only use this vignette to motivate the results in this section.

We begin with some definitions.

Definition 3.5.1 (*Curves and geodesics*) A curve on \mathcal{N} joining (X, ω_X) to (Y, ω_Y) is any continuous map $\gamma : [0, 1] \rightarrow \mathcal{N}$ such that $\gamma(0) = (X, \omega_X)$ and $\gamma(1) = (Y, \omega_Y)$. We will write *a curve on \mathcal{FN}* (resp. *a curve on \mathcal{CN}*) to mean that the image of γ is contained in \mathcal{FN} (resp. \mathcal{CN}). Such a curve is called a *geodesic* (Bridson and Haefliger 2011, §I.1) between X and Y if for all $s, t \in [0, 1]$ one has:

$$d_{\mathcal{N}}(\gamma(t), \gamma(s)) = |t - s| \cdot d_{\mathcal{N}}(X, Y).$$

A metric space is called a *geodesic space* if any two points can be connected by a geodesic.

The following theorem is a useful result about geodesics:

Theorem 3.5.2 (Burago et al. 2001, Theorem 2.4.16) *Let (X, d_X) be a complete metric space. If for any $x, x' \in X$ there exists a midpoint z such that $d_X(x, z) = d_X(z, y) = \frac{1}{2}d_X(x, y)$, then X is geodesic.*

The proofs of the next results are in Sect. 3.6.5. As a first step towards showing that \mathcal{CN}/\cong^w is geodesic, we show that the collection of finite networks forms a geodesic space.

Theorem 3.5.3 *The metric space $(\mathcal{FN}/\cong^w, d_{\mathcal{N}})$ is a geodesic space. More specifically, let $[X], [Y] \in (\mathcal{FN}/\cong^w, d_{\mathcal{N}})$. Then, for any $R \in \mathcal{R}^{\text{opt}}(X, Y)$, we can construct a geodesic $\gamma_R : [0, 1] \rightarrow \mathcal{FN}/\cong^w$ between $[X]$ and $[Y]$ as follows:*

$$\gamma_R(0) := [(X, \omega_X)], \gamma_R(1) := [(Y, d_Y)], \text{ and}$$

$$\gamma_R(t) := [(R, \omega_{\gamma_R(t)})] \text{ for } t \in (0, 1),$$

where for each $(x, y), (x', y') \in R$ and $t \in (0, 1)$,

$$\omega_{\gamma_R(t)}((x, y), (x', y')) := (1 - t) \cdot \omega_X(x, x') + t \cdot \omega_Y(y, y').$$

A key step in the proof of the preceding theorem is to choose an optimal correspondence between two finite networks. This may not be possible, in general, for compact networks. However, using the additional results on precompactness and completeness of \mathcal{CN}/\cong^w , we are able to obtain the desired geodesic structure in Theorem 3.5.4. The proof is similar to the one used by the authors of Ivanov et al. (2016) to prove that the metric space of isometry classes of compact metric spaces endowed with the Gromov–Hausdorff distance is geodesic.

Theorem 3.5.4 *The complete metric space $(\mathcal{CN}/\cong^w, d_{\mathcal{N}})$ is geodesic.*

Remark 3.5.5 Consider the collection of compact metric spaces endowed with the Gromov–Hausdorff distance. This collection can be viewed as a subspace of $(\mathcal{CN}/\cong^w, d_{\mathcal{N}})$. It is known (via a proof relying on Theorem 3.5.2) that this restricted metric space is geodesic (Ivanov et al. 2016). Furthermore, it was proved in Chowdhury and Mémoli (2018a) that an optimal correspondence always exists in this setting, and that such a correspondence can be used to construct explicit geodesics instead of resorting to Theorem 3.5.2. The key technique used in Chowdhury and Mémoli (2018a) was to take a convergent sequence of increasingly-optimal correspondences, use a result about compact metric spaces called Blaschke’s theorem (Burago et al. 2001, Theorem 7.3.8) to show that the limiting object is closed, and then use metric properties such as the Hausdorff distance to guarantee that this limiting object is indeed a correspondence. A priori, such techniques cannot be readily adapted to the network setting, and while one can obtain a convergent sequence of increasingly-optimal correspondences, the obstruction lies in showing that the limiting object is indeed a correspondence. This is why we use the indirect proof via Theorem 3.5.2.

Remark 3.5.6 (Branching and deviant geodesics) It is important to note that there exist geodesics in \mathcal{CN}/\cong^w that deviate from the straight-line form given by Theorem 3.5.3. Even in the setting of compact metric spaces, there exist infinite families of branching and deviant geodesics (Chowdhury and Mémoli 2018a).

3.6 Proofs from Section 3

3.6.1 Proof of Theorem 3.1.3

Theorem 3.1.3 (Sampling from a compact network) *Let (X, ω_X) be a compact network. Then for any $\varepsilon > 0$,*

- *there exists a refined $\varepsilon/4$ -system \mathcal{G} on X*

- for X' any finite subset of X having nonempty intersection with each element in a refined $\varepsilon/4$ system \mathcal{G} , one has

$$d_{\mathcal{N}}((X, \omega_X), (X', \omega_X|_{X' \times X'})) < \varepsilon.$$

In particular, one can always choose a finite subset $X' \subseteq X$ such that $d_{\mathcal{N}}(X, X') < \varepsilon$.

Proof of Theorem 3.1.3 The idea is to find a cover of X by open sets G_1, \dots, G_q and representatives $x_i \in G_i$ for each $1 \leq i \leq q$ such that whenever we have $(x, x') \in G_i \times G_j$, we know by continuity of ω_X that $|\omega_X(x, x') - \omega_X(x_i, x_j)| < \varepsilon/4$. Then we define a correspondence that associates each $x \in G_i$ to x_i , for $1 \leq i \leq q$. Such a correspondence has distortion bounded above by $\varepsilon/2$.

Let $\varepsilon > 0$. Let \mathcal{B} be a base for the topology on X .

Let $\{B(r, \varepsilon/4) : r \in \mathbb{R}\}$ be an open cover for \mathbb{R} . Then by continuity of ω_X , we get that

$$\left\{ \omega_X^{-1}[B(r, \varepsilon/4)] : r \in \mathbb{R} \right\}$$

is an open cover for $X \times X$. Each open set in this cover can be written as a union of open rectangles $U \times V$, for $U, V \in \mathcal{B}$. Thus the following set is an open cover of $X \times X$:

$$\mathcal{U} := \left\{ U \times V : U, V \in \mathcal{B}, U \times V \subseteq \omega_X^{-1}[B(r, \varepsilon/4)], r \in \mathbb{R} \right\}.$$

Claim 3 There exists a finite open cover $\mathcal{G} = \{G_1, \dots, G_q\}$ of X such that for any $1 \leq i, j \leq q$, we have $G_i \times G_j \subseteq U \times V$ for some $U \times V \in \mathcal{U}$.

Proof of Claim 3 The proof of the claim includes elements that are often used to provide a proof of the Tube Lemma (Munkres 2000, Lemma 26.8). Since $X \times X$ is compact, we take a finite subcover:

$$\mathcal{U}^f := \{U_1 \times V_1, \dots, U_n \times V_n\}, \quad \text{for some } n \in \mathbb{N}.$$

Let $x \in X$. Then we define:

$$\mathcal{U}_x^f := \{U \times V \in \mathcal{U}^f : x \in U\},$$

and write

$$\mathcal{U}_x^f = \left\{ U_{i_1}^x \times V_{i_1}^x, \dots, U_{i_{m(x)}}^x \times V_{i_{m(x)}}^x \right\}. \quad (5)$$

Here $m(x)$ is an integer depending on x , and $\{i_1, \dots, i_{m(x)}\}$ is a subset of $\{1, \dots, n\}$.

Since \mathcal{U}^f is an open cover of $X \times X$, we know that \mathcal{U}_x^f is an open cover of $\{x\} \times X$. Next define:

$$A_x := \bigcap_{k=1}^{m(x)} U_{i_k}^x. \quad (6)$$

Then A_x is open and contains x . In the literature (Munkres 2000, p. 167), the set $A_x \times X$ is called a *tube* around $\{x\} \times X$. Notice that $A_x \times X \subseteq \cup \mathcal{U}_x^f$. Since x was arbitrary in the preceding construction, we define \mathcal{U}_x^f and A_x for each $x \in X$. Then note that $\{A_x : x \in X\}$ is an open cover of X . Using compactness of X , we choose $\{s_1, \dots, s_p\} \subseteq X$, $p \in \mathbb{N}$, such that $\{A_{s_1}, \dots, A_{s_p}\}$ is a finite subcover of X .

Once again let $x \in X$, and let \mathcal{U}_x^f and A_x be defined as above. Define the following:

$$B_x := \left\{ A_x \times V_{i_k}^x : 1 \leq k \leq m(x) \right\}. \quad (7)$$

Since $x \in A_x$ and $X \subseteq \cup_{k=1}^{m(x)} V_{i_k}^x$, it follows that B_x is a cover of $\{x\} \times X$. Furthermore, since $\{A_{s_1}, \dots, A_{s_p}\}$ is a cover of X , it follows that the finite collection $\{B_{s_1}, \dots, B_{s_p}\}$ is a cover of $X \times X$.

Let $z \in X$. Since $X \subseteq \cup_{k=1}^{m(x)} V_{i_k}^x$, we pick $V_{i_k}^x$ for $1 \leq k \leq m(x)$ such that $z \in V_{i_k}^x$. Since x was arbitrary, such a choice exists for each $x \in X$. Therefore, we define:

$$C_z := \left\{ V \in \mathcal{B} : z \in V, A_{s_i} \times V \in B_{s_i} \text{ for some } 1 \leq i \leq p \right\}.$$

Since each B_{s_i} is finite and there are finitely many B_{s_i} , we know that C_z is a finite collection. Next define:

$$D_z := \bigcap_{V \in C_z} V.$$

Then D_z is open and contains z . Notice that $X \times D_z$ is a tube around $X \times \{z\}$. Next, using the fact that $\{A_{s_i} : 1 \leq i \leq p\}$ is an open cover of X , pick $A_{s_{i(z)}}$ such that $z \in A_{s_{i(z)}}$. Here $1 \leq i(z) \leq p$ is some integer depending on z . Then define

$$G_z := D_z \cap A_{s_{i(z)}}. \quad (8)$$

Then G_z is open and contains z . Since z was arbitrary, we define G_z for each $z \in X$. Then $\{G_z : z \in X\}$ is an open cover of X , and we take a finite subcover:

$$\mathcal{G} := \{G_1, \dots, G_q\}, \quad q \in \mathbb{N}.$$

Finally, we need to show that for any choice of $1 \leq i, j \leq q$, we have $G_i \times G_j \subseteq U \times V$ for some $U \times V \in \mathcal{U}$. Let $1 \leq i, j \leq q$. Note that we can write $G_i = G_w$ and $G_j = G_y$ for some $w, y \in X$. By the definition of G_w (Eq. 8), we then have $G_w \subseteq A_{s_{i(w)}}$ for some index $i(w)$ depending on w . Referring to Eqs. (5) and (6) and using the property that $\mathcal{U}_{s_{i(w)}}^f$ is a cover of $\{s_{i(w)}\} \times X$, we choose $U^{s_{i(w)}} \times V^{s_{i(w)}} \in \mathcal{U}_{s_{i(w)}}^f$ such that $V^{s_{i(w)}}$ contains y . Thus we obtain:

$$G_w \subseteq A_{s_{i(w)}} \subseteq U^{s_{i(w)}} \quad \text{for} \quad U^{s_{i(w)}} \times V^{s_{i(w)}} \in \mathcal{U}_{s_{i(w)}}^f, \quad 1 \leq i(w) \leq p,$$

From Eq. (7) we have that $A_{S_i(w)} \times V^{S_i(w)} \in B_{S_i(w)}$. Then $V^{S_i(w)} \in C_y$, and so we have:

$$G_y \subseteq D_y \subseteq V^{S_i(w)}.$$

It follows that $G_i \times G_j = G_w \times G_y \subseteq U^{S_i(w)} \times V^{S_i(w)} \in \mathcal{U}$. \square

From the definition of \mathcal{U} , it follows that \mathcal{G} is an $\varepsilon/4$ -system on X . By further partitioning \mathcal{G} against the connected components of X , we obtain a refined $\varepsilon/4$ -system. This concludes the first part of the proof.

For the next part, let \mathcal{G} be a refined $\varepsilon/4$ -system on X , and suppose X' is a finite set that intersects each element of \mathcal{G} . We first perform a disjointification step. Define:

$$\tilde{G}_1 := G_1, \tilde{G}_2 := G_2 \setminus \tilde{G}_1, \tilde{G}_3 := G_3 \setminus (\tilde{G}_1 \cup \tilde{G}_2), \dots, \tilde{G}_q := G_q \setminus \left(\bigcup_{k=1}^{q-1} \tilde{G}_k \right).$$

Next define a correspondence between X and X' as follows:

$$R := \{(x, x_i) : x \in \tilde{G}_i, 1 \leq i \leq q\}.$$

Let $(x, x_i), (x', x_j) \in R$. Then we have $(x, x'), (x_i, x_j) \in \tilde{G}_i \times \tilde{G}_j \subseteq G_i \times G_j$. By assumption we know that $G_i \times G_j \subseteq U \times V$, for some $U \times V \in \mathcal{U}$. Therefore $\omega_X(x, x'), \omega_X(x_i, x_j) \in B(r, \varepsilon/4)$ for some $r \in \mathbb{R}$. It follows that:

$$|\omega_X(x, x') - \omega_X(x_i, x_j)| < \varepsilon/2.$$

Since $(x, x_i), (x', x_j) \in R$ were arbitrary, we have $\text{dis}(R) \leq \varepsilon/2$. Hence $d_{\mathcal{N}}(X, X') \leq \varepsilon$. \square


3.6.2 Proof of Theorem 3.2.1

Theorem 3.2.1 (Weak isomorphism in \mathcal{CN}) *Let $X, Y \in \mathcal{CN}$. Then X and Y are Type II weakly isomorphic if and only if X and Y are Type I weakly isomorphic, i.e. there exists a set V and surjections $\varphi_X : V \rightarrow X$, $\varphi_Y : V \rightarrow Y$ such that:*

$$\omega_X(\varphi_X(v), \varphi_X(v')) = \omega_Y(\varphi_Y(v), \varphi_Y(v')) \text{ for all } v, v' \in V.$$

Proof of Theorem 3.2.1 By the definition of \cong_1^w , it is clear that if $X \cong_1^w Y$, then $d_{\mathcal{N}}(X, Y) = 0$, i.e. $X \cong_{\Pi}^w Y$ (cf. Theorem 2.3.7).

Conversely, suppose $d_{\mathcal{N}}(X, Y) = 0$. Our strategy is to obtain a set $Z \subseteq X \times Y$ with canonical projection maps $\pi_X : Z \rightarrow X, \pi_Y : Z \rightarrow Y$ and surjections $\psi_X : X \rightarrow \pi_X(Z), \psi_Y : Y \rightarrow \pi_Y(Z)$ as in the following diagram:



Furthermore, we will require:

$$\omega_X(\pi_X(z), \pi_X(z')) = \omega_Y(\pi_Y(z), \pi_Y(z')) \quad \text{for all } z, z' \in Z, \quad (9)$$

$$\omega_X(x, x') = \omega_X(\psi_X(x), \psi_X(x')) \quad \text{for all } x, x' \in X, \quad (10)$$

$$\omega_Y(y, y') = \omega_Y(\psi_Y(y), \psi_Y(y')) \quad \text{for all } y, y' \in Y. \quad (11)$$

As a consequence, we will obtain a chain of Type I weak isomorphisms

$$X \cong_1^w \pi_X(Z) \cong_1^w \pi_Y(Z) \cong_1^w Y.$$

Since Type I weak isomorphism is an equivalence relation (Proposition 2.3.5), it will follow that X and Y are Type I weakly isomorphic.

By applying Theorem 3.1.3, we choose sequences of finite subnetworks $\{X_n \subseteq X : n \in \mathbb{N}\}$ and $\{Y_n \subseteq Y : n \in \mathbb{N}\}$ such that $d_N(X_n, X) < 1/n$ and $d_N(Y_n, Y) < 1/n$ for each $n \in \mathbb{N}$. By the triangle inequality, $d_N(X_n, Y_n) < 2/n$ for each n .

For each $n \in \mathbb{N}$, let $T_n \in \mathcal{R}(X_n, X)$, $P_n \in \mathcal{R}(Y, Y_n)$ be such that $\text{dis}(T_n) < 2/n$ and $\text{dis}(P_n) < 2/n$. Define $\alpha_n := 4/n - \text{dis}(T_n) - \text{dis}(P_n)$, and notice that $\alpha_n \rightarrow 0$ as $n \rightarrow \infty$. Since $d_N(X, Y) = 0$ by assumption, for each $n \in \mathbb{N}$ we let $S_n \in \mathcal{R}(X, Y)$ be such that $\text{dis}(S_n) < \alpha_n$. Then,

$$\text{dis}(T_n \circ S_n \circ P_n) \leq \text{dis}(T_n) + \text{dis}(S_n) + \text{dis}(P_n) < 4/n. \quad (\text{cf. Remark 2.2.6})$$

Then for each $n \in \mathbb{N}$, we define $R_n := T_n \circ S_n \circ P_n \in \mathcal{R}(X_n, Y_n)$. By Remark 2.2.6, we know that R_n has the following expression:

$$R_n = \{(x_n, y_n) \in X_n \times Y_n : \text{there exist } \tilde{x} \in X, \tilde{y} \in Y \text{ such that } (x_n, \tilde{x}) \in T_n, \\ (\tilde{x}, \tilde{y}) \in S_n, (\tilde{y}, y_n) \in P_n\}.$$

Next define:

$$\mathcal{S} := \left\{ (\tilde{x}_n, \tilde{y}_n)_{n \in \mathbb{N}} \in (X \times Y)^{\mathbb{N}} : (\tilde{x}_n, \tilde{y}_n) \in S_n \text{ for each } n \in \mathbb{N} \right\}.$$

Since X, Y are first countable and compact, the product $X \times Y$ is also first countable and compact, hence sequentially compact. Any sequence in a sequentially compact space has a convergent subsequence, so for convenience, we replace each sequence in \mathcal{S} by a convergent subsequence. Next define:

$$Z := \{(x, y) \in X \times Y : (x, y) \text{ a limit point of some } (\tilde{x}_n, \tilde{y}_n)_{n \in \mathbb{N}} \in \mathcal{S}\}.$$

Claim 4 Z is a closed subspace of $X \times Y$. Hence it is compact and sequentially compact.

The second statement in the claim follows from the first: assuming that Z is a closed subspace of the compact space $X \times Y$, we obtain that Z is compact. Any subspace of a first countable space is first countable, so Z is also first countable. Next, observe that $\pi_X(Z)$ equipped with the subspace topology is compact, because it is a continuous

image of a compact space. It is also first countable because it is a subspace of the first countable space X . Furthermore, the restriction of ω_X to $\pi_X(Z)$ is continuous. Thus $\pi_X(Z)$ equipped with the restriction of ω_X is a compact network, and by similar reasoning, we get that $\pi_Y(Z)$ equipped with the restriction of ω_Y is also a compact network.

Proof of Claim 4 We will show that $Z \subseteq X \times Y$ contains all its limit points. Let $(x, y) \in X \times Y$ be a limit point of Z . Let $\{U_n \subseteq X : n \in \mathbb{N}, (x, y) \in U_n\}$ be a countable neighborhood base of (x, y) . For each $n \in \mathbb{N}$, the finite intersection $V_n := \bigcap_{i=1}^n U_i$ is an open neighborhood of (x, y) , and thus contains a point $(x_n, y_n) \in Z$ that is distinct from (x, y) (by the definition of a limit point). Pick such an (x_n, y_n) for each $n \in \mathbb{N}$. Then $(x_n, y_n)_{n \in \mathbb{N}}$ is a sequence in Z converging to (x, y) such that $(x_n, y_n) \in V_n$ for each $n \in \mathbb{N}$.

For each $n \in \mathbb{N}$, note that because $(x_n, y_n) \in Z$ and V_n is an open neighborhood of (x_n, y_n) , there exists a sequence in \mathcal{S} converging to (x_n, y_n) for which all but finitely many terms are contained in V_n . So for each $n \in \mathbb{N}$, let $(\tilde{x}_n, \tilde{y}_n) \in \mathcal{S}$ be such that $(\tilde{x}_n, \tilde{y}_n) \in V_n$. Then the sequence $(\tilde{x}_n, \tilde{y}_n)_{n \in \mathbb{N}} \in \mathcal{S}$ converges to (x, y) . Thus $(x, y) \in Z$. Since (x, y) was an arbitrary limit point of Z , it follows that Z is closed. \square

Proof of Equation 9. We now prove Eq. (9). Let $z = (x, y)$, $z' = (x', y') \in Z$, and let $(\tilde{x}_n, \tilde{y}_n)_{n \in \mathbb{N}}, (\tilde{x}'_n, \tilde{y}'_n)_{n \in \mathbb{N}}$ be elements of \mathcal{S} that converge to $(x, y), (x', y')$ respectively. We wish to show $|\omega_X(x, x') - \omega_Y(y, y')| = 0$. Let $\varepsilon > 0$, and observe that:

$$\begin{aligned} & |\omega_X(x, x') - \omega_Y(y, y')| \\ &= |\omega_X(x, x') - \omega_X(\tilde{x}_n, \tilde{x}'_n) + \omega_X(\tilde{x}_n, \tilde{x}'_n) - \omega_Y(\tilde{y}_n, \tilde{y}'_n) \\ &\quad + \omega_Y(\tilde{y}_n, \tilde{y}'_n) - \omega_Y(y, y')| \\ &\leq |\omega_X(x, x') - \omega_X(\tilde{x}_n, \tilde{x}'_n)| + |\omega_X(\tilde{x}_n, \tilde{x}'_n) - \omega_Y(\tilde{y}_n, \tilde{y}'_n)| \\ &\quad + |\omega_Y(\tilde{y}_n, \tilde{y}'_n) - \omega_Y(y, y')|. \end{aligned}$$

Claim 5 Suppose we are given sequences $(\tilde{x}_n, \tilde{y}_n)_{n \in \mathbb{N}}, (\tilde{x}'_n, \tilde{y}'_n)_{n \in \mathbb{N}}$ in Z converging to (x, y) and (x', y') in Z , respectively. Then there exists $N \in \mathbb{N}$ such that for all $n \geq N$, we have:

$$|\omega_X(x, x') - \omega_X(\tilde{x}_n, \tilde{x}'_n)| < \varepsilon/4, \quad |\omega_Y(\tilde{y}_n, \tilde{y}'_n) - \omega_Y(y, y')| < \varepsilon/4.$$

Proof of Claim 5 Write $a := \omega_X(x, x')$, $b := \omega_Y(y, y')$. Since ω_X, ω_Y are continuous, we know that $\omega_X^{-1}[B(a, \varepsilon/4)]$ and $\omega_Y^{-1}[B(b, \varepsilon/4)]$ are open neighborhoods of (x, x') and (y, y') . Since each open set in the product space $X \times X$ is a union of open rectangles of the form $A \times A'$ for A, A' open subsets of X , we choose an open set $A \times A' \subseteq \omega_X^{-1}[B(a, \varepsilon/4)]$ such that $(x, x') \in A \times A'$. Similarly, we choose an open set $B \times B' \subseteq \omega_Y^{-1}[B(b, \varepsilon/4)]$ such that $(y, y') \in B \times B'$. Then $A \times B, A' \times B'$ are open neighborhoods of $(x, y), (x', y')$ respectively. Since $(\tilde{x}_n, \tilde{y}_n)_{n \in \mathbb{N}}$ and $(\tilde{x}'_n, \tilde{y}'_n)_{n \in \mathbb{N}}$ converge to (x, y) and (x', y') , respectively, we choose $N \in \mathbb{N}$ such that for all $n \geq N$, we have $(\tilde{x}_n, \tilde{y}_n) \in A \times B$ and $(\tilde{x}'_n, \tilde{y}'_n) \in A' \times B'$. The claim now follows. \square

Now choose $N \in \mathbb{N}$ such that the property in Claim 5 is satisfied, as well as the additional property that $8/N < \varepsilon/4$. Then for any $n \geq N$, we have:

$$|\omega_X(x, x') - \omega_Y(y, y')| \leq \varepsilon/4 + |\omega_X(\tilde{x}_n, \tilde{x}'_n) - \omega_Y(\tilde{y}_n, \tilde{y}'_n)| + \varepsilon/4.$$

Separately note that for each $n \in \mathbb{N}$, having $(\tilde{x}_n, \tilde{y}_n), (\tilde{x}'_n, \tilde{y}'_n) \in S_n$ implies that there exist (x_n, y_n) and $(x'_n, y'_n) \in R_n$ such that $(x_n, \tilde{x}_n), (x'_n, \tilde{x}'_n) \in T_n$ and $(\tilde{y}_n, y_n), (\tilde{y}'_n, y'_n) \in P_n$. Thus we can bound the middle term above as follows:

$$\begin{aligned} & |\omega_X(\tilde{x}_n, \tilde{x}'_n) - \omega_Y(\tilde{y}_n, \tilde{y}'_n)| \\ &= |\omega_X(\tilde{x}_n, \tilde{x}'_n) - \omega_X(x_n, x'_n) + \omega_X(x_n, x'_n) - \omega_Y(y_n, y'_n) \\ &\quad + \omega_Y(y_n, y'_n) - \omega_Y(\tilde{y}_n, \tilde{y}'_n)| \\ &\leq |\omega_X(\tilde{x}_n, \tilde{x}'_n) - \omega_X(x_n, x'_n)| + |\omega_X(x_n, x'_n) - \omega_Y(y_n, y'_n)| \\ &\quad + |\omega_Y(y_n, y'_n) - \omega_Y(\tilde{y}_n, \tilde{y}'_n)| \\ &\leq \text{dis}(T_n) + \text{dis}(R_n) + \text{dis}(P_n) < 8/n \leq 8/N < \varepsilon/4. \end{aligned}$$

The preceding calculations show that:

$$|\omega_X(x, x') - \omega_Y(y, y')| < \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, it follows that $\omega_X(x, x') = \omega_Y(y, y')$. This proves Eq. 9.

It remains to define surjective maps $\psi_X : X \rightarrow \pi_X(Z)$, $\psi_Y : Y \rightarrow \pi_Y(Z)$ and to verify Eqs. (10) and (11). Both cases are similar, so we only show the details of constructing ψ_X and verifying Eq. (10).

Construction of ψ_X . Let $x \in X$. Suppose first that $x \in \pi_X(Z)$. Then we simply define $\psi_X(x) = x$. We also make the following observation, to be used later: for each $n \in \mathbb{N}$, letting $y \in Y$ be such that $(x, y) \in S_n$, there exists $x_n \in X_n$ and $y_n \in Y_n$ such that $(x_n, x) \in T_n$ and $(y, y_n) \in P_n$.

Next suppose $x \in X \setminus \pi_X(Z)$. For each $n \in \mathbb{N}$, let $x_n \in X_n$ be such that $(x_n, x) \in T_n$, and let $\tilde{x}_n \in X$ be such that $(x_n, \tilde{x}_n) \in T_n$. Also for each $n \in \mathbb{N}$, let $\tilde{y}_n \in Y$ be such that $(\tilde{x}_n, \tilde{y}_n) \in S_n$. Then for each $n \in \mathbb{N}$, let $y_n \in Y_n$ be such that $(\tilde{y}_n, y_n) \in P_n$. Then by sequential compactness of $X \times Y$, the sequence $(\tilde{x}_n, \tilde{y}_n)_{n \in \mathbb{N}}$ has a convergent subsequence which belongs to \mathcal{S} and converges to a point $(\tilde{x}, \tilde{y}) \in Z$. In particular, we obtain a sequence $(\tilde{x}_n)_{n \in \mathbb{N}}$ converging to a point \tilde{x} , such that (x_n, x) and $(x_n, \tilde{x}_n) \in T_n$ for each $n \in \mathbb{N}$. Define $\psi_X(x) = \tilde{x}$.

Since $x \in X$ was arbitrary, this construction defines $\psi_X : X \rightarrow \pi_X(Z)$. Note that ψ_X is simply the identity on $\pi_X(Z)$, hence is surjective.

Proof of Equation 10. Now we verify Eq. (10). Let $\varepsilon > 0$. There are three cases to check:

Case 1: $x, x' \in \pi_X(Z)$ In this case, we have:

$$|\omega_X(x, x') - \omega_X(\psi_X(x), \psi_X(x'))| = \omega_X(x, x') - \omega_X(x, x') = 0.$$

Case 2: $x, x' \in X \setminus \pi_X(Z)$ By continuity of ω_X , we obtain an open neighborhood $U := \omega_X^{-1}[B(\omega_X(\psi_X(x), \psi_X(x')), \varepsilon/2)]$ of (x, x') . By the definition of

ψ_X on $X \setminus \pi_X(Z)$, we obtain sequences $(\tilde{x}_n, \tilde{y}_n)_{n \in \mathbb{N}}$ and $(\tilde{x}'_n, \tilde{y}'_n)_{n \in \mathbb{N}}$ in S converging to $(\psi_X(x), \tilde{y})$ and $(\psi_X(x'), \tilde{y}')$ for some $\tilde{y}, \tilde{y}' \in Y$. By applying Claim 5, we obtain $N \in \mathbb{N}$ such that for all $n \geq N$, we have $(\tilde{x}_n, \tilde{x}'_n) \in U$. Note that we also obtain sequences $(x_n)_{n \in \mathbb{N}}$ and $(x'_n)_{n \in \mathbb{N}}$ such that $(x_n, x), (x_n, \tilde{x}_n) \in T_n$ and $(x'_n, x'), (x'_n, \tilde{x}'_n) \in T_n$. Choose N large enough so that it satisfies the property above and also that $4/N < \varepsilon/2$. Then for any $n \geq N$,

$$\begin{aligned} & |\omega_X(x, x') - \omega_X(\psi_X(x), \psi_X(x'))| \\ &= |\omega_X(x, x') - \omega_X(x_n, x'_n) + \omega_X(x_n, x'_n) - \omega_X(\tilde{x}_n, \tilde{x}'_n) + \omega_X(\tilde{x}_n, \tilde{x}'_n) \\ &\quad - \omega_X(\psi_X(x), \psi_X(x'))| \\ &\leq \text{dis}(T_n) + \text{dis}(T_n) + \varepsilon/2 < 4/n + \varepsilon/2 \leq 4/N + \varepsilon/2 < \varepsilon. \end{aligned}$$

Case 3: $x \in \pi_X(Z), x' \in X \setminus \pi_X(Z)$ By the definition of ψ_X on $X \setminus \pi_X(Z)$, we obtain: (1) a sequence $(\tilde{x}'_n)_{n \in \mathbb{N}}$ converging to $\psi_X(x')$, and (2) another sequence $(x'_n)_{n \in \mathbb{N}}$ such that (x'_n, x') and (x'_n, \tilde{x}'_n) both belong to T_n , for each $n \in \mathbb{N}$. By the definition of ψ_X on $\pi_X(Z)$, we obtain a sequence $(x_n)_{n \in \mathbb{N}}$ such that $(x_n, x) \in T_n$ for each $n \in \mathbb{N}$.

Let $U := \omega_X^{-1}[B(\omega_X(x, \psi_X(x')), \varepsilon/2)]$. Since $(\tilde{x}'_n)_{n \in \mathbb{N}}$ converges to $\psi_X(x')$, we know that all but finitely many terms of the sequence $(x, \tilde{x}'_n)_{n \in \mathbb{N}}$ belong to U . So we choose N large enough so that for each $n \geq N$, we have:

$$\begin{aligned} & |\omega_X(x, x') - \omega_X(x, \psi_X(x'))| \\ &= |\omega_X(x, x') - \omega_X(x_n, x'_n) + \omega_X(x_n, x'_n) - \omega_X(x, \tilde{x}'_n) + \omega_X(x, \tilde{x}'_n) \\ &\quad - \omega_X(x, \psi_X(x'))| \\ &\leq \text{dis}(T_n) + \text{dis}(T_n) + \varepsilon/2 < 4/n + \varepsilon/2 \leq 4/N + \varepsilon/2 < \varepsilon. \end{aligned}$$

Since $\varepsilon > 0$ was arbitrary, Eq. (10) follows. The construction of ψ_Y and proof for Eq. (11) are similar. This concludes the proof of the theorem. \square

3.6.3 Proofs related to Theorem 3.2.2

Our goal in this section is to prove that weak isomorphism, equality of motif sets, and strong isomorphism between skeleta are equivalent in the setting of compact networks with coherent topologies. Because the main result of this section will assume equality of motif sets, we will specifically only use the weak form of coherence (cf. Definition 2.5.1).

Proposition 3.6.1 *Let $(X, \omega_X), (Y, \omega_Y)$ be compact networks such that $M_n(X) = M_n(Y)$ for all $n \in \mathbb{N}$. Suppose X contains a countable subset S_X . Then there exists a weight-preserving map $f : S_X \rightarrow Y$.*

Proof of Proposition 3.6.1 We proceed via a diagonal argument. Write $S_X = \{x_1, x_2, \dots, x_n, \dots\}$. For each $n \in \mathbb{N}$, let $f_n : S_X \rightarrow Y$ be a map that preserves weights on $\{x_1, \dots, x_n\}$. Such a map exists by the assumption that $M_n(X) = M_n(Y)$.

Since Y is first countable and compact, hence sequentially compact, the sequence $(f_n(x_1))_n$ has a convergent subsequence; we write this as $(f_{1,n}(x_1))_n$. Since f_k is weight-preserving on $\{x_1, x_2\}$ for $k \geq 2$, we know that $f_{1,n}$ is weight-preserving on $\{x_1, x_2\}$ for $n \geq 2$. Using sequential compactness again, we have that $(f_{1,n}(x_2))_n$ has a convergent subsequence $(f_{2,n}(x_2))_n$. This sequence converges at both x_1 and x_2 , and $f_{2,n}$ is weight-preserving on $\{x_1, x_2\}$ for $n \geq 2$. Proceeding in this way, we obtain the diagonal sequence $(f_{n,n})_n$ which converges pointwise on S_X . Furthermore, for any $n \in \mathbb{N}$, $f_{k,k}$ is weight-preserving on $\{x_1, \dots, x_n\}$ for $k \geq n$.

Next define $f : S_X \rightarrow Y$ by setting $f(x) := \lim_n f_{n,n}(x)$ for each $x \in S_X$. It remains to show that f is weight-preserving. Let $x_n, x_m \in S_X$, and let $k \geq \max(m, n)$. Then $\omega_X(x_n, x_m) = \omega_Y(f_{k,k}(x_n), f_{k,k}(x_m))$. Using (sequential) continuity of ω_Y , we then have:

$$\begin{aligned} \omega_Y(f(x_n), f(x_m)) &= \omega_Y\left(\lim_k f_{k,k}(x_n), \lim_k f_{k,k}(x_m)\right) \\ &= \lim_k \omega_Y(f_{k,k}(x_n), f_{k,k}(x_m)) = \omega_X(x_n, x_m). \end{aligned}$$

In the second equality above, we used the fact that a sequence converges in the product topology iff the components converge. Since $x_n, x_m \in S_X$ were arbitrary, this concludes the proof. \square

Proposition 3.6.2 *Let $(X, \omega_X), (Y, \omega_Y)$ be compact networks. Suppose $f : S_X \rightarrow Y$ is a weight-preserving function defined on a countable dense subset $S_X \subseteq X$. Then f extends to a weight-preserving map on X .*

Proof of Proposition 3.6.2 Let $x \in X \setminus S_X$. By first countability, we take a countable neighborhood base $\{U_n : n \in \mathbb{N}\}$ of x such that $U_1 \supseteq U_2 \supseteq U_3 \dots$ (if necessary, we replace U_n by $\cap_{i=1}^n U_i$). For each $n \in \mathbb{N}$, let $x_n \in U_n \cap S_X$. Then $x_n \rightarrow x$. To see this, let U be any open set containing x . Then $U_n \subseteq U$ for some $n \in \mathbb{N}$, and so $x_k \in U_n \subseteq U$ for all $k \geq n$.

Because Y is compact and first countable, hence sequentially compact, the sequence $(f(x_n))_n$ has a convergent subsequence; let y be its limit. Define $f(x) = y$. Extend f to all of X this way.

We need to verify that f is weight-preserving. Let $x, x' \in X$. Invoking the definition of f , let $(x_n)_n, (x'_n)_n$ be sequences in S_X converging to x, x' such that $f(x_n) \rightarrow f(x)$ and $f(x'_n) \rightarrow f(x')$. By sequential continuity and the standard result that a sequence converges in the product topology iff the components converge, we have

$$\lim_n \omega_Y(f(x_n), f(x'_n)) = \omega_Y(f(x), f(x')); \quad \lim_n \omega_X(x_n, x'_n) = \omega_X(x, x').$$

Let $\varepsilon > 0$. By the previous observation, fix $N \in \mathbb{N}$ such that for all $n \geq N$, we have $|\omega_Y(f(x_n), f(x'_n)) - \omega_Y(f(x), f(x'))| < \varepsilon$ and $|\omega_X(x_n, x'_n) - \omega_X(x, x')| < \varepsilon$. Then,

$$\begin{aligned} |\omega_X(x, x') - \omega_Y(f(x), f(x'))| &= |\omega_X(x, x') - \omega_X(x_n, x'_n) + \omega_X(x_n, x'_n) \\ &\quad - \omega_Y(f(x), f(x'))| \end{aligned}$$

$$\begin{aligned} &\leq |\omega_X(x, x') - \omega_X(x_n, x'_n)| + |\omega_X(f(x_n), f(x'_n)) \\ &\quad - \omega_Y(f(x), f(x'))| < 2\varepsilon. \end{aligned}$$

Thus $\omega_X(x, x') = \omega_Y(f(x), f(x'))$. Since $x, x' \in X$ were arbitrary, this concludes the proof. \square

Definition 3.6.3 Let $(X, \omega_X), (Y, \omega_Y)$ be networks. We say that a map $f : X \rightarrow Y$ is l^∞ -intrinsic if we have $\delta_Y(f(x), f(x')) = \delta_X(x, x')$ for all $x, x' \in X$, i.e.

$$\begin{aligned} &\max \left(\|\omega_X(x, \cdot) - \omega_X(x', \cdot)\|_\infty, \|\omega_X(\cdot, x) - \omega_X(\cdot, x')\|_\infty \right) \\ &= \max \left(\|\omega_Y(f(x), \cdot) - \omega_Y(f(x'), \cdot)\|_\infty, \|\omega_Y(\cdot, f(x)) - \omega_Y(\cdot, f(x'))\|_\infty \right). \end{aligned}$$

Recall that a topological space is *separable* if it contains a countable dense subset.

Lemma 3.6.4 Let $(X, \omega_X), (Y, \omega_Y)$ be separable, compact, weakly coherent networks, and suppose $M_n(X) = M_n(Y)$ for all $n \in \mathbb{N}$. Let $f : X \rightarrow Y$ be a weight-preserving map. Then f is l^∞ -intrinsic.

The interpretation of this lemma is that when two networks have the same motif sets, they contain the same information, and thus a mapping between them does not add any extrinsic data.

Proof of Lemma 3.6.4 Let $x, x' \in X$. Because f is weight-preserving, we know

$$\|\omega_Y(f(x), \cdot) - \omega_Y(f(x'), \cdot)\|_\infty \geq \|\omega_X(x, \cdot) - \omega_X(x', \cdot)\|_\infty.$$

First we wish to show that this is actually an equality. Towards a contradiction, suppose the inequality above is strict. Let $y \in Y$ be such that

$$|\omega_Y(f(x), y) - \omega_Y(f(x'), y)| > \|\omega_X(x, \cdot) - \omega_X(x', \cdot)\|_\infty.$$

Using separability, fix a countable dense subset $\{x_1, x_2, \dots\} \subseteq X$. Next, for each $n \in \mathbb{N}$, define T_n to be the weight matrix obtained from $(y, f(x), f(x'), f(x_1), f(x_2), \dots, f(x_n))$, i.e. the corresponding motif matrix. Because all the motif sets are equal, for each $T_n \in M_{n+3}(Y)$ we have an equivalent copy in $M_{n+3}(X)$.

Because we do not maintain node labels when creating motif sets, it may be the case that even though T_n is a submatrix of T_{n+1} , the points chosen to realize T_n in X are different from the points in X used to realize T_{n+1} . Specifically, let (z^1, a^1, b^1, p_1^1) denote the points in X realizing T_1 , let $(z^2, a^2, b^2, p_1^2, p_2^2)$ denote the points in X realizing T_2 , and more generally, let $(z^n, a^n, b^n, p_1^n, p_2^n, \dots, p_n^n)$ denote the points in X realizing T_n for each $n \in \mathbb{N}$.

By sequential compactness and a diagonal argument as used in Proposition 3.6.1, we pass to a subsequence such that z^n, a^n, b^n , and p_i^n for $i \leq n$ all converge as $n \rightarrow \infty$. Denote the limits by z, a, b , and $\{p_i\}_{i \in \mathbb{N}}$, respectively. Crucially we have $\omega_X(a^n, z^n) = \omega_Y(f(x), y)$, $\omega_X(b^n, z^n) = \omega_Y(f(x'), y)$, and $\omega_X(a^n, b^n) = \omega_Y(f(x), f(x')) = \omega_X(x, x')$ for all $n \in \mathbb{N}$. We also have $\omega_X(a^n, p_i^n) = \omega_Y(f(x), f(x_i)) = \omega_X(x, x_i)$

for all $n \geq i$, and likewise for other combinations of z^n, a^n, b^n , and $\{p_i^n\}_{i \in \mathbb{N}, n \geq i}$. Consequently we have the same equalities for the limits of these sequences of points.

Next define a map $g : \{a, b, p_1, p_2, \dots\} \rightarrow X$ as $a \mapsto x, b \mapsto x'$, and $p_i \mapsto x_i$ for each $i \in \mathbb{N}$. Then g is a weight-preserving map defined on a countable dense subset of X , so by Proposition 3.6.2 it extends to a weight-preserving map $g : X \rightarrow X$. In particular we have $\omega_X(x, g(z)) = \omega_X(a, z) = \omega_Y(f(x), y)$ and $\omega_X(x', g(z)) = \omega_X(b, z) = \omega_Y(f(x'), y)$. Thus we have

$$\begin{aligned} \|\omega_X(x, \cdot) - \omega_X(x', \cdot)\|_\infty &\geq |\omega_X(x, g(z)) - \omega_X(x', g(z))| \\ &= |\omega_Y(f(x), y) - \omega_Y(f(x'), y)|, \end{aligned}$$

which is a contradiction. Similarly one proves

$$\|\omega_Y(\cdot, f(x)) - \omega_Y(\cdot, f(x'))\|_\infty = \|\omega_X(\cdot, x) - \omega_X(\cdot, x')\|_\infty.$$

Because $x, x' \in X$ were arbitrary, this concludes the proof. \square

The following proposition mirrors the familiar notion that isometric maps between metric spaces are continuous. As an immediate application, we will obtain *uniqueness* of the coherent topology, thus justifying its construction.

Proposition 3.6.5 *Let $(X, \omega_X), (Y, \omega_Y)$ be separable, compact, weakly coherent networks, and suppose $M_n(X) = M_n(Y)$ for all $n \in \mathbb{N}$. Let $f : X \rightarrow Y$ be a weight-preserving map. Then f is continuous.*

Proof Let $V \subseteq Y$ be open. We need to show that $U := f^{-1}(V)$ is open. By first-countability, it suffices to show that any sequence converging to a point of U is eventually inside U . Let $(x_n)_n$ be a sequence in X converging to $x \in U$. Because f is l^∞ -intrinsic (Lemma 3.6.4), we have

$$\|\omega_Y(f(x_n), \cdot) - \omega_Y(f(x), \cdot)\|_\infty = \|\omega_X(x_n, \cdot) - \omega_X(x, \cdot)\|_\infty \rightarrow 0.$$

Here the limit holds because X is coherent. Thus we have $f(x_n) \rightarrow f(x)$. But then there must exist $N \in \mathbb{N}$ such that $f(x_n) \in V$ for all $n \geq N$. Then $x_n \in U$ for all $n \geq N$. Since $(x_n)_n$ was arbitrary, it follows that U is open. This concludes the proof. \square

Corollary 3.6.6 (Uniqueness) *Let (X, ω_X) be a network with a separable, compact, weakly coherent topology τ_X . Let τ' be another separable, compact, coherent topology on X . Then $\tau' = \tau_X$.*

Proof The identity map $\text{id} : (X, \omega_X, \tau_X) \rightarrow (X, \omega_X, \tau')$ is weight-preserving, and since both the domain and codomain have the same motif set, we know by Proposition 3.6.5 that id is continuous. Thus τ_X is finer than τ' . By the same argument, we have that τ' is finer than τ_X . Thus the two topologies are equal. \square

The next result generalizes the result that an isometric embedding of a compact metric space into itself is automatically surjective (Burago et al. 2001, Theorem 1.6.14). However, before presenting the theorem we first discuss an auxiliary construction that is used in its proof.

Definition 3.6.7 (A variant of δ_X (cf. Eq. (3))) Let (X, ω_X) be any network. For any subset $A \subseteq X$, define $\delta_X^A : X \times X \rightarrow \mathbb{R}_+$ by

$$\delta_X^A(x, x') := \max \left(\sup_{a \in A} |\omega_X(x, a) - \omega_X(x', a)|, \sup_{a \in A} |\omega_X(a, x) - \omega_X(a, x')| \right).$$

Then δ_X is a special case of δ_X^A where $A = X$. More generally, A acts as a set of “observers”, and δ_X^A compares two points x, x' with respect to this observation set.

Next, for any $E \subseteq X$ and any $y \in X$, define $\delta_X^A(y, E) := \inf_{y' \in E} \delta_X^A(y, y')$. Then $\delta_X^A(\cdot, E)$ behaves as a proxy for the “distance to a set” function, where the set is fixed to be E and the “distance” is measured relative to the observation set A .

Theorem 3.6.8 Let (X, ω_X) be a compact network with a weakly coherent, Hausdorff topology. Suppose $f : X \rightarrow X$ is a weight-preserving map. Then f is surjective.

Proof of Theorem 3.6.8 Towards a contradiction, suppose $f(X) \neq X$. By Proposition 3.6.5, f is continuous. Define $X_0 := X$, and $X_n := f(X_{n-1})$ for each $n \in \mathbb{N}$. The continuous image of a compact space is compact, and compact subspaces of a Hausdorff space are closed. Thus we obtain a decreasing sequence of nonempty compact sets $X_0 \supseteq X_1 \supseteq X_2 \supseteq \dots$. Then $Z := \bigcap_{n \in \mathbb{N}} X_n$ is nonempty and compact, hence closed.

We now break up the proof up into several claims.

Claim 6 $f(Z) = Z$.

To see this, first note that $f(\bigcap_{n \in \mathbb{N}} X_n) \subseteq \bigcap_{n \in \mathbb{N}} f(X_n) \subseteq Z$. Next let $v \in Z$. For each $n \in \mathbb{N}$, let $u_n \in X_n$ be such that $f(u_n) = v$. Since singletons in a Hausdorff space are closed, we know that $\{v\}$ is closed. By continuity, it follows that $f^{-1}(\{v\})$ is closed.

By sequential compactness, the sequence $(u_n)_n$ has a convergent subsequence that converges to some limit u . Since each $u_n \in f^{-1}(\{v\})$ and a closed set contains its limit points, we then have $u \in f^{-1}(\{v\})$. Thus $f(u) = v$, and $v \in f(Z)$. Hence $Z = f(Z)$. This proves the claim.

Let $x \in X_0 \setminus X_1$. Define $x_0 := x$, and for each $n \in \mathbb{N}$, define $x_n := f(x_{n-1})$. Then $(x_n)_n$ is a sequence in the sequentially compact space X , and so it has a convergent subsequence $(x_{n_k})_k$. Let z be the limit of this subsequence.

Claim 7 $z \in Z$.

To see this, suppose towards a contradiction that $z \notin Z$. Then there exists $N \in \mathbb{N}$ such that $z \notin X_N$. Since X_N is closed, we have that $X \setminus X_N$ is open. By the definition of convergence, $X \setminus X_N$ contains all but finitely many terms of the sequence $(x_{n_k})_k$. But each x_{n_k} belongs to X_{n_k} , which is a subset of X_N for sufficiently large k . Thus infinitely many terms of the sequence $(x_{n_k})_k$ belong to X_N , a contradiction. Hence $z \in Z$.

Now we invoke the δ_X^\bullet construction as in Definition 3.6.7.

Claim 8 For any $E \subseteq X$ and any $y \in E$,

$$\delta_X^E(y, Z) = \delta_X^{f(E)}(f(y), f(Z)).$$

To see this claim, fix $y \in E$. Let $v \in f(Z)$. Then $v = f(y')$ for some $y' \in Z$, and $\delta_X^{f(E)}(f(y), v) = \delta_X^E(y, y')$. To see the latter assertion, let $u \in f(E)$; then $u = f(y'')$ for some $y'' \in E$. Because f is weight-preserving, we then have:

$$\begin{aligned} |\omega_X(f(y), u) - \omega_X(v, u)| &= |\omega_X(f(y), f(y'')) - \omega_X(f(y'), f(y''))| \\ &= |\omega_X(y, y'') - \omega_X(y', y'')|, \\ |\omega_X(u, f(y)) - \omega_X(u, v)| &= |\omega_X(f(y''), f(y)) - \omega_X(f(y''), f(y'))| \\ &= |\omega_X(y'', y) - \omega_X(y'', y')|. \end{aligned}$$

The preceding equalities show that for each $v \in f(Z)$, there exists $y' \in Z$ such that $\delta_X^{f(E)}(f(y), v) = \delta_X^E(y, y')$. Conversely, for any $y' \in Z$, we have $\delta_X^{f(E)}(f(y), f(y')) = \delta_X^E(y, y')$. It follows that $\delta_X^{f(E)}(f(y), f(Z)) = \delta_X^E(y, Z)$.

Claim 9 $\delta_X^X(x, Z) = 0$.

To see this, assume towards a contradiction that $\delta_X^X(x, Z) = \varepsilon > 0$ (δ_X^X is non-negative by definition). Since $f(Z) = Z$, we have by the preceding claim that $\delta_X^X(x, Z) = \delta_X^{f(X)}(f(x), Z) = \dots = \delta_X^{f^n(X)}(f^n(x), Z)$ for each $n \in \mathbb{N}$. In particular, for any $k \in \mathbb{N}$,

$$\varepsilon = \delta_X^{f^{nk}(X)}(f^{nk}(x), Z) \leq \delta_X^{f^{nk}(X)}(f^{nk}(x), z) \leq \delta_X^X(f^{nk}(x), z).$$

Here the first inequality follows because the left hand side includes an infimum over $z \in Z$, and the second inequality holds because the right hand side includes a supremum over a larger set.

Since $x_{n_k} \rightarrow z$, we have that $\delta_X^X(f^{nk}(x), z) = \delta_X^X(x^{n_k}, z) = \delta_X(x_{n_k}, z) \rightarrow 0$, where the limit holds by the definition of weak coherence (cf. Definition 2.5.1). This contradiction proves the claim.

Recall that by assumption, $x \notin Z$. For each $n \in \mathbb{N}$, let $z_n \in Z$ be such that $\delta_X^X(x, z_n) < 1/n$. Then the sequence $(z_n)_n$ converges to x by weak coherence. Hence any open set containing x also contains infinitely many points of Z that are distinct from x . Thus x is a limit point of the closed set Z , and so $x \in Z$. This is a contradiction. \square

Theorem 3.2.2 Suppose $(X, \omega_X), (Y, \omega_Y)$ are separable, compact networks with weakly coherent topologies. Then the following are equivalent:

1. $X \cong^w Y$.
2. $M_n(X) = M_n(Y)$ for all $n \in \mathbb{N}$.
3. $\text{sk}(X) \cong^s \text{sk}(Y)$.

Proof of Theorem 3.2.2 (2) follows from (1) by the stability of motif sets (Theorem 4.3.6). (1) follows from (3) by the triangle inequality of $d_{\mathcal{N}}$. We need to show that (2) implies (3).

First observe that $\text{sk}(X)$, being a continuous image of the separable space X , is separable, and likewise for $\text{sk}(Y)$. Let S_X, S_Y denote countable dense subsets of $\text{sk}(X)$ and $\text{sk}(Y)$. Next, because $d_{\mathcal{N}}(X, \text{sk}(X)) = 0$, an application of Theorem 4.3.6 shows that $M_n(X) = M_n(\text{sk}(X))$ for each $n \in \mathbb{N}$. The analogous result holds for $\text{sk}(Y)$. Thus $M_n(\text{sk}(X)) = M_n(\text{sk}(Y))$ for each $n \in \mathbb{N}$. Since X and Y have weakly coherent topologies, so do $\text{sk}(X)$ and $\text{sk}(Y)$, by Proposition 2.12.2. By Propositions 3.6.1 and 3.6.2, there exist weight-preserving maps $\varphi : \text{sk}(X) \rightarrow \text{sk}(Y)$ and $\psi : \text{sk}(Y) \rightarrow \text{sk}(X)$. Define $X^{(1)} := \psi(\text{sk}(Y))$ and $Y^{(1)} := \varphi(\text{sk}(X))$. Also define φ_1 and ψ_1 to be the restrictions of φ and ψ to $X^{(1)}$ and $Y^{(1)}$, respectively. Finally define $X^{(2)} := \psi_1(Y^{(1)})$ and $Y^{(2)} := \varphi_1(X^{(1)})$. Then we have the following diagram.

$$\begin{array}{ccccc}
 \text{sk}(X) & & \supseteq & X^{(1)} & & \supseteq & X^{(2)} \\
 & \searrow \varphi & & & \searrow \varphi_1 & & \\
 & & \psi & & & \psi_1 & \\
 \text{sk}(Y) & & \supseteq & Y^{(1)} & & \supseteq & Y^{(2)}
 \end{array}$$

Now $\psi \circ \varphi$ is a weight-preserving map from $\text{sk}(X)$ into itself. Furthermore, it is continuous by Proposition 3.6.5. Since $\text{sk}(X)$ is Hausdorff (Proposition 2.12.3), an application of Theorem 3.6.8 now shows that $\psi \circ \varphi : \text{sk}(X) \rightarrow \text{sk}(X)$ is surjective. It follows from Definition 2.4.3 that $\psi \circ \varphi$ is an automorphism of $\text{sk}(X)$, hence a bijection. It follows that φ is injective. The dual argument for $\varphi \circ \psi$ shows that ψ is also injective.

Since $\psi \circ \varphi(\text{sk}(X)) = X^{(2)} = \text{sk}(X)$ and $X^{(2)} \subseteq X^{(1)} \subseteq \text{sk}(X)$, we must have $X^{(1)} = \text{sk}(X)$. Similarly, $Y^{(1)} = \text{sk}(Y)$. Thus $\varphi : \text{sk}(X) \rightarrow \text{sk}(Y)$ and $\psi : \text{sk}(Y) \rightarrow \text{sk}(X)$ are a weight-preserving bijections. In particular, we have $\text{sk}(X) \cong^s \text{sk}(Y)$. This concludes the proof. \square

3.6.4 Proofs related to Theorem 3.3.1

Lemma 3.6.9 Let $X_1, \dots, X_n \in \mathcal{FN}$, and for each $i = 1, \dots, n-1$, let $R_i \in \mathcal{R}(X_i, X_{i+1})$. Define

$$\begin{aligned}
 R &:= R_1 \circ R_2 \circ \dots \circ R_n \\
 &= \left\{ (x_1, x_n) \in X_1 \times X_n \mid \exists (x_i)_{i=2}^{n-1}, (x_i, x_{i+1}) \in R_i \text{ for all } i \right\}.
 \end{aligned}$$

Then $\text{dis}(R) \leq \sum_{i=1}^n \text{dis}(R_i)$.

Proof We proceed by induction, beginning with the base case $n = 2$. For convenience, write $X := X_1$, $Y := X_2$, and $Z := X_3$. Let $(x, z), (x', z') \in R_1 \circ R_2$. Let $y \in Y$ be such that $(x, y) \in R_1$ and $(y, z) \in R_2$. Let $y' \in Y$ be such that $(x', y') \in R_1, (y', z') \in$

R_2 . Then we have:

$$\begin{aligned} |\omega_X(x, x') - \omega_Z(z, z')| &= |\omega_X(x, x') - \omega_Y(y, y') + \omega_Y(y, y') - \omega_Z(z, z')| \\ &\leq |\omega_X(x, x') - \omega_Y(y, y')| + |\omega_Y(y, y') - \omega_Z(z, z')| \\ &\leq \text{dis}(R) + \text{dis}(S). \end{aligned}$$

This holds for any $(x, z), (x', z') \in R \circ S$, and proves the claim.

Suppose that the result holds for $n = N \in \mathbb{N}$. Write $R' = R_1 \circ \dots \circ R_N$ and $R = R' \circ R_{N+1}$. Since R' is itself a correspondence, applying the base case yields:

$$\begin{aligned} \text{dis}(R) &\leq \text{dis}(R') + \text{dis}(R_{N+1}) \\ &\leq \sum_{i=1}^N \text{dis}(R_i) + \text{dis}(R_{N+1}) \quad \text{by induction} \\ &= \sum_{i=1}^{N+1} \text{dis}(R_i). \end{aligned}$$

This proves the lemma. \square

Theorem 3.3.1 $(\mathcal{CN}/\cong^w, d_N)$ is a complete metric space containing $(\mathcal{FN}/\cong^w, d_N)$.

Proof of Theorem 3.3.1 Let $([X_i])_{i \in \mathbb{N}}$ be a Cauchy sequence in \mathcal{FN}/\cong^w . First we wish to show this sequence converges in \mathcal{CN}/\cong^w . Note that $(X_i)_{i \in \mathbb{N}}$ is a Cauchy sequence in \mathcal{FN} , since the distance between two equivalence classes is given by the distance between any representatives. To show $(X_i)_i$ converges, it suffices to show that a subsequence of $(X_i)_i$ converges, so without loss of generality, suppose $d_N(X_i, X_{i+1}) < 2^{-i}$ for each i . Then for each i , there exists $R_i \in \mathcal{R}(X_i, X_{i+1})$ such that $\text{dis}(R_i) \leq 2^{-i+1}$. Fix such a sequence $(R_i)_{i \in \mathbb{N}}$. For $j > i$, define

$$R_{ij} := R_i \circ R_{i+1} \circ R_{i+2} \circ \dots \circ R_{j-1}.$$

By Lemma 3.6.9, $\text{dis}(R_{ij}) \leq \text{dis}(R_i) + \text{dis}(R_{i+1}) + \dots + \text{dis}(R_{j-1}) \leq 2^{-i+2}$. Next define:

$$\overline{X} := \{(x_j) : (x_j, x_{j+1}) \in R_j \text{ for all } j \in \mathbb{N}\} \subseteq \prod_{i \in \mathbb{N}} X_i.$$

To see $\overline{X} \neq \emptyset$, let $x_1 \in X_1$, and use the (nonempty) correspondences to pick a sequence (x_1, x_2, x_3, \dots) . By construction, $(x_i) \in \overline{X}$.

Define $\omega_{\overline{X}}((x_j), (x'_j)) = \limsup_{j \rightarrow \infty} \omega_{X_j}(x_j, x'_j)$. We claim that $\omega_{\overline{X}}$ is bounded, and thus is a real-valued weight function. To see this, let $(x_j), (x'_j) \in \overline{X}$. Let $j \in \mathbb{N}$. Then we have:

$$|\omega_{X_j}(x_j, x'_j)| = |\omega_{X_j}(x_j, x'_j) - \omega_{X_{j-1}}(x_{j-1}, x'_{j-1}) + \omega_{X_{j-1}}(x_{j-1}, x'_{j-1}) - \dots|$$

$$\begin{aligned}
& -\omega_{X_1}(x_1, x'_1) + \omega_{X_1}(x_1, x'_1)| \\
& \leq |\omega_{X_1}(x_1, x'_1)| + \text{dis}(R_1) + \text{dis}(R_2) + \cdots + \text{dis}(R_{j-1}) \\
& \leq |\omega_{X_1}(x_1, x'_1)| + 2
\end{aligned}$$

But j was arbitrary. Thus we obtain:

$$|\omega_{\overline{X}}(x_j), (x'_j))| = \limsup_{j \rightarrow \infty} |\omega_{X_j}(x_j, x'_j)| \leq |\omega_{X_1}(x_1, x'_1)| + 2 < \infty.$$

Claim 10 $(\overline{X}, \omega_{\overline{X}}) \in \mathcal{CN}$. More specifically, \overline{X} is a first countable compact topological space, and $\omega_{\overline{X}}$ is continuous with respect to the product topology on $\overline{X} \times \overline{X}$.

Proof of Claim 10 We equip $\prod_{i \in \mathbb{N}} X_i$ with the product topology. First note that the countable product $\prod_{i \in \mathbb{N}} X_i$ of first countable spaces is first countable. Any subspace of a first countable space is first countable, so $\overline{X} \subseteq \prod_{i \in \mathbb{N}} X_i$ is first countable. By Tychonoff's theorem, $\prod_{i \in \mathbb{N}} X_i$ is compact. So to show that \overline{X} is compact, we only need to show that it is closed.

If $\overline{X} = \prod_{i \in \mathbb{N}} X_i$, we would automatically know that \overline{X} is compact. Suppose not, and let $(x_i)_{i \in \mathbb{N}} \in (\prod_{i \in \mathbb{N}} X_i) \setminus \overline{X}$. Then there exists $N \in \mathbb{N}$ such that $(x_N, x_{N+1}) \notin R_N$. Define:

$$U := X_1 \times X_2 \times \cdots \times \{x_N\} \times \{x_{N+1}\} \times X_{N+2} \times \cdots.$$

Since X_i has the discrete topology for each $i \in \mathbb{N}$, it follows that $\{x_N\}$ and $\{x_{N+1}\}$ are open. Hence U is an open neighborhood of $(x_i)_{i \in \mathbb{N}}$ and is disjoint from $\prod_{i \in \mathbb{N}} X_i$. It follows that $(\prod_{i \in \mathbb{N}} X_i) \setminus \overline{X}$ is open, hence \overline{X} is closed and thus compact.

It remains to show that $\omega_{\overline{X}}$ is continuous. We will show that preimages of open sets in \mathbb{R} under $\omega_{\overline{X}}$ are open. Let $(a, b) \subseteq \mathbb{R}$, and suppose $\omega_{\overline{X}}^{-1}[(a, b)]$ is nonempty (otherwise, there is nothing to show). Let $(x_i)_{i \in \mathbb{N}}, (x'_i)_{i \in \mathbb{N}} \in \overline{X} \times \overline{X}$ be such that

$$\alpha := \omega_{\overline{X}}((x_i)_i, (x'_i)_i) \in (a, b).$$

Write $r' := \min(|\alpha - a|, |b - \alpha|)$, and define $r := \frac{1}{2}r'$.

Let $N \in \mathbb{N}$ be such that $2^{-N+3} < r$. Consider the following open sets:

$$\begin{aligned}
U &:= \{x_1\} \times \{x_2\} \times \cdots \times \{x_N\} \times X_{N+1} \times X_{N+2} \times \cdots \subseteq \prod_{i \in \mathbb{N}} X_i, \\
V &:= \{x'_1\} \times \{x'_2\} \times \cdots \times \{x'_N\} \times X_{N+1} \times X_{N+2} \times \cdots \subseteq \prod_{i \in \mathbb{N}} X_i.
\end{aligned}$$

Next write $A := \overline{X} \cap U$ and $B := \overline{X} \cap V$. Then A and B are open with respect to the subspace topology on \overline{X} . Thus $A \times B$ is open in $\overline{X} \times \overline{X}$. Note that $(x_i)_{i \in \mathbb{N}} \in A$ and $(x'_i)_{i \in \mathbb{N}} \in B$. We wish to show that $A \times B \subseteq \omega_{\overline{X}}^{-1}[(a, b)]$, so it suffices to show that $\omega_{\overline{X}}(A, B) \subseteq (a, b)$.

Let $(z_i)_{i \in \mathbb{N}} \in A$ and $(z'_i)_{i \in \mathbb{N}} \in B$. Notice that $z_i = x_i$ and $z'_i = x'_i$ for each $i \leq N$. So for $n \leq N$, we have $|\omega_{X_n}(z_n, z'_n) - \omega_{X_n}(x_n, x'_n)| = 0$.

Next let $n \in \mathbb{N}$, and note that:

$$\begin{aligned} & |\omega_{X_{N+n}}(z_{N+n}, z'_{N+n}) - \omega_{X_{N+n}}(x_{N+n}, x'_{N+n})| \\ &= |\omega_{X_{N+n}}(z_{N+n}, z'_{N+n}) - \omega_{X_N}(z_N, z'_N) + \omega_{X_N}(z_N, z'_N) - \omega_{X_{N+n}}(x_{N+n}, x'_{N+n})| \\ &= |\omega_{X_{N+n}}(z_{N+n}, z'_{N+n}) - \omega_{X_N}(z_N, z'_N) + \omega_{X_N}(x_N, x'_N) - \omega_{X_{N+n}}(x_{N+n}, x'_{N+n})| \\ &\leq \text{dis}(R_{N,N+n}) + \text{dis}(R_{N,N+n}) \leq 2^{-N+2} + 2^{-N+2} = 2^{-N+3} < r. \end{aligned}$$

Here the second to last inequality follows from Lemma 3.6.9. The preceding calculation holds for arbitrary $n \in \mathbb{N}$. It follows that:

$$\limsup_{i \rightarrow \infty} \omega_{X_i}(x_i, x'_i) - \limsup_{i \rightarrow \infty} \omega_{X_i}(z_i, z'_i) \leq \limsup_{i \rightarrow \infty} (\omega_{X_i}(x_i, x'_i) - \omega_{X_i}(z_i, z'_i)) < r,$$

and similarly $\limsup_{i \rightarrow \infty} \omega_{X_i}(z_i, z'_i) - \limsup_{i \rightarrow \infty} \omega_{X_i}(x_i, x'_i) < r$. Thus we have $\omega_{\overline{X}}((z_i)_i, (z'_i)_i) \in (a, b)$. This proves continuity of $\omega_{\overline{X}}$. \square

Next we claim that $X_i \xrightarrow{d_{\mathcal{N}}} \overline{X}$ as $i \rightarrow \infty$. Fix $i \in \mathbb{N}$. We wish to construct a correspondence $S \in \mathcal{R}(X_i, \overline{X})$. Let $y \in X_i$. We write $x_i = y$ and pick $x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots$ such that $(x_j, x_{j+1}) \in R_j$ for each $j \in \mathbb{N}$. We denote this sequence by $(x_j)^{x_i=y}$, and note that by construction, it lies in \overline{X} . Conversely, for any $(x_j) \in \overline{X}$, we simply pick its i th coordinate x_i as a corresponding element in X_i . We define:

$$\begin{aligned} S &:= A \cup B, \quad \text{where} \\ A &:= \{(y, (x_j)^{x_i=y}) : y \in X_i\} \\ B &:= \{(x_i, (x_k)) : (x_k) \in \overline{X}\} \end{aligned}$$

Then $S \in \mathcal{R}(X_i, \overline{X})$. We claim that $\text{dis}(S) \leq 2^{-i+2}$. Let $z = (y, (x_k))$, $z' = (y', (x'_k)) \in B$. Let $n \in \mathbb{N}$, $n \geq i$. Then we have:

$$\begin{aligned} |\omega_{X_i}(y, y') - \omega_{X_n}(x_n, x'_n)| &= |\omega_{X_i}(y, y') - \omega_{X_{i+1}}(x_{i+1}, x'_{i+1}) + \omega_{X_{i+1}}(x_{i+1}, x'_{i+1}) \\ &\quad - \dots + \omega_{X_{n-1}}(x_{n-1}, x'_{n-1}) + \omega_{X_n}(x_n, x'_n)| \\ &\leq \text{dis}(R_i) + \text{dis}(R_{i+1}) + \dots + \text{dis}(R_{n-1}) \\ &\leq 2^{-i+1} + 2^{-i} + \dots + 2^{-n+2} \\ &\leq 2^{-i+2}. \end{aligned}$$

This holds for arbitrary $n \geq i$. It follows that we have:

$$|\omega_{X_i}(y, y') - \omega_{\overline{X}}((x_k), (x'_k))| \leq 2^{-i+2}.$$

Similar inequalities hold for $z, z' \in A$, and for $z \in A, z' \in B$. Thus $\text{dis}(S) \leq 2^{-i+2}$. It follows that $d_{\mathcal{N}}(X_i, \bar{X}) \leq 2^{-i+1}$. Thus the sequence $([X_i])_i$ converges to $[\bar{X}] \in \mathcal{CN}/\cong^w$.

Finally, we need to check that $(\mathcal{CN}/\cong^w, d_{\mathcal{N}})$ is complete. Let $([Y_n])_n$ be a Cauchy sequence in \mathcal{CN}/\cong^w . By invoking Theorem 3.1.3, for each n let $[X_n] \in \mathcal{FN}/\cong^w$ be such that $d_{\mathcal{N}}([X_n], [Y_n]) < \frac{1}{n}$. Let $\varepsilon > 0$. Then for sufficiently large m and n , we have:

$$d_{\mathcal{N}}([X_n], [X_m]) \leq d_{\mathcal{N}}([X_n], [Y_n]) + d_{\mathcal{N}}([Y_n], [Y_m]) + d_{\mathcal{N}}([Y_m], [X_m]) < \varepsilon.$$

Thus $([X_n])_n$ is a Cauchy sequence in \mathcal{FN}/\cong^w . By applying what we have shown above, this sequence converges to some $[\bar{X}] \in \mathcal{CN}/\cong^w$. By applying the triangle inequality, we see that the sequence $([Y_n])_n$ also converges to $[\bar{X}]$. This shows completeness, and concludes the proof. \square

Remark 3.6.10 In the proof of Theorem 3.3.1, note that the construction of the limit is dependent upon the initial choice of optimal correspondences. However, all such limits obtained from different choices of optimal correspondences belong to the same weak isomorphism class.

3.6.5 Proofs related to Theorem 3.5.4

Theorem 3.5.3 *The metric space $(\mathcal{FN}/\cong^w, d_{\mathcal{N}})$ is a geodesic space. More specifically, let $[X], [Y] \in (\mathcal{FN}/\cong^w, d_{\mathcal{N}})$. Then, for any $R \in \mathcal{R}^{\text{opt}}(X, Y)$, we can construct a geodesic $\gamma_R : [0, 1] \rightarrow \mathcal{FN}/\cong^w$ between $[X]$ and $[Y]$ as follows:*

$$\begin{aligned} \gamma_R(0) &:= [(X, \omega_X)], \quad \gamma_R(1) := [(Y, \omega_Y)], \quad \text{and} \\ \gamma_R(t) &:= [(R, \omega_{\gamma_R(t)})] \quad \text{for } t \in (0, 1), \end{aligned}$$

where for each $(x, y), (x', y') \in R$ and $t \in (0, 1)$,

$$\omega_{\gamma_R(t)}((x, y), (x', y')) := (1 - t) \cdot \omega_X(x, x') + t \cdot \omega_Y(y, y').$$

Proof of Theorem 3.5.3 Let $[X], [Y] \in \mathcal{FN}/\cong^w$. We will show the existence of a curve $\gamma : [0, 1] \rightarrow \mathcal{FN}$ such that $\gamma(0) = (X, \omega_X)$, $\gamma(1) = (Y, \omega_Y)$, and for all $s, t \in [0, 1]$,

$$d_{\mathcal{N}}(\gamma(s), \gamma(t)) = |t - s| \cdot d_{\mathcal{N}}(X, Y).$$

Note that this yields $d_{\mathcal{N}}([\gamma(s)], [\gamma(t)]) = |t - s| \cdot d_{\mathcal{N}}([X], [Y])$ for all $s, t \in [0, 1]$, which is what we need to show.

Let $R \in \mathcal{R}^{\text{opt}}(X, Y)$, i.e. let R be a correspondence such that $\text{dis}(R) = 2d_{\mathcal{N}}(X, Y)$. For each $t \in (0, 1)$ define $\gamma(t) := (R, \omega_{\gamma(t)})$, where

$$\omega_{\gamma(t)}((x, y), (x', y')) := (1 - t) \cdot \omega_X(x, x') + t \cdot \omega_Y(y, y') \quad \text{for all } (x, y), (x', y') \in R.$$

Also define $\gamma(0) = (X, \omega_X)$ and $\gamma(1) = (Y, \omega_Y)$.

Claim 11 For any $s, t \in [0, 1]$,

$$d_{\mathcal{N}}(\gamma(s), \gamma(t)) \leq |t - s| \cdot d_{\mathcal{N}}(X, Y).$$

Suppose for now that Claim 11 holds. We further claim that this implies, for all $s, t \in [0, 1]$,

$$d_{\mathcal{N}}(\gamma(s), \gamma(t)) = |t - s| \cdot d_{\mathcal{N}}(X, Y).$$

To see this, assume towards a contradiction that there exist $s_0 < t_0$ such that:

$$d_{\mathcal{N}}(\gamma(s_0), \gamma(t_0)) < |t_0 - s_0| \cdot d_{\mathcal{N}}(X, Y).$$

$$\begin{aligned} \text{Then } d_{\mathcal{N}}(X, Y) &\leq d_{\mathcal{N}}(X, \gamma(s_0)) + d_{\mathcal{N}}(\gamma(s_0), \gamma(t_0)) + d_{\mathcal{N}}(\gamma(t_0), Y) \\ &< |s_0 - 0| \cdot d_{\mathcal{N}}(X, Y) + |t_0 - s_0| \cdot d_{\mathcal{N}}(X, Y) + |1 - t_0| \cdot d_{\mathcal{N}}(X, Y) \\ &= d_{\mathcal{N}}(X, Y), \text{ a contradiction.} \end{aligned}$$

Thus it suffices to show Claim 11. There are three cases: (i) $s, t \in (0, 1)$, (ii) $s = 0, t \in (0, 1)$, and (iii) $s \in (0, 1), t = 1$. The latter two cases are similar, so we just prove (i) and (ii). For (i), fix $s, t \in (0, 1)$. Notice that $\Delta := \text{diag}(R \times R) := \{(r, r) : r \in R\}$ is a correspondence in $\mathcal{R}(R, R)$. Then we obtain:

$$\begin{aligned} \text{dis}(\Delta) &= \max_{(a,a),(b,b) \in \Delta} |\omega_{\gamma(t)}(a, b) - \omega_{\gamma(s)}(a, b)| \\ &= \max_{(x,y),(x',y') \in R} |\omega_{\gamma(t)}((x, y), (x', y')) - \omega_{\gamma(s)}((x, y), (x', y'))| \\ &= \max_{(x,y),(x',y') \in R} |(1-t)\omega_X(x, x') + t \cdot \omega_Y(y, y') - (1-s)\omega_X(x, x') \\ &\quad - s \cdot \omega_Y(y, y')| \\ &= \max_{(x,y),(x',y') \in R} |(s-t)\omega_X(x, x') - (s-t)\omega_Y(y, y')| \\ &= |t-s| \cdot \max_{(x,y),(x',y') \in R} |\omega_X(x, x') - \omega_Y(y, y')| \\ &\leq 2|t-s| \cdot d_{\mathcal{N}}(X, Y). \end{aligned}$$

Finally $d_{\mathcal{N}}(\gamma(t), \gamma(s)) \leq \frac{1}{2} \text{dis}(\Delta) \leq |t-s| \cdot d_{\mathcal{N}}(X, Y)$.

For (ii), fix $s = 0, t \in (0, 1)$. Define $R_X = \{(x, (x, y)) : (x, y) \in R\}$. Then R_X is a correspondence in $\mathcal{R}(X, R)$.

$$\begin{aligned} \text{dis}(R_X) &= \max_{(x,(x,y)),(x',(x',y')) \in R_X} |\omega_X(x, x') - (1-t) \cdot \omega_X(x, x') - t \cdot \omega_Y(y, y')| \\ &= \max_{(x,(x,y)),(x',(x',y')) \in R_X} t \cdot |\omega_X(x, x') - \omega_Y(y, y')| \\ &= t \text{dis}(R) = 2t \cdot d_{\mathcal{N}}(X, Y). \end{aligned}$$

Thus $d_{\mathcal{N}}(X, \gamma(t)) \leq t \cdot d_{\mathcal{N}}(X, Y)$. The proof for case (iii), i.e. that $d_{\mathcal{N}}(\gamma(s), Y) \leq |1 - s| \cdot d_{\mathcal{N}}(X, Y)$, is similar. This proves Claim 11, and the result follows. \square

Theorem 3.5.4 *The complete metric space $(\mathcal{CN}/\cong^w, d_{\mathcal{N}})$ is geodesic.*

Proof of Theorem 3.5.4 Let $[X], [Y] \in \mathcal{CN}/\cong^w$. It suffices to find a geodesic between X and Y , because the distance between any two equivalence classes is given by the distance between any two representatives, and hence we will obtain a geodesic between $[X]$ and $[Y]$.

Let $(X_n)_n, (Y_n)_n$ be sequences in \mathcal{FN} such that $d_{\mathcal{N}}(X_n, X) < \frac{1}{n}$ and $d_{\mathcal{N}}(Y_n, Y) < \frac{1}{n}$ for each n . For each n , let R_n be an optimal correspondence between X_n and Y_n , endowed with the weight function

$$\omega_n((x, y), (a, b)) = \frac{1}{2}\omega_{X_n}(x, a) + \frac{1}{2}\omega_{Y_n}(y, b).$$

By the proof of Theorem 3.5.3, the network (R_n, ω_n) is a midpoint of X_n and Y_n .

Claim 12 The collection $\{R_n : n \in \mathbb{N}\}$ is precompact.

Assume for now that Claim 12 is true. Then we can pick a sequence (R_n) that converges to some $R \in \mathcal{CN}$. Then we obtain:

$$\begin{aligned} d_{\mathcal{N}}(X, R) &\leq d_{\mathcal{N}}(X, X_n) + d_{\mathcal{N}}(X_n, R_n) + d_{\mathcal{N}}(R_n, R) \\ &= d_{\mathcal{N}}(X, X_n) + \frac{1}{2}d_{\mathcal{N}}(X_n, Y_n) + d_{\mathcal{N}}(R_n, R) \rightarrow \frac{1}{2}d_{\mathcal{N}}(X, Y). \end{aligned}$$

Similarly $d_{\mathcal{N}}(R, Y) \leq \frac{1}{2}d_{\mathcal{N}}(X, Y)$. Furthermore, equality holds in both inequalities, because we would get a contradiction otherwise. Thus R is a midpoint of X and Y , and moreover, $[R]$ is a midpoint of $[X]$ and $[Y]$. The result now follows by an application of Theorem 3.5.2.

It remains to prove Claim 12. By Theorem 3.4.4, it suffices to show that $\{R_n\}$ is uniformly approximable.

Since $d_{\mathcal{N}}(X_n, X) \rightarrow 0$ and $d_{\mathcal{N}}(Y_n, Y) \rightarrow 0$, we can choose $D > 0$ large enough so that $\text{diam}(X_n) \leq \frac{D}{2}$ and $\text{diam}(Y_n) \leq \frac{D}{2}$ for all n . Then $\text{diam}(R_n) \leq D$ for all n .

Let $\varepsilon > 0$. Fix N large enough so that $\frac{1}{N} < \frac{\varepsilon}{2}$, and write $N(\varepsilon) = \max_{n \leq N} |R_n|$. We wish to show that every R_n is ε -approximable by a finite network with cardinality up to $N(\varepsilon)$. For any $n \leq N$, we know R_n approximates itself, and $|R_n| \leq N(\varepsilon)$. Next let $n > N$. It will suffice to show that R_n is ε -approximable by R_N .

Let S, T be optimal correspondences between X_n, X_N and Y_n, Y_N respectively. Note that $d_{\mathcal{N}}(X_N, X_n) \leq d_{\mathcal{N}}(X_N, X) + d_{\mathcal{N}}(X, X_n) \leq \frac{1}{N} + \frac{1}{N} = \frac{2}{N}$, and similarly $d_{\mathcal{N}}(Y_N, Y_n) \leq \frac{2}{N}$. Thus $\text{dis}(S) \leq \frac{4}{N}$ and $\text{dis}(T) \leq \frac{4}{N}$. Next write

$$Q = \{(x, y, x', y') \in R_N \times R_n : (x, x') \in S, (y, y') \in T\}.$$

Observe that since S and T are correspondences, Q is a correspondence between R_N and R_n . Next we calculate $\text{dis}(Q)$:

$$\text{dis}(Q) = \max_{\substack{(x, y, x', y'), \\ (a, b, a', b') \in Q}} |\omega_N((x, y), (a, b)) - \omega_n((x', y'), (a', b'))|$$

$$\begin{aligned}
 &= \max_{\substack{(x,y,x',y'), \\ (a,b,a',b') \in Q}} \left| \frac{1}{2} \omega_{X_N}(x, a) + \frac{1}{2} \omega_{Y_N}(y, b) - \frac{1}{2} \omega_{X_n}(x', a') - \frac{1}{2} \omega_{Y_n}(y', b') \right| \\
 &\leq \frac{1}{2} \max_{(x,x'), (a,a') \in S} |\omega_{X_N}(x, a) - \omega_{X_n}(x', a')| \\
 &\quad + \frac{1}{2} \max_{(y,y'), (b,b') \in S} |\omega_{Y_N}(y, b) - \omega_{Y_n}(y', b')| \\
 &= \frac{1}{2} \text{dis}(S) + \frac{1}{2} \text{dis}(T) \leq \frac{4}{N}.
 \end{aligned}$$

Thus $d_{\mathcal{N}}(R_N, R_n) \leq \frac{2}{N} < \varepsilon$. This shows that any R_n can be ε -approximated by a network having up to $N(\varepsilon)$ points. Thus $\{R_n\}$ is uniformly approximable, hence precompact. Thus Claim 12 and the result follow. \square

4 Invariants, stability, and convergence

At this point, we have computed $d_{\mathcal{N}}$ between several examples of networks, as in Example 2.7.2 and Remark 2.7.3. We also asserted in Remark 2.7.6 that $d_{\mathcal{N}}$ is in general difficult to compute. A standard approach that is used in practice to circumvent the problem of computing $d_{\mathcal{N}}$ is to compute lower bounds instead. We frame this approach in terms of defining *quantitatively stable invariants* of networks, and then comparing the invariants instead of comparing the networks directly. In this section, we set up the framework of stable network invariants, explain how this incorporates existing stability results on persistent homology and hierarchical clustering, and go on to prove convergence of these methods in the infinite setting. In particular, the setting of compact networks gives a suitable family of infinite networks with well-defined persistence diagrams.

Throughout this section, we will reserve the notation X for random variables, and switch to using letters Y, Z to denote networks.

4.1 Preliminary setup

Throughout this section, all the underlying topological spaces will be assumed to be compact, unless specified otherwise. We use the following convention for an open cover of a topological space: An *open cover* of a topological space Z is a collection of open sets $\{U_i \subseteq Z : i \in I\}$ indexed by some set I such that each U_i is *nonempty* and $\bigcup_{i \in I} U_i = Z$. Given a topological space Z , we will write $\text{Borel}(Z)$ to denote the Borel σ -field on Z . We often write $(\Omega, \mathcal{F}, \mathbb{P})$ to denote a probability space. The support supp of a measure μ_Z on a topological space Z is defined as:

$$\text{supp}(\mu_Z) := \{z \in Z : \text{for each open neighborhood } N_z \ni z, \text{ we have } \mu_Z(N_z) > 0\}.$$

The complement of $\text{supp}(\mu_Z)$ is the union of open sets of measure zero. It follows that $\text{supp}(\mu_Z)$ is closed, hence compact.

Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a measurable space (Z, \mathcal{G}) , a *random variable* (defined on Ω with values in Z) is a measurable function $\mathbf{z} : \Omega \rightarrow Z$. The

pushforward or *image measure* of \mathbf{z} is defined to be the measure $\mathbf{z}_\# \mathbb{P}$ on \mathcal{G} given by writing $\mathbf{z}_\# \mathbb{P}(A) := \mathbb{P}(\mathbf{z}^{-1}[A])$ for all $A \in \mathcal{G}$. The pushforward is often called the *distribution* of \mathbf{z} .

We recall an important corollary of the existence of infinite products of probability measures. For any probability space (Z, \mathcal{E}, μ_Z) , there exists a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which there are independent random variables $\mathbf{z}_1, \mathbf{z}_2, \dots$ taking values in Z with distribution μ_Z (Dudley 2002, §8.2). This is done by letting $\Omega := \prod_{n \in \mathbb{N}} Z$ and taking each \mathbf{z}_i to be the canonical projection map $(z_i)_{i \in \mathbb{N}} \mapsto z_i$.

Definition 4.1.1 (*Measure network*, Chowdhury and Mémoli 2019) A *measure network* is a triple (Z, ω_Z, μ_Z) where Z is a compact Polish space, μ_Z is a Borel probability measure on Z , and $\omega_Z : Z \times Z \rightarrow \mathbb{R}$ is a bounded measurable function.

We now specify what it means to sample from a measure network.

Definition 4.1.2 (*Sampling from a measure network*) Let (Z, ω_Z, μ_Z) be a measure network, and let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A measurable function $\mathbf{z} : \Omega \rightarrow Z$ is said to be a random variable taking values in Z . For each $i \in \mathbb{N}$, let $\mathbf{z}_i : \Omega \rightarrow Z$ be an independent random variable with distribution μ_Z . For each $n \in \mathbb{N}$, let $\mathcal{Z}_n = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n\}$. To say that we sample n points randomly from (Z, ω_Z, μ_Z) means that we take a selection $\mathcal{Z}_n(\omega) = \{z_1 := \mathbf{z}_1(\omega), z_2 := \mathbf{z}_2(\omega), \dots, z_n := \mathbf{z}_n(\omega)\}$ and equip the collection $Z_n := \{z_1, \dots, z_n\}$ with the restriction of ω_Z .

A related concept is that of *almost sure convergence w.r.t. $d_{\mathcal{N}}$* . To say that \mathcal{Z}_n converges almost surely to Z with respect to $d_{\mathcal{N}}$ means that for each $\varepsilon > 0$,

$$\mathbb{P} \left(\limsup_{n \rightarrow \infty} \{\omega \in \Omega : d_{\mathcal{N}}(\mathcal{Z}_n(\omega), Z) \geq \varepsilon\} \right) = 0.$$

This notion of sampling is the form used when sampling from metric spaces, and it assumes that selecting the points $Z_n = \{z_1, \dots, z_n\}$ is enough to obtain the restriction of ω_Z to $Z_n \times Z_n$. This is often the case, such as when ω_Z is given by a Gaussian kernel. Note that this is quite different from sampling from a combinatorial graph, where one has to use techniques such as Markov chain Monte Carlo so as to not sample very sparse graphs (Lyu et al. 2019).

The presence of μ_Z allows us to ask questions such as the following: Given a measure network (Z, ω_Z, μ_Z) and some $\varepsilon > 0$, what should it mean to take an *optimal ε -system* (cf. Definition 3.1.2) on Z ? The next definition sheds some light on this question.

Definition 4.1.3 (*Minimal mass function \mathfrak{m}*) Let (Z, ω_Z, μ_Z) be a measure network. Let \mathcal{U} be any ε -system on Z for $\varepsilon > 0$. We define the *minimal mass function* $\mathfrak{m}(\mathcal{U}) := \min \{\mu_Z(U) : U \in \mathcal{U}, \mu_Z(U) > 0\}$. Note that \mathfrak{m} returns the minimal non-zero mass of an element in \mathcal{U} .

Next let $\varepsilon > 0$. Define an *optimal minimal mass function* $\mathfrak{M}_\varepsilon : \mathcal{CN} \rightarrow (0, 1]$ as follows:

$$\mathfrak{M}_\varepsilon(Z) := \sup \{\mathfrak{m}(\mathcal{U}) : \mathcal{U} \text{ a refined } \varepsilon\text{-system on } Z\}.$$

Recall that a *refined* ε -system was defined in Definition 3.1.2. Since \mathcal{U} covers Z , we know that the total mass of \mathcal{U} is 1. Thus the set of elements in \mathcal{U} with positive mass is nonempty, and so $m(\mathcal{U})$ is strictly positive. It follows that $\mathfrak{M}_\varepsilon(Z)$ is strictly positive. More is true when μ_Z is fully supported on Z : given any ε -system \mathcal{U} on Z and any $U \in \mathcal{U}$, we automatically have $\mu_Z(U) > 0$ (recall our convention above that empty elements are excluded from an open cover).

In the preceding definition, for a given $\varepsilon > 0$, the function $\mathfrak{M}_\varepsilon(Z)$ considers the collection of all *refined* ε -systems on Z , and then maximizes the minimal mass of any element in such an ε -system. The functions in Definition 4.1.3 are crucial to the next result, which shows that as we sample points from a distribution on a network, the sampled subnetwork converges almost surely to the support of the distribution.

Theorem 4.1.4 (Probabilistic network approximation) *Let (Z, ω_Z, μ_Z) be a measure network. For each $i \in \mathbb{N}$, let $z_i : \Omega \rightarrow Z$ be an independent random variable defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with distribution μ_Z . For each $n \in \mathbb{N}$, let $\mathcal{Z}_n = \{z_1, z_2, \dots, z_n\}$. Let $\varepsilon > 0$. Then we have:*

$$\mathbb{P}(\{\omega \in \Omega : d_{\mathcal{N}}(\text{supp}(\mu_Z), Z_n) \geq \varepsilon\}) \leq \frac{(1 - \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)))^n}{\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z))},$$

where $Z_n = \mathcal{Z}_n(\omega)$ is an n -point sample as in Definition 4.1.2. In particular, the subnetwork \mathcal{Z}_n converges almost surely to Z w.r.t. $d_{\mathcal{N}}$.

Remark 4.1.5 (Relation to the metric setting) The invariants in Definition 4.1.3 are abstractions of invariants that are standard in the metric setting. For a concrete example, note that the mass of an ε -ball in d -dimensional Euclidean space scales as ε^d . Thus in the setting of Euclidean space \mathbb{R}^d , the quantity on the right of Theorem 4.1.4 would scale as $\varepsilon^{-d}(1 - \varepsilon^d)^n$. For stronger results in the metric setting, see Carlsson and Mémoli (2010, Theorem 30).

Before proving the theorem, we prove the following useful lemma:

Lemma 4.1.6 *Assume the setup of (Z, ω_Z, μ_Z) , $(\Omega, \mathcal{F}, \mathbb{P})$, and \mathcal{Z}_n for each $n \in \mathbb{N}$ as in Theorem 4.1.4. Fix $\varepsilon > 0$, and let $\mathcal{U} = \{U_1, \dots, U_m\}$ be a refined ε -system on $\text{supp}(\mu_Z)$. For each $1 \leq i \leq m$ and each $n \in \mathbb{N}$, define the following event:*

$$A_i := \bigcap_{k=1}^n \{\omega \in \Omega : z_k(\omega) \notin U_i\} \subseteq \Omega.$$

Then we have $\mathbb{P}(\bigcup_{k=1}^m A_k) \leq \frac{1}{m(\mathcal{U})}(1 - m(\mathcal{U}))^n$.

Proof of Lemma 4.1.6 Here we are considering the probability that at least one of the U_i has empty intersection with \mathcal{Z}_n . As the points are sampled independently, we can invoke independence and write $\mathbb{P}(A_i) = (1 - \mu_Z(U_i))^n$. Then we have:

$$\mathbb{P}\left(\bigcup_{k=1}^m A_k\right) \leq \sum_{k=1}^m \mathbb{P}(A_k) = \sum_{k=1}^m (1 - \mu_Z(U_k))^n$$

$$\leq m \cdot \max_{1 \leq k \leq m} (1 - \mu_Z(U_k))^n \leq \frac{(1 - m(\mathcal{U}))^n}{m(\mathcal{U})}.$$

Here the first inequality follows by subadditivity of measure, and the last inequality follows because the total mass $\mu_Z(\text{supp}(\mu_Z)) = 1$ is an upper bound for $m \cdot m(\mathcal{U})$. Note also that each $U \in \mathcal{U}$ has nonzero mass, by the observation in Definition 4.1.3. \square

Proof of Theorem 4.1.4 By endowing $\text{supp}(\mu_Z)$ with the restriction of ω_Z , it may itself be viewed as a network with full support. So for notational convenience, we assume $Z = \text{supp}(\mu_Z)$.

First observe that $\mathfrak{M}_{\varepsilon/2}(Z) \in (0, 1]$. Let $r \in (0, \mathfrak{M}_{\varepsilon/2}(Z))$, and let \mathcal{U}_r be an $\varepsilon/2$ -system on Z such that $m(\mathcal{U}_r) \in (r, \mathfrak{M}_{\varepsilon/2}(Z)]$. For convenience, write $m := |\mathcal{U}_r|$, and also write $\mathcal{U}_r = \{U_1, \dots, U_m\}$. For each $1 \leq i \leq m$, define A_i as in the statement of Lemma 4.1.6. Then by Lemma 4.1.6, the probability that at least one U_i has empty intersection with Z_n is bounded as $\mathbb{P}(\bigcup_{k=1}^m A_k) \leq \frac{1}{m(\mathcal{U}_r)}(1 - m(\mathcal{U}_r))^n$. On the other hand, if U_i has nonempty intersection with Z_n for each $1 \leq i \leq m$, then by Theorem 3.1.3 we obtain $d_{\mathcal{N}}(Z, Z_n) < \varepsilon$. For each $n \in \mathbb{N}$, define: $B_n := \{\omega \in \Omega : d_{\mathcal{N}}(Z, Z_n(\omega)) \geq \varepsilon\}$. Then we have:

$$\mathbb{P}(B_n) \leq \mathbb{P}\left(\bigcup_{k=1}^m A_k\right) \leq \frac{(1 - m(\mathcal{U}_r))^n}{m(\mathcal{U}_r)}.$$

Since $r \in (0, \mathfrak{M}_{\varepsilon/2}(Z))$ was arbitrary, letting r approach $\mathfrak{M}_{\varepsilon/2}(Z)$ shows that $\mathbb{P}(B_n) \leq \frac{(1 - \mathfrak{M}_{\varepsilon/2}(Z))^n}{\mathfrak{M}_{\varepsilon/2}(Z)}$. We have by Definition 4.1.3 that $\mathfrak{M}_{\varepsilon/2}(Z)$ is strictly positive. Thus the term on the right side of the inequality is an element of a convergent geometric series, so

$$\sum_{n=1}^{\infty} \mathbb{P}(B_n) \leq \frac{1}{\mathfrak{M}_{\varepsilon/2}(Z)} \sum_{n=1}^{\infty} (1 - \mathfrak{M}_{\varepsilon/2}(Z))^n < \infty.$$

By the Borel–Cantelli lemma, we have $\mathbb{P}(\limsup_{n \rightarrow \infty} B_n) = 0$. The result follows. \square

In the upcoming sections, we will see applications of these results to two different fields of exploratory data analysis: *persistent homology* and *hierarchical clustering*. These methods are classically defined for metric spaces, but generalizations have recently been proposed to allow for meaningful constructions on generalized datasets such as the networks in this work.

4.2 Network invariants and basic examples

An \mathbb{R} -invariant of a network is a map $\iota : \mathcal{CN} \rightarrow \mathbb{R}$ such that for any $Y, Z \in \mathcal{CN}$, if $Y \cong^w Z$ then $\iota(Y) = \iota(Z)$. Any \mathbb{R} -invariant is an example of a *pseudometric* (and in particular, a *metric*) *space valued invariant*, which we define next. Recall that a pseudometric space (V, d_V) is a metric space where we allow $d_V(v, v') = 0$ even if $v \neq v'$.

Definition 4.2.1 Let (V, d_V) be any metric or pseudometric space. A V -valued invariant is any map $\iota : \mathcal{CN} \rightarrow V$ such that $\iota(X, \omega_X) = \iota(Y, \omega_Y)$ whenever $X \cong^w Y$.

In what follows, we will construct several maps that give rise to pseudometric space valued invariants (Proposition 4.2.7). We will eventually prove that our proposed invariants are *quantitatively stable*. This notion is made precise in Sect. 4.3.

Example 4.2.2 Define the *diameter* map to be the map

$$\text{diam} : \mathcal{CN} \rightarrow \mathbb{R} \text{ given by } (Z, \omega_Z) \mapsto \max_{z, z' \in Z} |\omega_Z(z, z')|.$$

Then diam is an \mathbb{R} -invariant. Observe that the maximum is achieved for $(Z, \omega_Z) \in \mathcal{CN}$ because Z (hence $Z \times Z$) is compact and $\omega : Z \times Z \rightarrow \mathbb{R}$ is continuous. An application of diam to Example 2.2.10 gives an upper bound on $d_{\mathcal{N}}(Z, Y)$ for $Z, Y \in \mathcal{CN}$ in the following way:

$$\begin{aligned} d_{\mathcal{N}}(Z, Y) &\leq d_{\mathcal{N}}(Z, N_1(0)) + d_{\mathcal{N}}(N_1(0), Y) \\ &= \frac{1}{2}(\text{diam}(Z) + \text{diam}(Y)) \text{ for any } Z, Y \in \mathcal{CN}. \end{aligned}$$

Recall that $\text{pow}(\mathbb{R})$, the nonempty elements of the power set of \mathbb{R} , is a pseudometric space when endowed with the Hausdorff distance (Burago et al. 2001, Proposition 7.3.3).

Example 4.2.3 Define the *spectrum* map

$$\text{spec} : \mathcal{CN} \rightarrow \text{pow}(\mathbb{R}) \text{ by } (Z, \omega_Z) \mapsto \{\omega_Z(z, z') : z, z' \in Z\}.$$

The spectrum also has two local variants. Define the out-local spectrum of Z by $z \mapsto \text{spec}_Z^{\text{out}}(z) := \{\omega_Z(z, z') : z' \in Z\}$. Notice that $\text{spec}(Z) = \bigcup_{z \in Z} \text{spec}_Z^{\text{out}}(z)$ for any network Z , thus justifying the claim that this construction localizes spec . Similarly, we define the in-spectrum of Z as the map $z \mapsto \text{spec}_Z^{\text{in}}(z) := \{\omega_Z(z', z) : z' \in Z\}$. Notice that one still has $\text{spec}(Z) = \bigcup_{z \in Z} \text{spec}_Z^{\text{in}}(z)$ for any network Z . Finally, we observe that the two local versions of spec do not necessarily coincide in an asymmetric network.

The spectrum is closely related to the *multisets* used by Boutin and Kemper (2007) to produce invariants of weighted undirected graphs. For an undirected graph G , they considered the collection of all subgraphs with three nodes, along with the edge weights for each subgraph (compare to our notion of spectrum). Then they proved that the distribution of edge weights of these subgraphs is an invariant when G belongs to a certain class of graphs.

Example 4.2.4 Define the *trace* map $\text{tr} : \mathcal{CN} \rightarrow \text{pow}(\mathbb{R})$ by $(Z, \omega_Z) \mapsto \text{tr}(Z) := \{\omega_Z(z, z) : z \in Z\}$. This also defines an associated map $z \mapsto \text{tr}_Z(z) := \omega_Z(z, z)$. An example is provided in Fig. 13: in this case, we have $(Z, \text{tr}_Z) = (\{p, q\}, (\alpha, \beta))$.

Fig. 13 The trace map erases data between pairs of nodes





Fig. 14 The out map applied to each node yields the greatest weight of an arrow *leaving* the node, and the in map returns the greatest weight entering the node



Example 4.2.5 (*The out and in maps*) Let $(Z, \omega_Z) \in \mathcal{CN}$, and let $z \in Z$. Now define $\text{out} : \mathcal{CN} \rightarrow \text{pow}(\mathbb{R})$ and $\text{in} : \mathcal{CN} \rightarrow \text{pow}(\mathbb{R})$ by

$$\begin{aligned} \text{out}(Z) &= \left\{ \max_{z' \in Z} |\omega_Z(z, z')| : z \in Z \right\} \quad \text{for all } (Z, \omega_Z) \in \mathcal{CN} \\ \text{in}(Z) &= \left\{ \max_{z' \in Z} |\omega_Z(z', z)| : z \in Z \right\} \quad \text{for all } (Z, \omega_Z) \in \mathcal{CN}. \end{aligned}$$

For each $z \in Z$, $\max_{z' \in Z} |\omega_Z(z, z')|$ and $\max_{z' \in Z} |\omega_Z(z', z)|$ are achieved because $\{z\} \times Z$ and $Z \times \{z\}$ are compact. We also define the associated maps out_Z and in_Z by writing, for any $(Z, \omega_Z) \in \mathcal{CN}$ and any $z \in Z$,

$$\text{out}_Z(z) = \max_{z' \in Z} |\omega_Z(z, z')| \quad \text{in}_Z(z) = \max_{z' \in Z} |\omega_Z(z', z)|.$$

To see how these maps operate on a network, let $Z = \{p, q, r\}$ and consider the weight matrix $\Sigma = \begin{pmatrix} 1 & 2 & 3 \\ 0 & 0 & 4 \\ 0 & 0 & 5 \end{pmatrix}$. The network corresponding to this matrix is shown in Fig. 14. We ascertain the following directly from the matrix:

$$\begin{aligned} \text{out}_Z(p) &= 3 & \text{in}_Z(p) &= 1 \\ \text{out}_Z(q) &= 4 & \text{in}_Z(q) &= 2 \\ \text{out}_Z(r) &= 5 & \text{in}_Z(r) &= 5. \end{aligned}$$

So the out map returns the maximum (absolute) value in each row, and the in map pulls out the maximum (absolute) value in each column of the weight matrix. As in the preceding example, we may use the Hausdorff distance to compare the images of networks under the out and in maps.

Example 4.2.6 (min-out and min-in) Define the maps $m^{\text{out}} : \mathcal{CN} \rightarrow \mathbb{R}$ and $m^{\text{in}} : \mathcal{CN} \rightarrow \mathbb{R}$ by

$$\begin{aligned} m^{\text{out}}((Z, \omega_Z)) &= \min_{z \in Z} \text{out}_Z(z) \quad \text{for all } (Z, \omega_Z) \in \mathcal{CN} \\ m^{\text{in}}((Z, \omega_Z)) &= \min_{z \in Z} \text{in}_Z(z) \quad \text{for all } (Z, \omega_Z) \in \mathcal{CN}. \end{aligned}$$

Then both m^{in} and m^{out} are \mathbb{R} -invariants. We take the minimum when defining $m^{\text{out}}, m^{\text{in}}$ because for any network (Z, ω_Z) , we have $\max_{z \in Z} \text{out}_Z(z) = \max_{z \in Z} \text{in}_Z(z) = \text{diam}(Z)$. Also observe that the minima are achieved above because Z is compact.

Proposition 4.2.7 *The maps out, in, tr, spec, and spec $^\bullet$ are $\text{pow}(\mathbb{R})$ -invariants. Similarly, diam, m^{out} , and m^{in} are \mathbb{R} -invariants.*

This follows from the stronger statement of Proposition 4.3.2.

Definition 4.2.8 (*Motif sets are metric space valued invariants*) Recall the motif sets described in Definition 2.1.13. Our use of motif sets is motivated by the following observation, which appeared in Mémoli (2012, Section 5). For any $n \in \mathbb{N}$, let $\mathcal{C}(\mathbb{R}^{n \times n})$ denote the set of closed subsets of $\mathbb{R}^{n \times n}$. Under the Hausdorff distance induced by the ℓ^∞ metric on $\mathbb{R}^{n \times n}$, this set becomes a valid metric space (Burago et al. 2001, Proposition 7.3.3). The motif sets defined in Definition 2.1.13 define a metric space valued invariant as follows: for each $n \in \mathbb{N}$, let $M_n : \mathcal{CN} \rightarrow \mathcal{C}(\mathbb{R}^{n \times n})$ be the map $Z \mapsto M_n(Z)$. We call this the *motif set invariant*. So for $(Z, \omega_Z), (Y, \omega_Y) \in \mathcal{CN}$, for each $n \in \mathbb{N}$, we let $(U, d_U) = (\mathbb{R}^{n \times n}, \ell^\infty)$ and consider the following distance between the n -motif sets of Z and Y :

$$d_n(M_n(Z), M_n(Y)) := d_H^U(M_n(Z), M_n(Y)).$$

Since d_H is a proper distance between closed subsets, $d_n(M_n(Z), M_n(Y)) = 0$ if and only if $M_n(Z) = M_n(Y)$.

4.3 Quantitative stability of basic invariants

Let (V, d_V) be a given pseudometric space. The V -valued invariant $\iota : \mathcal{CN} \rightarrow V$ is said to be *quantitatively stable* if there exists a constant $L > 0$ such that

$$d_V(\iota(Z), \iota(Y)) \leq L \cdot d_{\mathcal{N}}(Z, Y)$$

for all networks Z and Y . The least constant L such that the above holds for all $Z, Y \in \mathcal{CN}$ is the Lipschitz constant of ι and is denoted $\mathbf{L}(\iota)$.

Note that by identifying a non-constant quantitatively stable V -valued invariant ι , we immediately obtain a lower bound for the $d_{\mathcal{N}}$ distance between any two compact networks (Z, ω_Z) and (Y, ω_Y) . Furthermore, given a finite family $\iota_\alpha : \mathcal{CN} \rightarrow V$,

$\alpha \in A$, of non-constant quantitatively stable invariants, we may obtain the following lower bound for the distance between compact networks Z and Y :

$$d_{\mathcal{N}}(Z, Y) \geq \max_{\alpha \in A} \mathbf{L}(\iota_{\alpha})^{-1} d_V(\iota_{\alpha}(Z), \iota_{\alpha}(Y)).$$

It is often the case that computing $d_V(\iota(Z), \iota(Y))$ is substantially simpler than computing the $d_{\mathcal{N}}$ distance between Z and Y (which leads to a possibly NP-hard problem). The invariants described in the previous section are quantitatively stable. The proofs of the next results are in Sect. 4.8.1.

Lemma 4.3.1 *Let $(Z, \omega_Z), (Y, \omega_Y) \in \mathcal{CN}$. Let f represent any of the maps tr , out , and in , and let f_Z (resp. f_Y) represent the corresponding map $\text{tr}_Z, \text{out}_Z, \text{in}_X$ (resp. $\text{tr}_Y, \text{out}_Y, \text{in}_Y$). Then we obtain:*

$$d_{\mathbf{H}}^{\mathbb{R}}(f(Z), f(Y)) = \inf_{R \in \mathcal{R}(Z, Y)} \sup_{(z, y) \in R} |f_Z(z) - f_Y(y)|.$$

Proposition 4.3.2 *The invariants diam , tr , out , in , m^{out} , and m^{in} are quantitatively stable, with Lipschitz constant $\mathbf{L} = 2$.*

Example 4.3.3 Proposition 4.3.2 provides simple lower bounds for the $d_{\mathcal{N}}$ distance between compact networks. One application is the following: for all networks Z and Y , we have $d_{\mathcal{N}}(Z, Y) \geq \frac{1}{2} |\text{diam}(Z) - \text{diam}(Y)|$. For example, for the networks $Z = N_2(\begin{pmatrix} 1 & 5 \\ 2 & 4 \end{pmatrix})$ and $Y = N_k(\mathbb{1}_{k \times k})$ (the all-ones matrix—also see Example 2.1.12 to recall the N_2 and N_k notation) we have $d_{\mathcal{N}}(Z, Y) \geq \frac{1}{2} |5 - 1| = 2$, for all $k \in \mathbb{N}$. For another example, consider the weight matrices

$$\Sigma := \begin{pmatrix} 0 & 5 & 2 \\ 3 & 1 & 4 \\ 1 & 4 & 3 \end{pmatrix} \quad \text{and} \quad \Sigma' := \begin{pmatrix} 3 & 4 & 2 \\ 3 & 1 & 5 \\ 3 & 3 & 4 \end{pmatrix}.$$


Let $Z = N_3(\Sigma)$ and $Y = N_3(\Sigma')$. By comparing the diagonals, we can easily see that $Z \not\cong^s Y$, but let us see how the invariants we proposed can help. Note that $\text{diam}(Z) = \text{diam}(Y) = 5$, so the lower bound provided by diameter ($\frac{1}{2} |5 - 5| = 0$) does not help in telling the networks apart. However, $\text{tr}(Z) = \{0, 1, 3\}$ and $\text{tr}(Y) = \{3, 1, 4\}$, and Proposition 4.3.2 then yields

$$d_{\mathcal{N}}(Z, Y) \geq \frac{1}{2} d_{\mathbf{H}}^{\mathbb{R}}(\{0, 1, 3\}, \{1, 3, 4\}) = \frac{1}{2}.$$

Consider now the out and in maps. Note that one has $\text{out}(Z) = \{5, 4\}$, $\text{out}(Y) = \{4, 5\}$, $\text{in}(Z) = \{3, 5, 4\}$, and $\text{in}(Y) = \{3, 4, 5\}$. Then $d_{\mathbf{H}}^{\mathbb{R}}(\text{out}(Z), \text{out}(Y)) = 0$, and $d_{\mathbf{H}}^{\mathbb{R}}(\text{in}(Z), \text{in}(Y)) = 0$. Thus in both cases, we obtain $d_{\mathcal{N}}(Z, Y) \geq 0$. So in this particular example, the out and in maps are not useful for obtaining a lower bound to $d_{\mathcal{N}}(X, Y)$ via Proposition 4.3.2.

Now we state a proposition regarding the stability of global and local spectrum invariants. These will be of particular interest for computational purposes as we explain in Sect. A.1.

Fig. 15 Lower-bounding $d_{\mathcal{N}}$ by using global spectra (cf. Example 4.3.5)



Proposition 4.3.4 Let spec^\bullet refer to either the out or in version of local spectrum. Then, for all $(Z, \omega_Z), (Y, \omega_Y) \in \mathcal{CN}$ we have

$$\begin{aligned} d_{\mathcal{N}}(Z, Y) &\geq \frac{1}{2} \inf_{R \in \mathcal{R}} \sup_{(z, y) \in R} d_{\mathbb{H}}^{\mathbb{R}}(\text{spec}_Z^\bullet(z), \text{spec}_Y^\bullet(y)) \\ &\geq \frac{1}{2} d_{\mathbb{H}}^{\mathbb{R}}(\text{spec}(Z), \text{spec}(Y)). \end{aligned}$$

As a corollary, we get $L(\text{spec}^\bullet) = L(\text{spec}) = 2$.

Example 4.3.5 (An application of Proposition 4.3.4) Consider the networks in Fig. 15. By Proposition 4.3.4, we may calculate a lower bound for $d_{\mathcal{N}}(Z, Y)$ by simply computing the Hausdorff distance between $\text{spec}(Z)$ and $\text{spec}(Y)$, and dividing by 2. In this example, $\text{spec}(Z) = \{1, 2\}$ and $\text{spec}(Y) = \{1, 2, 3\}$. Thus $d_{\mathbb{H}}^{\mathbb{R}}(\text{spec}(Z), \text{spec}(Y)) = 1$, and $d_{\mathcal{N}}(Z, Y) \geq \frac{1}{2}$.

Computing the lower bound involving local spectra requires solving a bottleneck linear assignment problem over the set of all correspondences between Z and Y . This can be solved in polynomial time; details are provided in Sect. A.1. The second lower bound stipulates computing the Hausdorff distance on \mathbb{R} between the (global) spectra of Z and Y —a computation which can be carried out in (smaller) polynomial time as well.

To conclude this section, we state a theorem asserting that motif sets form a family of quantitatively stable invariants. Recall that we used this result in the proof of Theorem 3.2.2.

Theorem 4.3.6 For each $n \in \mathbb{N}$, M_n is a stable invariant with $L(M_n) = 2$.

4.4 Persistent homology invariants and their convergence

Persistent homology is a technique for studying multiscale features of data by creating geometric complexes, associating algebraic invariants to these complexes, and studying the structure of these invariants (Robins 1999; Carlsson 2009; Edelsbrunner and Harer 2010; Edelsbrunner and Morozov 2014; Frosini 1992). In the conventional setting, persistent homology takes Euclidean or metric data as input and produces topological summaries called *persistence diagrams* or *barcodes* as output, one in each dimension $k \in \mathbb{Z}_+$. In an intermediate step, this technique first constructs a collection of vector spaces with linear maps, and these vector spaces are implicitly built on top of geometric complexes. The structure of these linear maps is then studied via matrix

reductions and summarized in persistence barcodes. Persistent homology research has seen rapid progress in recent years, and in particular, the notion of persistent homology for general, possibly asymmetric settings such as that of networks has been studied in Turner (2019), Chowdhury and Mémoli (2018b, c), Chowdhury et al. (2019), Pinto (2020), Huntsman (2020), and Dey et al. (2020). In these works, persistent homology methods are defined for finite networks.

Here and in the next section, we use our machinery of ε -systems (cf. Definition 3.1.2 and subsequent constructions) to prove that persistence diagrams of infinite, compact networks are well-defined. This answers a question posed by Turner in Turner (2019).

We begin with the definition of a *persistent vector space*. Throughout this section, all our vector spaces are assumed to be over a fixed ground field \mathbb{F} .

Definition 4.4.1 A *persistent vector space* \mathcal{V} (also denoted **PVec**) is a family $\{V^\delta \xrightarrow{v_{\delta,\delta'}} V^{\delta'}\}_{\delta \leq \delta' \in \mathbb{R}}$ of vector spaces and linear maps such that: (1) $v_{\delta,\delta}$ is the identity map for any $\delta \in \mathbb{R}$, and (2) $v_{\delta,\delta''} = v_{\delta',\delta''} \circ v_{\delta,\delta'}$ whenever $\delta \leq \delta' \leq \delta''$.

Conventional hierarchical clustering methods take in metric data as input and produce ultrametrics as output that are in turn faithfully visualized as dendrograms (Carlsson and Mémoli 2010). A conventional persistent homology method (e.g. Vietoris–Rips, defined below) yields a higher dimensional analogue of this process: it takes a metric dataset as input, and outputs a persistent vector space \mathcal{V} that is faithfully represented as a *persistence diagram* $\text{Dgm}(\mathcal{V})$. A classification result in Carlsson et al. (2005, §5.2) shows that the persistence diagram is a full invariant of a persistent vector space. This completes the analogy with the setting of hierarchical clustering.

Persistence diagrams can be compared using the *bottleneck distance*, which we denote by d_B . We point the reader to Chazal et al. (2016) and references therein for details.

While the persistence diagram and bottleneck distance are the primary tools in practical applications, theoretical proofs are often made simpler through the language of *interleavings* and *interleaving distance*. We present this next.

Definition 4.4.2 (ε -interleaving, Chazal et al. 2009a) Let $\mathcal{U} = \{U^\delta \xrightarrow{s_{\delta,\delta'}} U^{\delta'}\}_{\delta \leq \delta' \in \mathbb{R}}$ and $\mathcal{V} = \{V^\delta \xrightarrow{t_{\delta,\delta'}} V^{\delta'}\}_{\delta \leq \delta' \in \mathbb{R}}$ be two persistent vector spaces. Given $\varepsilon \geq 0$, \mathcal{U} and \mathcal{V} are said to be ε -interleaved if there exist two families of linear maps $\{\varphi_\delta : U^\delta \rightarrow V^{\delta+\varepsilon}\}_{\delta \in \mathbb{R}}$ and $\{\psi_\delta : V^\delta \rightarrow U^{\delta+\varepsilon}\}_{\delta \in \mathbb{R}}$ such that: (1) $\varphi_{\delta'} \circ s_{\delta,\delta'} = t_{\delta+\varepsilon,\delta'+\varepsilon} \circ \varphi_\delta$, (2) $\psi_{\delta'} \circ t_{\delta,\delta'} = s_{\delta+\varepsilon,\delta'+\varepsilon} \circ \psi_\delta$, (3) $s_{\delta,\delta+2\varepsilon} = \psi_{\delta+\varepsilon} \circ \varphi_\delta$, and (4) $t_{\delta,\delta+2\varepsilon} = \varphi_{\delta+\varepsilon} \circ \psi_\delta$ for each $\delta \leq \delta' \in \mathbb{R}$.

The *interleaving distance* d_I between \mathcal{U} and \mathcal{V} is then defined as:

$$d_I(\mathcal{U}, \mathcal{V}) := \inf\{\varepsilon \geq 0 : \mathcal{U} \text{ and } \mathcal{V} \text{ are } \varepsilon\text{-interleaved}\}.$$

The interleaving and bottleneck distances are connected by the *Isometry Theorem*, which states that the two distances are in fact equivalent. Various forms of this theorem have appeared in the literature; we will end this section with a statement of this result that appears in Chazal et al. (2016).

Our aim in this work is to describe the *convergence* of persistent homology methods applied to network data. When dealing with finite networks, the vector spaces resulting from applying a persistent homology method will necessarily be finite dimensional. However, our setting is that of infinite (more specifically, *compact*) networks, and so we need additional machinery to ensure that our methods output well-defined persistent vector spaces. The following definition and theorem are provided in full detail in Chazal et al. (2016).

Definition 4.4.3 (§2.1, Chazal et al. 2016) A persistent vector space $\mathcal{V} = \{V^\delta \xrightarrow{v_{\delta,\delta'}} V^{\delta'}\}_{\delta \leq \delta' \in \mathbb{R}}$ is *q-tame* if $v_{\delta,\delta'}$ has finite rank whenever $\delta < \delta'$.

Theorem 4.4.4 (Chazal et al. 2016, also Chazal et al. 2014 Theorem 2.3) Any *q-tame persistent vector space* \mathcal{V} has a well-defined persistence diagram $\text{Dgm}(\mathcal{V})$. If \mathcal{U}, \mathcal{V} are ε -interleaved *q-tame persistent vector spaces*, then $d_B(\text{Dgm}(\mathcal{U}), \text{Dgm}(\mathcal{V})) \leq \varepsilon$.

We conclude this section with a statement of the isometry theorem.

Theorem 4.4.5 (Theorem 5.14, Chazal et al. 2016) Let \mathcal{U}, \mathcal{V} be *q-tame persistent vector spaces*. Then,

$$d_1(\mathcal{U}, \mathcal{V}) = d_B(\mathcal{U}, \mathcal{V}).$$

4.5 Vietoris–Rips, Dowker, path, and ordered tuple persistent homology methods on networks

We now present a collection of methods for producing persistent vector spaces from network data. For finite networks, these methods have appeared in Chowdhury and Mémoli (2018b), and Turner (2019). Here our goal is to define these methods for compact networks and to establish their convergence properties. First we will focus on defining the different types of *filtered complexes*—i.e. complexes nested by inclusion—from which persistent vector spaces can be derived. Once the complexes are defined, the persistent vector spaces are obtained by applying the *homology functor* (cf. Munkres 1984 for details on homology). Recall that an *abstract simplicial complex* on a set Z is a collection Σ of non-empty finite subsets of Z such that whenever $\sigma \in \Sigma$ and $\tau \subseteq \sigma$, we also have $\tau \in \Sigma$. A *filtered simplicial complex* is a nested collection of simplicial complexes $\{\Sigma_\delta \subseteq \Sigma_{\delta'}\}_{\delta \leq \delta'}$.

Definition 4.5.1 (*Vietoris–Rips complexes (VR)*) Given a compact network (Z, ω_Z) and $\delta \in \mathbb{R}$, the *Vietoris–Rips complex at resolution δ* is defined as:

$$\text{VR}_\delta(Z) := \left\{ \sigma \in \text{pow}(Z) : \sigma \text{ finite, } \max_{z, z' \in \sigma} \omega_Z(z, z') \leq \delta \right\}.$$

Practitioners of persistent homology will recognize that the Vietoris–Rips complex construction given above is a direct generalization of the Vietoris–Rips complex of a metric space. This definition yields a simplicial filtration $\{\text{VR}_\delta(Z) \hookrightarrow$

$\mathbf{VR}_{\delta'}(Z)\}_{\delta \leq \delta' \in \mathbb{R}}$. Applying the simplicial homology functor in dimension k (for $k \in \mathbb{Z}_+$) to this filtration yields the *Vietoris–Rips persistent vector space* $\mathbf{PVec}_k^{\mathbf{VR}}(Z)$.

Next we describe two constructions—the Dowker source and sink complexes—that practitioners will recognize as asymmetric generalizations of the Čech complex of a metric space (Chazal et al. 2014; Chowdhury and Mémoli 2018b).

Definition 4.5.2 (*Čech complex (\check{C}) of a metric space*) Given a metric space (Z, d_Z) and $\delta \in \mathbb{R}$, the *Čech complex at resolution δ* is defined as:

$$\check{C}_\delta(Z) := \left\{ \sigma \in \text{pow}(Z) : \sigma \text{ finite, } \min_{p \in Z} \max_{z \in \sigma} d_Z(z, p) \leq \delta \right\}.$$

Definition 4.5.3 (*Dowker complexes (**Dow**)*) Given a compact network (Z, ω_Z) and $\delta \in \mathbb{R}$, the *Dowker sink-complex at resolution δ* is defined as:

$$\mathbf{Dow}_\delta^{\text{si}}(Z) := \left\{ \sigma \in \text{pow}(Z) : \sigma \text{ finite, } \min_{p \in Z} \max_{z \in \sigma} \omega_Z(z, p) \leq \delta \right\}.$$

Similarly, the *Dowker source-complex at resolution δ* is defined as:

$$\mathbf{Dow}_\delta^{\text{so}}(Z) := \left\{ \sigma \in \text{pow}(Z) : \sigma \text{ finite, } \min_{p \in Z} \max_{z \in \sigma} \omega_Z(p, z) \leq \delta \right\}.$$

The Dowker sink and source complexes are different in general when X is asymmetric. Surprisingly, the persistent vector spaces obtained from the sink and source filtrations *are equivalent*. This result was established in Chowdhury and Mémoli (2018b) in the setting of finite networks. For compact networks, the statement is as follows. Here we momentarily denote the Dowker source and sink persistent vector spaces as $\mathbf{PVec}^{\text{so}}$ and $\mathbf{PVec}^{\text{si}}$, respectively.

Theorem 4.5.4 (*Dowker duality*) Let (Z, ω_Z) be a compact network, and let $k \in \mathbb{Z}_+$. Then,

$$\mathbf{PVec}_k^{\text{si}}(Z) = \mathbf{PVec}_k^{\text{so}}(Z).$$

The proof is via a functorial generalization of Dowker’s Theorem (1952), which holds in the case of infinite sets. Alternatively, a functorial generalization of the Nerve Lemma can also be used to prove this result, as suggested in Chazal et al. (2014). Hence we denote the resulting persistent vector space (in dimension $k \in \mathbb{Z}_+$) as $\mathbf{PVec}_k^{\mathbf{Dow}}(X)$, without distinguishing between sink and source constructions.

Whereas the Vietoris–Rips and Dowker complexes generate conventional simplicial complexes, it is possible to build more general complexes that better capture asymmetry in the function ω_Z . One example is via *ordered-tuple (OT) complexes* (Turner 2019). An **OT**-complex Σ over a set Z is a collection of nonempty, finite tuples (z_1, \dots, z_p) such that whenever $(z_1, \dots, z_p) \in \Sigma$, we also have $(z_1, \dots, \hat{z}_i, \dots, z_p) \in \Sigma$ for all $1 \leq i \leq p$, where the hat denotes omission from the sequence. Notions related to homology can be defined for OT complexes without

difficulty, and thus a persistent vector space can be immediately obtained after defining a filtered OT complex. The following construction was phrased in Turner (2019) in terms of set-function pairs, which are just finite networks in our setting.

Definition 4.5.5 (*Vietoris–Rips OT complexes (OT)*) Given a compact network (Z, ω_Z) and $\delta \in \mathbb{R}$, the *Vietoris–Rips OT complex* at resolution δ is defined as:

$$\mathbf{OT}_\delta^{\mathbf{VR}}(Z) := \{(z_1, z_2, \dots, z_p) : \omega_Z(z_i, z_j) \leq \delta \text{ for all } i \leq j\}.$$

We denote the corresponding persistent vector space by $\mathbf{PVec}_k^{\mathbf{OT}}(Z)$.

Yet another approach is via the notion of *path homology* (Grigor'yan et al. 2014; Chowdhury and Mémoli 2018c). A preliminary notion that we will need is that of a *boundary operator*. Suppose we are given a set Z and the free vector space of finite-length tuples in Z , i.e. the vector space $\mathbb{F}[\{(z_0, z_1, \dots, z_k) : k \in \mathbb{Z}_+, z_i \in Z \text{ for all } i\}]$. Then one defines a boundary operator via the following alternating sum:

$$\partial_k((z_0, z_1, \dots, z_k)) = \sum_{i=0}^k (-1)^i (z_0, z_1, \dots, \hat{z}_i, \dots, z_k).$$

Definition 4.5.6 (*Persistent path homology (Path)*) Given a compact network (Z, ω_Z) and $\delta \in \mathbb{R}$, first construct the digraph G^δ with $V(G^\delta) := Z$ and $E(G^\delta) := \{(v, v') : \omega_Z(v, v') \leq \delta\}$. Next, for each $k \in \mathbb{Z}_+$, define the free vector space of *allowed k -paths* $\mathcal{A}_k(G^\delta) := \mathbb{F}[(v_0, v_1, \dots, v_k) : (v_i, v_{i+1}) \in E(G^\delta) \text{ for each } 0 \leq i \leq k-1]$. Finally define $\Omega_k(G^\delta) := \mathbb{F}[\{u \in \mathcal{A}_k(G^\delta) : \partial_k(u) \in \mathcal{A}_{k-1}(G^\delta)\}]$. Then the degree- k path homology vector space of G^δ is given by $H_k^\pm(G^\delta) := \ker(\partial_k)/\text{im}(\partial_{k+1})$, and the inclusions $\{G^\delta \subseteq G^{\delta'}\}_{\delta \leq \delta'}$ are transformed into a persistent vector space by functoriality of homology. We denote the corresponding persistent vector space by $\mathbf{PVec}_k^{\mathbf{Path}}(Z)$.

The following lemma collects results from Chowdhury and Mémoli (2018b, c), and Turner (2019) on the *stability of persistent homology invariants* of networks.

Lemma 4.5.7 (Quantitative stability between d_N and d_1) *Let (Z, ω_Z) and (Y, ω_Y) be two networks. Let $\varepsilon > 2d_N(X, Y)$. Then $\mathbf{PVec}_k^*(X)$ and $\mathbf{PVec}_k^*(Y)$ are ε -interleaved, where $*$ denotes any of the methods in $\{\mathbf{VR}, \mathbf{Dow}, \mathbf{Path}, \mathbf{OT}\}$.*

Theorem 4.5.8 *Let $(Z, \omega_Z) \in \mathcal{CN}$, $k \in \mathbb{Z}_+$. Then $\mathbf{PVec}_k^*(Z)$ is q -tame, where $*$ denotes any of the methods in $\{\mathbf{VR}, \mathbf{Dow}, \mathbf{Path}, \mathbf{OT}\}$.*

The metric space analogue of Theorem 4.5.8 appeared in Chazal et al. (2014, Proposition 5.1); the same proof structure works in the setting of networks after applying our results on approximation via ε -systems.

Proof of Theorem 4.5.8 For convenience, write $\mathbf{PVec}_k^*(Z) = \{V^\delta \xrightarrow{\nu_{\delta, \delta'}} V^{\delta'}\}_{\delta \leq \delta' \in \mathbb{R}}$. Let $\delta < \delta'$. We need to show $\nu_{\delta, \delta'}$ has finite rank. Write $\varepsilon := (\delta' - \delta)/2$. Let \mathcal{U} be an $\varepsilon/4$ -system on Z (this requires Theorem 3.1.3). Then by Theorem 3.1.3 we pick

a finite subset $Z' \subseteq Z$ such that $d_{\mathcal{N}}(Z, Z') < \varepsilon/2$. Then $\mathbf{PVec}_k^*(Z')$ and $\mathbf{PVec}_k^*(Z)$ are ε -interleaved by Lemma 4.5.7. For convenience, write $\mathbf{PVec}_k^*(Z') = \{U^\delta \xrightarrow{\mu_{\delta, \delta'}} U^{\delta'}\}_{\delta \leq \delta' \in \mathbb{R}}$. Then the map $v_{\delta, \delta'} : V^\delta \rightarrow V^{\delta'}$ factorizes through $U^{\delta+\varepsilon}$ via interleaving maps $V^\delta \rightarrow U^{\delta+\varepsilon} \rightarrow V^{\delta+2\varepsilon} = V^{\delta'}$. Since $U^{\delta+\varepsilon}$ is finite dimensional, it follows that $v_{\delta, \delta'}$ has finite rank. This concludes the proof. \square

Corollary 4.5.9 (Stability) *Let $(Z, \omega_Z), (Y, \omega_Y) \in \mathcal{CN}$, $k \in \mathbb{Z}_+$. Then for any $*$ in $\{\mathbf{VR}, \mathbf{Dow}, \mathbf{Path}, \mathbf{OT}\}$,*

$$d_B(\mathrm{Dgm}_k^*(Z), \mathrm{Dgm}_k^*(Y)) \leq 2d_{\mathcal{N}}(X, Y).$$

Proof By Theorem 4.5.8, $\mathbf{PVec}^*(Z), \mathbf{PVec}^*(Y)$ are both tame. Thus they have well-defined persistence diagrams (Theorem 4.4.4). The result follows by Lemma 4.5.7 and Theorem 4.4.4. \square

The following result presents a probabilistic version of the prior result. Probabilistic approaches have practical value due to the inevitable measurement error found in network data (Newman 2018b, a).

Theorem 4.5.10 (Convergence) *Let (Z, ω_Z, μ_Z) be a measure network. For each $i \in \mathbb{N}$, let $z_i : \Omega \rightarrow Z$ be an independent random variable defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with distribution μ_Z . For each $n \in \mathbb{N}$, let $\mathcal{Z}_n = \{z_1, z_2, \dots, z_n\}$. Let $\varepsilon > 0$. Then we have:*

$$\mathbb{P}(\{\omega \in \Omega : d_B(\mathrm{Dgm}^*(\mathrm{supp}(\mu_Z)), \mathrm{Dgm}^*(\mathcal{Z}_n(\omega))) \geq \varepsilon\}) \leq \frac{(1 - \mathfrak{M}_{\varepsilon/4}(\mathrm{supp}(\mu_Z)))^n}{\mathfrak{M}_{\varepsilon/4}(\mathrm{supp}(\mu_Z))},$$


where $\mathcal{Z}_n(\omega)$ is the subnetwork induced by $\{z_1(\omega), \dots, z_n(\omega)\}$ and $*$ belongs to $\{\mathbf{VR}, \mathbf{Dow}, \mathbf{Path}, \mathbf{OT}\}$. In particular, $\mathbf{PVec}^*(\mathcal{Z}_n)$ converges almost surely to $\mathbf{PVec}^*(\mathrm{supp}(\mu_Z))$ w.r.t. bottleneck distance.

Proof of Theorem 4.5.10 We can consider $\mathrm{supp}(\mu_Z)$ as a network with full support by endowing it with the restriction of ω_Z to $\mathrm{supp}(\mu_Z) \times \mathrm{supp}(\mu_Z)$. So for convenience, assume $Z = \mathrm{supp}(\mu_Z)$. For any $\omega \in \Omega$ such that $d_{\mathcal{N}}(Z, \mathcal{Z}_n(\omega)) < \varepsilon/2$, we have by Corollary 4.5.9 that $d_B(\mathrm{Dgm}^*(Z), \mathrm{Dgm}^*(\mathcal{Z}_n)) < \varepsilon$. Thus by applying Theorem 4.1.4, we have:

$$\begin{aligned} \mathbb{P}(\{\omega \in \Omega : d_B(\mathrm{Dgm}^*(Z), \mathrm{Dgm}^*(\mathcal{Z}_n(\omega))) \geq \varepsilon\}) &\leq \mathbb{P}(\{\omega \in \Omega : d_{\mathcal{N}}(Z, \mathcal{Z}_n(\omega)) \geq \varepsilon/2\}) \\ &\leq \frac{(1 - \mathfrak{M}_{\varepsilon/4}(Z))^n}{\mathfrak{M}_{\varepsilon/4}(Z)}. \end{aligned}$$

We conclude the proof with an application of the Borel–Cantelli lemma, as in the proof of Theorem 4.1.4. \square

Fig. 16 The framework of quantitatively stable invariants allows us to produce easily computable proxies for $d_{\mathcal{N}}$. Apart from the hierarchical clustering map \mathcal{H} , all the other maps embed network features into simple spaces with metrics that can be computed in polynomial time. The \mathcal{H} map embeds an input compact network into $\mathcal{CN}^{\text{ult}}$ (cf. Definition 2.1.7), which can then be displayed as a dendrogram-like object (Smith et al. 2016)



4.6 Hierarchical clustering methods on asymmetric networks

We now provide similar results in the special case of hierarchical clustering for asymmetric networks. First we state some definitions relating to connectivity and induce networks from preexisting partitions of a network. Then we will be ready to define hierarchical clustering methods on asymmetric networks and to prove related convergence results. The reciprocal and nonreciprocal clustering methods described below were introduced in a slightly restricted setting (finite networks (Z, ω_Z) with zero diagonal and nonnegative ω_Z) in Carlsson et al. (2013). The stability of these methods, given in the form $d_{\mathcal{N}}(\mathcal{H}(Y, \omega_Y), \mathcal{H}(Z, \omega_Z)) \leq d_{\mathcal{N}}(Y, Z)$ where \mathcal{H} denotes either one of these hierarchical clustering methods, was also shown in Carlsson et al. (2013).

Applying either of the hierarchical clustering methods we describe next to a network (Z, ω_Z) produces a new weight function u_Z that satisfies the ultrametric inequality (cf. Definition 2.1.7). Thus we get maps of the form $\mathcal{H} : \mathcal{CN} \rightarrow \mathcal{CN}^{\text{ult}}$ as shown in Fig. 16.

Before proceeding, we make some historical remarks and explain the connection to the next few sections. The methods described next will be reminiscent of single linkage hierarchical clustering (SLHC). In the Euclidean case, it has been known for some time that SLHC (Hartigan 1981, 1985) is not sensitive to modes of the data distribution. This observation was refined in Carlsson and Mémoli (2010) where it was shown that in the case of metric measure spaces, SLHC depended on only the support of the measure. Despite this drawback, SLHC has long been of mathematical interest (Hartigan 1985). In more recent years, it has been shown that under certain axioms that characterize the well-behavedness of an HC method for metric spaces, every well-behaved HC method factors through SLHC (Carlsson and Mémoli 2013). This thread of axiomatic characterization has been further developed in the setting of finite, asymmetric networks in Carlsson et al. (2013, 2018) and most recently in Carlsson et al. (2021). The reciprocal and nonreciprocal clustering methods we describe next are key players in these axiomatic treatments of directed networks, as they constitute extremal examples that bound an infinite family of axiomatic HC methods on directed networks. Over the next few sections, we treat the reciprocal and

nonreciprocal methods in the infinite network setting and connect back to results in Carlsson and Mémoli (2010) by showing that such methods remain sensitive only to the support of the data distribution.

4.6.1 Chain cost and path-connectedness

Definition 4.6.1 (*The modified ω -weights and ω -path-connectedness*) Given a network (Z, ω_Z) , one defines new weight functions $\tilde{\omega}_Z : Z \times Z \rightarrow \text{pow}(\mathbb{R})$ and $\overline{\omega}_Z, \underline{\omega}_Z : Z \times Z \rightarrow \mathbb{R}$ by writing the following for $z, z' \in Z$:

$$\begin{aligned}\tilde{\omega}_Z(z, z') &:= \{\omega_Z(z, z), \omega_Z(z, z'), \omega_Z(z', z')\}, \\ \overline{\omega}_Z(z, z') &:= \max \tilde{\omega}_Z(z, z'), \quad \text{and} \\ \underline{\omega}_Z(z, z') &:= \min \tilde{\omega}_Z(z, z').\end{aligned}$$

To say that (Z, ω_Z) is ω -path-connected means that given any $x, x' \in X$, there exists $r_{x,x'} \in X$ and a continuous function $\gamma : [0, 1] \rightarrow X$ such that $\gamma(0) = x$, $\gamma(1) = x'$, and for any $\varepsilon > 0$, there exist $0 = t_0 \leq t_1 \leq t_2 \leq \dots \leq t_n = 1$ for $n \in \mathbb{Z}_+$ such that:

$$\tilde{\omega}_X(\gamma(t_i), \gamma(t_{i+1})) \subseteq (r_{x,x'} - \varepsilon, r_{x,x'} + \varepsilon) \text{ for each } 0 \leq i \leq n-1. \quad (12)$$

Remark 4.6.2 When (Z, ω_Z) is a metric space, all the self weights $\omega_Z(z, z)$ are zero. So one would have $\overline{\omega}_Z = \omega_Z$ and $\underline{\omega}_Z = 0$. The notion of ω -path-connectedness agrees with the standard notion of path-connectedness in a metric space, as we elaborate in Lemma 4.6.4 below. In general, however, if $\gamma(t)$ connecting x to x' satisfies (12), because of the asymmetry of ω_X , it does not follow that the reverse curve $\gamma(1-t)$ connecting x' to x will satisfy (12).

Example 4.6.3 Similarity kernels, e.g. the cosine similarity or the Gaussian kernel, constitute simple examples where $\omega_Z(z, z)$ is nonzero. Such examples abound in machine learning applications. The unit circle equipped with a Gaussian kernel is an example of a ω -path-connected space with nonzero self weights.

Lemma 4.6.4 *Let (Z, ω_Z) be a ω -path-connected network. Then there exists a unique $r_Z \in \mathbb{R}$ such that $\omega_Z(z, z) = r_Z$ for all $z \in Z$. In the case of metric spaces, one has $r_Z = 0$.*

Proof of Lemma 4.6.4 Let $z, z' \in Z$ and let $r_{z,z'} \in \mathbb{R}$ be as in Definition 4.6.1. Let $(\varepsilon_n)_{n \in \mathbb{N}}$ be a sequence decreasing to 0. Fix $n \in \mathbb{N}$, and let $\gamma : [0, 1] \rightarrow X$ be a continuous function such that $\gamma(0) = z$, $\gamma(1) = z'$, and there exist $t_0^n = 0 \leq t_1^n \leq t_2^n \leq \dots, t_{k_n}^n = 1$ such that:

$$\tilde{\omega}_Z(\gamma(t_i^n), \gamma(t_{i+1}^n)) \subseteq B(r_{z,z'}, \varepsilon_n) \quad \text{for each } 0 \leq i \leq k_n - 1.$$

In particular, because $\tilde{\omega}_Z(z, \gamma(t_1^n)) \subseteq B(r_{z,z'}, \varepsilon_n)$ and $\tilde{\omega}_Z(\gamma(t_{k_n-1}^n), z') \subseteq B(r_{z,z'}, \varepsilon_n)$, we have:

$$\{\omega_Z(z, z), \omega_Z(z', z')\} \subseteq B(r_{z,z'}, \varepsilon_n).$$

Thus $|\omega_Z(z, z) - \omega_Z(z', z')| \leq 2\varepsilon_n$. Letting $n \rightarrow \infty$, we obtain $\omega_Z(z, z) = \omega_Z(z', z')$. Since $z' \in Z$ was arbitrary, we get that $\omega_Z(z', z') = \omega_Z(z, z)$ for all $z' \in Z$. The result now follows. \square

Definition 4.6.5 (*Path-connectivity constant*) Let (Z, ω_Z) be a ω -path-connected network. Then we define its *path-connectivity constant* pc_Z to be the real number r_Z obtained via Lemma 4.6.4.

Definition 4.6.6 (*Networks arising from disconnected networks*) Let (Z, ω_Z) be a network such that Z is a finite, disjoint union of ω -path-connected components $\{U_a : a \in A\}$, where A is a (finite) indexing set and each U_a is compact. Let $\nu_A : A \times A \rightarrow \mathbb{R}$ be the map given by writing, for each $a, a' \in A$,

$$\nu_A(a, a') := \min \{ \overline{\omega}_Z(z, z') : z \in U_a, z' \in U_{a'} \}.$$

Then (A, ν_A) is a network. Analogously, one induces a *symmetric* network by defining λ_A as follows:

$$\lambda_A(a, a') := \min \{ \max(\overline{\omega}_Z(z, z'), \overline{\omega}_Z(z', z)) : z \in U_a, z' \in U_{a'} \}.$$

The following definition will be useful in the next section.

Definition 4.6.7 (*Chains and directed cost*) A *chain* c from z to z' is defined to be a finite ordered set of points starting at z and reaching z' :

$$c = \{z_0, z_1, z_2, \dots, z_n : z_0 = z, z_n = z', z_i \in Z \text{ for all } i\}.$$

The collection of all chains from z to z' will be denoted $C_Z(z, z')$. The (*directed*) *cost* of a chain $c \in C_Z(z, z')$ is defined as follows: $\text{cost}_Z(c) := \max_{z_i, z_{i+1} \in c} \overline{\omega}_z(z_i, z_{i+1})$.

Remark 4.6.8 (*Equivalence of ultrametrics and dendrograms*) Before proceeding to the next section, we remind the reader that any ultrametric has a lossless representation as a dendrogram, and conversely, any dendrogram has a lossless representation as an ultrametric (Jardine and Sibson 1971). By virtue of this result, we write the outputs of hierarchical clustering methods as ultrametrics. As shown in Smith et al. (2016), a similar duality holds even in the setting of (asymmetric) networks, up to a small modification of definitions. In particular, the output of an HC method on a network is a network in itself, along with some special structure that allows it to be visualized as a (generalized) dendrogram.

4.6.2 The nonreciprocal clustering method: definition and convergence

We now present the nonreciprocal hierarchical clustering method for directed networks.

Definition 4.6.9 (Nonreciprocal clustering) The *nonreciprocal clustering method* (\mathcal{H}^{NR}) is a map $\mathcal{H}^{\text{NR}} : \mathcal{CN} \rightarrow \mathcal{CN}$ given by $(Z, \omega_Z) \mapsto (Z, u_Z^{\text{NR}})$, where $u_Z^{\text{NR}} : Z \times Z \rightarrow \mathbb{R}$ is defined by writing, for each $z, z' \in Z$,

$$u_Z^{\text{NR}}(z, z') := \max \left(\inf_{c \in C_Z(z, z')} \text{cost}_Z(c), \inf_{c \in C_Z(z', z)} \text{cost}_Z(c) \right).$$

The output u_X^{NR} is symmetric and satisfies the *ultrametric inequality*, so it can be represented as a tree (Semple and Steel 2003, §7.2). Compare this to the *cluster trees* discussed by Hartigan (1975). The idea behind this definition is easily summarized: two points z and z' belong to the same cluster at resolution δ if there are directed paths $z \rightarrow z'$ and $z' \rightarrow z$, each with cost $\leq \delta$.

The next result shows that nonreciprocal clustering essentially recovers only the structure of the support of the data distribution, and no information about the data density. Such a result is also known for metric spaces (Carlsson and Mémoli 2010); the proof for the current result uses the machinery of ε -systems. Proofs related to the next result are provided in Sect. 4.8.2.

Theorem 4.6.10 (Convergence of nonreciprocal clustering) *Let (Z, ω_Z, μ_Z) be a measure network. Suppose $\text{supp}(\mu_Z)$ is a finite, disjoint union of compact, ω -path-connected components $\{Z_a : a \in A\}$. Let (A, v_A) be as in Definition 4.6.6, and let $(A, u_A^{\text{NR}}) = \mathcal{H}^{\text{NR}}(A, v_A)$. For each $i \in \mathbb{N}$, let $z_i : \Omega \rightarrow Z$ be an independent random variable defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with distribution μ_Z . For each $n \in \mathbb{N}$, let $Z_n = \{z_1, z_2, \dots, z_n\}$, and for each $\omega \in \Omega$, let $Z_n(\omega)$ denote the subnetwork induced by $\{z_1(\omega), \dots, z_n(\omega)\}$. Let $\varepsilon > 0$. Then,*

$$\mathbb{P}(\{\omega \in \Omega : d_{\mathcal{N}}((A, u_A^{\text{NR}}), \mathcal{H}^{\text{NR}}(Z_n(\omega))) \geq \varepsilon\}) \leq \frac{(1 - \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)))^n}{\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z))}.$$

In particular, the output of the nonreciprocal clustering method applied to the sampled network Z_n converges almost surely to (A, u_A^{NR}) in the sense of $d_{\mathcal{N}}$ as the sample size increases.

We end this section with an application of nonreciprocal clustering to Finsler manifolds.

Proposition 4.6.11 (Nonreciprocal clustering on Finsler manifolds) *Let (M, F_M, ω_M) be a compact, connected Finsler manifold without boundary, where ω_M is the asymmetric weight function induced by the Finsler function F_M . Then $u_M^{\text{NR}}(z, z') = 0$ for all $z, z' \in M$.*

Proof of Proposition 4.6.11 Let $z, z' \in M$. Let $\varepsilon > 0$, and let $\gamma, \gamma' : [0, 1] \rightarrow M$ be curves from z to z' and from z' to z , respectively. By choosing n uniformly separated points on $\gamma([0, 1])$ and $\gamma'([0, 1])$ for sufficiently large n , we obtain finite chains c and c' on $\gamma([0, 1])$ and $\gamma'([0, 1])$ such that $\max(\text{cost}_M(c), \text{cost}_M(c')) < \varepsilon$. Since $\varepsilon > 0$ was arbitrary, we obtain $u_M^{\text{NR}}(z, z') = 0$. \square

4.6.3 The reciprocal clustering method: definition and convergence

Definition 4.6.12 (*Reciprocal clustering*) The *reciprocal clustering method* (\mathcal{H}^R) is a map $\mathcal{H}^R : \mathcal{CN} \rightarrow \mathcal{CN}$ given by $(Z, \omega_Z) \mapsto (Z, u_Z^R)$, where $u_Z^R : Z \times Z \rightarrow \mathbb{R}$ is defined by writing, for each $z, z' \in Z$,

$$u_Z^R(z, z') := \inf_{c \in C_Z(z, z')} \max_{z_i, z_{i+1} \in c} (\max(\bar{\omega}_Z(z_i, z_{i+1}), \bar{\omega}_Z(z_{i+1}, z_i))). \quad (13)$$

The function u_Z^R satisfies the ultrametric inequality, so it can be represented as a tree (Semple and Steel 2003, §7.2).

Our convergence result for reciprocal clustering requires two additional assumption on the underlying network: (1) the weight function is a dissimilarity measure (i.e. self-weights are 0), and (2) the asymmetry is bounded, i.e. the network has finite reversibility (cf. Sect. 2.8).

Under these assumptions, we are able to show that with increasing sample size, reciprocal clustering converges to an object that is insensitive to the shape of the data distribution. However, the limiting network is slightly different from that of the nonreciprocal clustering case. The proof of the next result is provided in Sect. 4.8.3.

Theorem 4.6.13 (Convergence of reciprocal clustering) *Let (Z, ω_Z, μ_Z) be a measure network with dissimilarity weights and finite reversibility. Suppose $\text{supp}(\mu_Z)$ is a finite, disjoint union of compact, ω -path-connected components $\{Z_a : a \in A\}$. Let (A, λ_A) be as in Definition 4.6.6, and let $(A, u_A^R) = \mathcal{H}^R(A, \lambda_A)$. For each $i \in \mathbb{N}$, let $z_i : \Omega \rightarrow Z$ be an independent random variable defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with distribution μ_Z . For each $n \in \mathbb{N}$, let $Z_n = \{z_1, z_2, \dots, z_n\}$, and for each $\omega \in \Omega$, let $Z_n(\omega)$ denote the subnetwork induced by $\{z_1(\omega), \dots, z_n(\omega)\}$. Let $\varepsilon > 0$. Then,*

$$\mathbb{P}(\{\omega \in \Omega : d_{\mathcal{N}}((A, u_A^R), \mathcal{H}^R(Z_n(\omega))) \geq \varepsilon\}) \leq \frac{(1 - \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z), A))^n}{\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z), A)}.$$

In particular, the output of the reciprocal clustering method applied to the sampled network Z_n converges almost surely to (A, u_A^R) in the sense of $d_{\mathcal{N}}$ as the sample size increases.

In the case of Finsler manifolds with finite reversibility, we can also recover the result of Proposition 4.6.11.

Proposition 4.6.14 (Reciprocal clustering on Finsler manifolds with finite reversibility) *Let (M, F_M, ω_m) be a compact, connected finitely-reversible Finsler manifold without boundary. Here ω_m is the asymmetric weight function induced by the Finsler function F_M . Then $u_M^R(z, z') = 0$ for all $z, z' \in M$.*

Proof of Proposition 4.6.14 Let $z, z' \in M$. Let $\varepsilon > 0$, and let $\gamma : [0, 1] \rightarrow M$ be a curve from z to z' . By invoking the finite reversibility of M , choose n uniformly separated points $\{z_1, \dots, z_n\}$ on $\gamma([0, 1])$ for sufficiently large n such that $\max(\bar{\omega}_M(z_i, z_{i+1}), \bar{\omega}_M(z_{i+1}, z_i)) < \varepsilon$ for each $i = 1, \dots, n-1$. Here $z_1 = z$ and $z_n = z'$. Then $u_M^R(z, z') < \varepsilon$. Since $\varepsilon > 0$ was arbitrary, the result follows. \square

4.7 Clustering and persistence on the directed circle with finite reversibility

Recall the directed circle with finite reversibility $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$ presented in Sect. 2.9. Also consider the *directed circle with finite reversibility on n nodes* $(\mathbb{S}_n^1, \omega_{n, \rho})$ obtained by writing

$$\mathbb{S}_n^1 := \left\{ e^{2\pi i \frac{k}{n}} \in \mathbb{C} : k \in \{0, 1, \dots, n-1\} \right\}$$

and defining $\omega_{n, \rho}$ to be the restriction of $\omega_{\mathbb{S}^1, \rho}$ on this set.

Directed circles on n nodes without the finite reversibility condition were introduced in Chowdhury and Mémoli (2018b), and their Dowker persistent homology was fully characterized by drawing on a connection to the Čech complex of a standard unit circle. The persistent homology of this complex had in turn been fully characterized in Adamaszek et al. (2016), and Adamaszek and Adams (2017). In this section we further elaborate on this connection by showing that the Dowker complexes of directed circles on n nodes with reversibility ρ actually *interpolate* between the Čech complex of the standard circle and the Dowker complex of the directed, irreversible circle. We will utilize the language of interleavings of simplicial complexes, and we describe this construction briefly. For more details we direct the reader to Chowdhury and Mémoli (2018b).

Definition 4.7.1 (*Interleavings of simplicial complexes*) Simplicial maps between simplicial complexes $f, g : U \rightarrow V$ are said to be *contiguous* if $f(\sigma) \cup g(\sigma)$ is a simplex in V whenever σ is a simplex in U (Munkres 1984). Two filtered simplicial complexes $\{U_\delta \xrightarrow{\iota_{\delta, \delta'}} U_{\delta'}\}_{\delta \in \mathbb{R}}, \{V_\delta \xrightarrow{\iota_{\delta, \delta'}} V_{\delta'}\}_{\delta \in \mathbb{R}}$ are said to be ε -interleaved for $\varepsilon \geq 0$ if there exist two families of simplicial maps $\{\phi_\delta : U_\delta \rightarrow V_{\delta+\varepsilon}\}_{\delta \in \mathbb{R}}$ and $\{\psi_\delta : V_\delta \rightarrow U_{\delta+\varepsilon}\}_{\delta \in \mathbb{R}}$ such that the following pairs of maps are contiguous for any $\delta \leq \delta' \in \mathbb{R}$:

- $\phi_{\delta'} \circ \iota_{\delta, \delta'}$ and $\iota_{\delta+\varepsilon, \delta'+\varepsilon} \circ \phi_\delta$
- $\psi_{\delta'} \circ \iota_{\delta, \delta'}$ and $\iota_{\delta+\varepsilon, \delta'+\varepsilon} \circ \psi_\delta$
- $\psi_{\delta+\varepsilon} \circ \phi_\delta$ and $\iota_{\delta, \delta+2\varepsilon}$
- $\phi_{\delta+\varepsilon} \circ \psi_\delta$ and $\iota_{\delta, \delta+2\varepsilon}$.

Importantly, if two filtered simplicial complexes are ε -interleaved, then their corresponding persistent vector spaces are ε -interleaved in the sense of Definition 4.4.2 (cf. Dey et al. 2016 and also Chowdhury and Mémoli 2018b for details on these concepts using a similar setup).

Now we set up some notation. To use the results of Adamaszek et al. (2016), and Adamaszek and Adams (2017), we need the circle parametrized by the unit interval $[0, 1)$, which we denote as S^1 . Let (S^1, d) denote the $[0, 1)$ -circle with geodesic distance. Then define:

$$S_n^1 := \left\{ 0, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n} \right\}.$$

Also define:

$$\begin{aligned}\varphi : S_n^1 &\rightarrow \mathbb{S}_n^1 & \psi : \mathbb{S}_n^1 &\rightarrow S_n^1 \\ \frac{k}{n} &\mapsto \exp(2\pi i \frac{k}{n}) & \exp(2\pi i \frac{k}{n}) &\mapsto \frac{k}{n}\end{aligned}$$

In what follows we replace $\check{\mathbf{C}}_\delta(-)$, $\mathbf{Dow}_\delta(-)$ notation with $\check{\mathbf{C}}(-; \delta)$, $\mathbf{Dow}(-; \delta)$ for readability. For each $\delta \in \mathbb{R}$, define:

$$U_\delta := \check{\mathbf{C}}(S_n^1; \frac{1}{2}(\frac{\delta}{2\pi} + \frac{\delta}{2\pi\rho})), \quad V_\delta := \mathbf{Dow}^{\text{so}}(\mathbb{S}_n^1; \delta).$$

To motivate this definition, first we consider different cases of δ . For $\delta < 0$, both U_δ and V_δ are empty. For $\delta = 0$, U_δ and V_δ comprise n 0-simplices. For $\delta \geq 2\pi$, both U_δ and V_δ contain the full simplex on n points. Let $\delta \in (0, 2\pi)$, and let $\tau \in \mathbf{Dow}^{\text{so}}(\mathbb{S}_n^1; \delta)$. Then τ has the form

$$\tau = \left[\exp\left(2\pi i \left(\frac{j}{n}\right)\right), \exp\left(2\pi i \left(\frac{j}{n} + \frac{1}{n}\right)\right), \dots, \right. \\ \left. \exp\left(2\pi i \left(\frac{j}{n} + \frac{k}{n}\right)\right), \dots, \exp\left(2\pi i \left(\frac{j}{n} + \frac{m}{n}\right)\right) \right],$$

where $\exp(2\pi i(\frac{j}{n} + \frac{k}{n}))$ acts as the source. This means that we have $2\pi \frac{m-k}{n} < \delta$ and $2\pi \frac{k-0}{n} < \frac{\delta}{\rho}$. In particular, the total geodesic distance on the unit circle from $\exp(2\pi i(\frac{j}{n}))$ to $\exp(2\pi i(\frac{j}{n} + \frac{m}{n}))$ is $2\pi \frac{m-0}{n} < \delta + \frac{\delta}{\rho}$. Rescaling to the $[0, 1)$ -circle, we find the total geodesic distance from $\frac{j}{n}$ to $\frac{j+m}{n}$ to be $\frac{\delta}{2\pi} + \frac{\delta}{2\pi\rho}$. Finally consider $\sigma := \psi(\tau) = [\frac{j}{n}, \frac{j}{n} + \frac{1}{n}, \dots, \frac{j}{n} + \frac{\lfloor m/2 \rfloor}{n}, \dots, \frac{j}{n} + \frac{m}{n}]$. If m is even, then each $\frac{l}{n} \in \sigma$ is contained in the closed ball

$$\overline{B}\left(\frac{j}{n} + \frac{m/2}{n}, \frac{\delta}{4\pi} + \frac{\delta}{4\pi\rho}\right).$$

In the odd m case, we may need to extend the radius of the closed ball by $1/n$. This can be done via the update $\delta \leftarrow \delta + 4\pi/n$, as we have $U_{\delta+4\pi/n} = \check{\mathbf{C}}(S_n^1; \frac{\delta}{4\pi} + \frac{1}{n} + \frac{\delta}{4\pi\rho} + \frac{1}{n\rho}) \supseteq \check{\mathbf{C}}(S_n^1; \frac{\delta}{4\pi} + \frac{\delta}{4\pi\rho} + \frac{1}{n})$. Thus for any simplex $\tau \in V_\delta$, we have that $\psi(\tau)$ is a simplex in $U_{\delta+4\pi/n}$.

Conversely, suppose $\sigma = [\frac{j}{n}, \frac{j}{n} + \frac{1}{n}, \dots, \frac{j}{n} + \frac{m}{n}]$ is a simplex in U_δ . Then the geodesic distance from $\frac{j}{n}$ to $\frac{j+m}{n}$ is bounded above by $\frac{\delta}{2\pi} + \frac{\delta}{2\pi\rho}$. Rescaling to the $[0, 2\pi)$ -circle, the geodesic distance from $\exp(2\pi i(\frac{j}{n}))$ to $\exp(2\pi i(\frac{j}{n} + \frac{m}{n}))$ is bounded above by $\delta + \delta/\rho$. Thus there is a point $\xi \in \mathbb{S}^1$ such that the geodesic distance from $\exp(2\pi i(\frac{j}{n}))$ to ξ is δ/ρ and the distance from ξ to $\exp(2\pi i(\frac{j}{n} + \frac{m}{n}))$ is δ . We also know that ξ is in a $\frac{1}{2} \cdot \frac{2\pi}{n}$ -neighborhood of \mathbb{S}_n^1 . Let k be such that $\exp(2\pi i(\frac{j}{n} + \frac{k}{n}))$ is the closest element of \mathbb{S}_n^1 clockwise of ξ . Then we have:

$$\begin{aligned}\omega_{n,\rho}\left(e^{2\pi i \frac{j}{n}} \cdot e^{2\pi i \frac{k}{n}}, e^{2\pi i \frac{j}{n}} \cdot e^{2\pi i \frac{m}{n}}\right) &< \delta + 2\pi/n, \\ \omega_{n,\rho}\left(e^{2\pi i \frac{j}{n}}, e^{2\pi i \frac{j}{n}} \cdot e^{2\pi i \frac{k}{n}}\right) &< \delta/\rho.\end{aligned}$$

In particular, $\tau = \varphi(\sigma)$ is a simplex in $V_{\delta+4\pi/n}$.

It follows from the preceding work that the rescaling maps φ, ψ , which are essentially identity maps on simplices, induce interleavings. We record this below:

Claim 13 Let U_δ and V_δ be the simplicial complexes on n points as defined above. Then they are $4\pi/n$ interleaved via the simplicial maps induced by φ and ψ .

Finally we invoke a characterization of the Čech complex of S^1 .

Theorem 4.7.2 (Adamaszek et al. 2016; Adamaszek and Adams 2017) *Let $n \in \mathbb{N}$, and let $0 \leq k \leq n - 2$ be an integer. Then,*

$$\check{C}(S_n^1, \frac{k}{2n}) = \begin{cases} \bigvee^{n-k-1} S^{2l} & : \frac{k}{n} = \frac{l}{l+1}, \\ S^{2l+1} & : \frac{l}{l+1} < \frac{k}{n} < \frac{l+1}{l+2}, \end{cases}$$

where $l \in \mathbb{Z}_+$. Moreover, if $\delta < \delta' < \frac{1}{2}$ and $\check{C}(S_n^1, \delta), \check{C}(S_n^1, \delta')$ have the same homotopy type, then the induced inclusion map is a homotopy equivalence. Consequently the persistent homology of $\check{C}(S_n^1; -)$ is fully characterized.

The latter part of the preceding theorem is not contained in Adamaszek et al. (2016), but follows from work in Adamaszek and Adams (2017). An explicit verification of this portion, due to Henry Adams, can be found in Chowdhury and Mémoli (2018b, Theorem 49).

Theorem 4.7.3 *Let $1 \leq \rho < \infty$ be a finite reversibility parameter. The directed circle $(S^1, \omega_{S^1, \rho})$ with reversibility ρ has the following Dowker persistence barcodes in each dimension $2l + 1$ for $l \in \mathbb{Z}_+$:*

$$\text{Dgm}_{2l+1}^{\text{Dow}}(S^1, \omega_{S^1, \rho}) = \left\{ \left(\frac{l(4\pi\rho)}{2(l+1)(1+\rho)}, \frac{(l+1)(4\pi\rho)}{2(l+2)(1+\rho)} \right) \right\}.$$

In even dimensions, the persistence diagram is trivial. Moreover, in the fully reversible case $\rho = 1$, rescaling from the $[0, 2\pi)$ -circle to the $[0, 1)$ -circle recovers the persistent homology of the Čech complex of the circle (Adamaszek and Adams 2017, Section 9). An illustration of this result is provided in Fig. 3.

Proof By Theorems 3.1.3 and 3.3.1, we know that it suffices to characterize the Dowker persistence for S_n^1 and then take the limit as $n \rightarrow \infty$. First we observe from Theorem 4.7.2 that $\check{C}(S_n^1; \varepsilon) \simeq S^{2l+1}$ on $(\frac{l}{2(l+1)}, \frac{l+1}{2(l+2)})$. By the work preceding Claim 13, we know that $\check{C}(S_n^1; \varepsilon)$ is related to $\text{Dow}(S_n^1; \delta)$ via the affine transformation $\varepsilon = \frac{\delta}{4\pi} + \frac{\delta}{4\pi\rho}$. Thus letting $\varepsilon = \frac{\delta}{4\pi} + \frac{\delta}{4\pi\rho}$ and solving for δ then yields:

$$\delta \in \left(\frac{l(4\pi\rho)}{2(l+1)(1+\rho)}, \frac{(l+1)(4\pi\rho)}{2(l+2)(1+\rho)} \right).$$

From the interleaving result in Claim 13 and Definition 4.7.1, we know that as $n \rightarrow \infty$, $\text{Dgm}^{\text{Dow}}(S_n^1) \rightarrow \text{Dgm}^{\check{C}}(S_n^1)$. Taking the affine transformation of scale parameters into account then yields the result. \square

We now specialize to the case of hierarchical clustering on $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$.

Theorem 4.7.4 (Nonreciprocal clustering on $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$)

$$u_{\mathbb{S}^1, \rho}^{\text{NR}}(x, x') = 0 \text{ for all } x, x' \in \mathbb{S}^1.$$

Proof of Theorem 4.7.4 We claim that \mathbb{S}^1 is ω -path-connected, with path-connectivity constant $\text{pc}_{\mathbb{S}^1} = 0$ (invoking Lemma 4.6.4). The result then follows by Lemma 4.8.1. Let $x, y \in \mathbb{S}^1$. Without loss of generality, suppose $\omega_{\mathbb{S}^1, \rho}(x, y) = \omega_{\mathbb{S}^1}(x, y)$ (otherwise, switch x and y). Let $\varepsilon > 0$. Pick $x_0 = x, x_1, x_2, \dots, x_n = y$ such that $\omega_{\mathbb{S}^1, \rho}(x_i, x_{i+1}) = \omega_{\mathbb{S}^1}(x_i, x_{i+1}) < \varepsilon$ for each $0 \leq i \leq n-1$. We automatically have $\omega_{\mathbb{S}^1, \rho}(x_i, x_i) = 0$ for all $0 \leq i \leq n$, and hence $\tilde{\omega}_{\mathbb{S}^1, \rho}(x_i, x_{i+1}) \subseteq B(0, \varepsilon)$ for all i . Since $x, y \in \mathbb{S}^1$ were arbitrary, it follows by Definition 4.6.1 that \mathbb{S}^1 is ω -path-connected. The preceding work shows that $\text{pc}_{\mathbb{S}^1} = 0$. The result follows. \square

The case of nonreciprocal clustering essentially followed from an application of Theorem 4.6.10 (i.e. the special case of Lemma 4.8.1. Next we consider the application of reciprocal clustering to $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$ via Theorem 4.6.13.

Theorem 4.7.5 (Reciprocal clustering on $(\mathbb{S}^1, \omega_{\mathbb{S}^1, \rho})$)

$$u_{\mathbb{S}^1, \rho}^{\text{R}}(x, x') = 0 \text{ for all } x, x' \in \mathbb{S}^1.$$

Proof of Theorem 4.7.5 The proof is exactly analogous to that of Theorem 4.7.4, except that we utilize the finite reversibility to obtain reciprocal connections. \square

4.8 Proofs from Section 4

4.8.1 Proofs related to stability of basic invariants

Proof of Lemma 4.3.1 Observe that $f(Z) = \{f_Z(z) : z \in Z\} = f_Z(Z)$, so we need to show

$$d_{\text{H}}^{\mathbb{R}}(f_Z(Z), f_Y(Y)) = \inf_{R \in \mathcal{R}(Z, Y)} \sup_{(z, y) \in R} |f_Z(z) - f_Y(y)|.$$

Recall that by the definition of Hausdorff distance on \mathbb{R} , we have

$$d_{\text{H}}^{\mathbb{R}}(f_Z(Z), f_Y(Y)) = \max \left\{ \sup_{z \in Z} \inf_{y \in Y} |f_Z(z) - f_Y(y)|, \sup_{y \in Y} \inf_{z \in Z} |f_Z(z) - f_Y(y)| \right\}.$$

Let $a \in Z$ and let $R \in \mathcal{R}(Z, Y)$. Then there exists $b \in Y$ such that $(a, b) \in R$. Then we have:

$$|f_Z(a) - f_Y(b)| \leq \sup_{(z, y) \in R} |f_Z(z) - f_Y(y)|, \text{ and so}$$

$$\inf_{b \in Y} |f_Z(a) - f_Y(b)| \leq \sup_{(z,y) \in R} |f_Z(z) - f_Y(y)|.$$

This holds for all $a \in Z$. Then,

$$\sup_{a \in Z} \inf_{b \in Y} |f_Z(a) - f_Y(b)| \leq \sup_{(z,y) \in R} |f_Z(z) - f_Y(y)|.$$

This holds for all $R \in \mathcal{R}(Z, Y)$. So we have

$$\sup_{a \in Z} \inf_{b \in Y} |f_Z(a) - f_Y(b)| \leq \inf_{R \in \mathcal{R}} \sup_{(z,y) \in R} |f_Z(z) - f_Y(y)|.$$

By a similar argument, we also have

$$\sup_{b \in Y} \inf_{a \in Z} |f_Z(a) - f_Y(b)| \leq \inf_{R \in \mathcal{R}} \sup_{(z,y) \in R} |f_Z(z) - f_Y(y)|.$$

$$\text{Thus } d_H^{\mathbb{R}}(f_Z(Z), f_Y(Y)) \leq \inf_{R \in \mathcal{R}} \sup_{(z,y) \in R} |f_Z(z) - f_Y(y)|.$$

Now we show the reverse inequality. Let $z \in Z$, and let $\eta > d_H^{\mathbb{R}}(f_Z(Z), f_Y(Y))$. Then there exists $y \in Y$ such that $|f_Z(z) - f_Y(y)| < \eta$. Define $\varphi(z) = y$, and extend φ to all of Z in this way. Let $y \in Y$. Then there exists $z \in Z$ such that $|f_Z(z) - f_Y(y)| < \eta$. Define $\psi(y) = z$, and extend ψ to all of Y in this way. Let $R = \{(z, \varphi(z)) : z \in Z\} \cup \{(\psi(y), y) : y \in Y\}$. Then for each $(a, b) \in R$, we have $|f_Z(a) - f_Y(b)| < \eta$. Thus we have $\inf_{R \in \mathcal{R}} \sup_{(z,y) \in R} |f_Z(z) - f_Y(y)| < \eta$. Since $\eta > d_H^{\mathbb{R}}(f_Z(Z), f_Y(Y))$ was arbitrary, it follows that

$$\inf_{R \in \mathcal{R}(Z, Y)} \sup_{(z,y) \in R} |f_Z(z) - f_Y(y)| \leq d_H^{\mathbb{R}}(f_Z(Z), f_Y(Y)).$$

□

Proof of Proposition 4.3.2 Let $\eta > d_{\mathcal{N}}(Z, Y)$. We break this proof into three parts.

The diam case. Recall that diam is an \mathbb{R} -valued invariant, so we wish to show $|\text{diam}(Z) - \text{diam}(Y)| \leq 2d_{\mathcal{N}}(Z, Y)$. Let $R \in \mathcal{R}(Z, Y)$ be such that for any $(a, b), (a', b') \in R$, we have $|\omega_Z(a, a') - \omega_Y(b, b')| < 2\eta$.

Let $z, z' \in Z$ such that $|\omega_Z(z, z')| = \text{diam}(Z)$, and let y, y' be such that $(z, y), (z', y') \in R$. Then we have:

$$\begin{aligned} |\omega_Z(z, z') - \omega_Y(y, y')| &< 2\eta \\ |\omega_Z(z, z') - \omega_Y(y, y')| + |\omega_Y(y, y')| &< 2\eta + |\omega_Y(y, y')| \\ |\omega_Z(z, z')| &< \text{diam}(Y) + 2\eta. \end{aligned}$$

Thus $\text{diam}(Z) < \text{diam}(Y) + 2\eta$

Similarly, we get $\text{diam}(Y) < \text{diam}(Z) + 2\eta$. It follows that $|\text{diam}(Z) - \text{diam}(Y)| < 2\eta$. Since $\eta > d_{\mathcal{N}}(Z, Y)$ was arbitrary, it follows that:

$$|\text{diam}(Z) - \text{diam}(Y)| \leq 2d_{\mathcal{N}}(Z, Y).$$

For tightness, consider the networks $Z = N_1(1)$ and $Y = N_1(2)$. By Example 2.2.10, we have that $d_{\mathcal{N}}(Z, Y) = \frac{1}{2}$. On the other hand, $\text{diam}(Z) = 1$ and $\text{diam}(Y) = 2$ so that $|\text{diam}(Z) - \text{diam}(Y)| = 1 = 2d_{\mathcal{N}}(Z, Y)$.

The cases tr, out, and in. First we show $\mathbf{L}(\text{tr}) = 2$. By Lemma 4.3.1, it suffices to show

$$\inf_{R \in \mathcal{R}(Z, Y)} \sup_{(z, y) \in R} |\text{tr}_Z(z) - \text{tr}_Y(y)| < 2\eta.$$

Let $R \in \mathcal{R}(Z, Y)$ be such that for any $(a, b), (a', b') \in R$, we have $|\omega_Z(a, a') - \omega_Y(b, b')| < 2\eta$. Then we obtain $|\omega_Z(a, a) - \omega_Y(b, b)| < 2\eta$. Thus $|\text{tr}_Z(a) - \text{tr}_Y(b)| < 2\eta$. Since $(a, b) \in R$ was arbitrary, it follows that $\sup_{(a, b) \in R} |\text{tr}_Z(a) - \text{tr}_Y(b)| < 2\eta$. It follows that $\inf_{R \in \mathcal{R}} \sup_{(a, b) \in R} |\text{tr}_Z(a) - \text{tr}_Y(b)| < 2\eta$. The result now follows because $\eta > d_{\mathcal{N}}(Z, Y)$ was arbitrary. The proofs for out and in are similar, so we just show the former. By Lemma 4.3.1, it suffices to show

$$\inf_{R \in \mathcal{R}(Z, Y)} \sup_{(z, y) \in R} |\text{out}_Z(z) - \text{out}_Y(y)| < 2\eta.$$

Recall that $\text{out}_Z(z) = \max_{z' \in Z} |\omega_Z(z, z')|$. Let $R \in \mathcal{R}(Z, Y)$ be such that $|\omega_Z(z, z') - \omega_Y(y, y')| < 2\eta$ for any $(z, y), (z', y') \in R$. By triangle inequality, it follows that $|\omega_Z(z, z')| < |\omega_Y(y, y')| + 2\eta$. In particular, for $(z', y') \in R$ such that $|\omega_Z(z, z')| = \text{out}_Z(z)$, we have $\text{out}_Z(z) < |\omega_Y(y, y')| + 2\eta$. Hence $\text{out}_Z(z) < \text{out}_Y(y) + 2\eta$. Similarly, $\text{out}_Y(y) < \text{out}_Z(z) + 2\eta$. Thus we have $|\text{out}_Z(z) - \text{out}_Y(y)| < 2\eta$. This holds for all $(z, y) \in R$, so we have:

$$\sup_{(z, y) \in R} |\text{out}_Z(z) - \text{out}_Y(y)| < 2\eta.$$

Minimizing over all correspondences, we get:

$$\inf_{R \in \mathcal{R}} \sup_{(a, b) \in R} |\text{out}_Z(a) - \text{out}_Y(b)| < 2\eta.$$

The result follows because $\eta > d_{\mathcal{N}}(Z, Y)$ was arbitrary.

Finally, we need to show that our bounds for the Lipschitz constant are tight. Let $Z = N_1(1)$ and let $Y = N_1(2)$. Then $d_{\mathcal{N}}(Z, Y) = \frac{1}{2}$. We also have $d_{\mathbb{H}}^{\mathbb{R}}(\text{tr}(Z), \text{tr}(Y)) = |1 - 2| = 1$, and similarly $d_{\mathbb{H}}^{\mathbb{R}}(\text{out}(Z), \text{out}(Y)) = d_{\mathbb{H}}^{\mathbb{R}}(\text{in}(Z), \text{in}(Y)) = 1$.

The cases \mathbf{m}^{out} and \mathbf{m}^{in} . The two cases are similar, so let's just prove $\mathbf{L}(\mathbf{m}^{\text{out}}) = 2$. Since \mathbf{m}^{out} is an \mathbb{R} -invariant, we wish to show $|\mathbf{m}^{\text{out}}(Z) - \mathbf{m}^{\text{out}}(Y)| < 2\eta$. It suffices to show:

$$|\mathbf{m}^{\text{out}}(Z) - \mathbf{m}^{\text{out}}(Y)| \leq d_{\mathbb{H}}^{\mathbb{R}}(\text{out}(Z), \text{out}(Y)),$$

because we have already shown

$$d_H^{\mathbb{R}}(\text{out}(Z), \text{out}(Y)) = \inf_{R \in \mathcal{R}(Z, Y)} \sup_{(z, y) \in R} |\text{out}_Z(z) - \text{out}_Y(y)| < 2\eta.$$

Here we have used Lemma 4.3.1 for the first equality above.

Let $\varepsilon > d_H^{\mathbb{R}}(\text{out}(Z), \text{out}(Y))$. Then for any $z \in Z$, there exists $y \in Y$ such that:

$$|\text{out}_Z(z) - \text{out}_Y(y)| < \varepsilon.$$

Let $a \in Z$ be such that $m^{\text{out}}(Z) = \text{out}_Z(a)$. Then we have:

$$|\text{out}_Z(a) - \text{out}_Y(y)| < \varepsilon,$$

for some $y \in Y$. In particular, we have:

$$m^{\text{out}}(Y) \leq \text{out}_Y(y) < \varepsilon + \text{out}_Z(a) = \varepsilon + m^{\text{out}}(Z).$$

Similarly, we obtain:

$$m^{\text{out}}(Z) < \varepsilon + m^{\text{out}}(Y).$$

Thus we have $|m^{\text{out}}(Z) - m^{\text{out}}(Y)| < \varepsilon$. Since $\varepsilon > d_H^{\mathbb{R}}(\text{out}(Z), \text{out}(Y))$ was arbitrary, we have:

$$|m^{\text{out}}(Z) - m^{\text{out}}(Y)| \leq d_H^{\mathbb{R}}(\text{out}(Z), \text{out}(Y)).$$

The inequality now follows by Lemma 4.3.1 and our proof in the case of the out map.

For tightness, note that $|m^{\text{out}}(N_1(1)) - m^{\text{out}}(N_1(2))| = |1 - 2| = 1 = 2 \cdot \frac{1}{2} = 2d_{\mathcal{N}}(N_1(1), N_1(2))$. The same example works for the m^{in} case. \square

Proof of Proposition 4.3.4 (First inequality.) Let $Z, Y \in \mathcal{CN}$ and let $\eta > d_{\mathcal{N}}(Z, Y)$. Let $R \in \mathcal{R}(Z, Y)$ be such that $\sup_{(z, y), (z', y') \in R} |\omega_Z(z, z') - \omega_Y(y, y')| < 2\eta$. Let $(z, y) \in R$, and let $\alpha \in \text{spec}_Z(z)$. Then there exists $z' \in Z$ such that $\omega_Z(z, z') = \alpha$. Let $y' \in Y$ be such that $(z', y') \in R$. Let $\beta = \omega_Y(y, y')$. Note $\beta \in \text{spec}_Y(y)$. Also note that $|\alpha - \beta| < 2\eta$. By a symmetric argument, for each $\beta \in \text{spec}_Y(y)$, there exists $\alpha \in \text{spec}_Z(z)$ such that $|\alpha - \beta| < 2\eta$. So $d_H^{\mathbb{R}}(\text{spec}_Z(z), \text{spec}_Y(y)) < 2\eta$. This is true for any $(z, y) \in R$, and so we have $\sup_{(z, y) \in R} d_H^{\mathbb{R}}(\text{spec}_Z(z), \text{spec}_Y(y)) \leq 2\eta$. Then we have:

$$\inf_{R \in \mathcal{R}} \sup_{(z, y) \in R} d_H^{\mathbb{R}}(\text{spec}_Z(z), \text{spec}_Y(y)) \leq 2\eta.$$

Since $\eta > d_{\mathcal{N}}(Z, Y)$ was arbitrary, the first inequality follows.

(Second inequality.) Let $R \in \mathcal{R}(Z, Y)$. Let $\eta(R) = \sup_{(z, y) \in R} d_H^{\mathbb{R}}(\text{spec}_Z(z), \text{spec}_Y(y))$. Let $\alpha \in \text{spec}(Z)$. Then $\alpha \in \text{spec}_Z(z)$ for some $z \in Z$. Let $y \in Y$ such that $(z, y) \in R$.

Then there exists $\beta \in \text{spec}_Y(y)$ such that $|\alpha - \beta| \leq d_H^{\mathbb{R}}(\text{spec}_Z(z), \text{spec}_Y(y))$, and in particular, $|\alpha - \beta| \leq \eta(R)$. In other words, for each $\alpha \in \text{spec}(Z)$, there exists $\beta \in \text{spec}(Y)$ such that $|\alpha - \beta| \leq \eta(R)$. By a symmetric argument, for each $\beta \in \text{spec}(Y)$, there exists $\alpha \in \text{spec}(Z)$ such that $|\alpha - \beta| \leq \eta(R)$. Thus $d_H^{\mathbb{R}}(\text{spec}(Z), \text{spec}(Y)) \leq \eta(R)$. This holds for any $R \in \mathcal{R}$. Thus we have

$$d_H^{\mathbb{R}}(\text{spec}(Z), \text{spec}(Y)) \leq \inf_{R \in \mathcal{R}} \sup_{(z, y) \in R} d_H^{\mathbb{R}}(\text{spec}_Z(z), \text{spec}_Y(y)).$$

This proves the second inequality. \square

Proof of Theorem 4.3.6 Let $n \in \mathbb{N}$. We wish to show $d_n(M_n(Z), M_n(Y)) \leq 2d_{\mathcal{N}}(Z, Y)$. Let $R \in \mathcal{R}(Z, Y)$. Let $(z_i) \in Z^n$, and let $(y_i) \in Y^n$ be such that for each i , we have $(z_i, y_i) \in R$. Then for all $j, k \in \{1, \dots, n\}$, $|\omega_Z(z_i, z_j) - \omega_Y(y_i, y_j)| \leq \text{dis}(R)$.

Thus $\inf_{(y_i) \in Y^n} |\omega_Z(z_i, z_j) - \omega_Y(y_i, y_j)| \leq \text{dis}(R)$. This is true for any $(z_i) \in X^n$. Thus we get:

$$\sup_{(z_i) \in Z^n} \inf_{(y_i) \in Y^n} |\omega_Z(z_i, z_j) - \omega_Y(y_i, y_j)| \leq \text{dis}(R).$$

By a symmetric argument, we get $\sup_{(y_i) \in Y^n} \inf_{(z_i) \in Z^n} |\omega_Z(z_i, z_j) - \omega_Y(y_i, y_j)| \leq \text{dis}(R)$. Thus $d_n(M_n(Z), M_n(Y)) \leq \text{dis}(R)$. This holds for any $R \in \mathcal{R}(Z, Y)$. Thus $d_n(M_n(Z), M_n(Y)) \leq \inf_{R \in \mathcal{R}(Z, Y)} \text{dis}(R) = 2d_{\mathcal{N}}(Z, Y)$.

For tightness, let $Z = N_1(1)$ and let $Y = N_1(2)$. Then $d_{\mathcal{N}}(Z, Y) = \frac{1}{2}$, so we wish to show $d_n(M_n(Z), M_n(Y)) = 1$ for each $n \in \mathbb{N}$. Let $n \in \mathbb{N}$. Let $\mathbb{1}_{n \times n}$ denote the $n \times n$ matrix with 1 in each entry. Then $M_n(Z) = \{\mathbb{1}_{n \times n}\}$ and $M_n(Y) = \{2 \cdot \mathbb{1}_{n \times n}\}$. Thus $d_n(M_n(Z), M_n(Y)) = 1$. Since n was arbitrary, we conclude that equality holds for each $n \in \mathbb{N}$. \square

4.8.2 Proofs related to Theorem 4.6.10

Lemma 4.8.1 (Nonreciprocal clustering on an ω -path-connected network) *Let (Z, ω_Z) be a ω -path-connected network with path-connectivity constant pc_Z (cf. Definition 4.6.1). Then $(Z, u_Z^{\text{NR}}) = \mathcal{H}^{\text{NR}}(Z, \omega_Z)$ is given by writing $u_Z^{\text{NR}}(z, z') = \text{pc}_Z$ for all $z, z' \in Z$.*

Proof of Lemma 4.8.1 Let $z, z' \in Z$, and let $\varepsilon > 0$. By Definition 4.6.1, there exist chains $c \in C_Z(z, z')$ and $c' \in C_Z(z', z)$ such that $\max(\text{cost}_Z(c), \text{cost}_Z(c')) < \text{pc}_Z + \varepsilon$. Thus $u_Z^{\text{NR}}(z, z') < \text{pc}_Z + \varepsilon$. This holds for each $\varepsilon > 0$, and for any $z, z' \in Z$. Thus $u_Z^{\text{NR}} \leq \text{pc}_Z$. We also have $u_Z^{\text{NR}} \geq \text{pc}_Z$ by the definition of chain cost and $\bar{\omega}_Z$. This concludes the proof. \square

Remark 4.8.2 (Distortion and $\bar{\omega}_Z$) Let $(Z, \omega_Z), (Y, \omega_Y) \in \mathcal{N}$ and let $R \in \mathcal{R}(Z, Y)$. Then,

$$\sup_{(z, y), (z', y') \in R} |\bar{\omega}_Z(z, z') - \bar{\omega}_Y(y, y')| \leq \text{dis}(R).$$

To see this, fix $(z, y), (z', y') \in R$. Suppose $u, u' \in \{z, z'\}$ are such that $\bar{\omega}_Z(z, z') = \omega_Z(u, u')$. Let $v, v' \in \{y, y'\}$ be such that $(u, v), (u', v') \in R$. Then we have

$$\bar{\omega}_Z(z, z') = \omega_Z(u, u') \leq \text{dis}(R) + \omega_Y(v, v') \leq \text{dis}(R) + \bar{\omega}_Y(y, y').$$

Now let $v, v' \in \{y, y'\}$ be such that $\bar{\omega}_Y(y, y') = \omega_Y(v, v')$. Let $u, u' \in \{z, z'\}$ be such that $(u, v), (u', v') \in R$. Then,

$$\bar{\omega}_Y(y, y') = \omega_Y(v, v') \leq \text{dis}(R) + \omega_Z(u, u') \leq \text{dis}(R) + \bar{\omega}_Z(z, z').$$

It follows that $|\bar{\omega}_Z(z, z') - \bar{\omega}_Y(y, y')| \leq \text{dis}(R)$.

Lemma 4.8.3 (Nonreciprocal clustering collapses ω -path-connected subsets) *Let (Z, ω_Z) be a network such that Z can be written as a finite, disjoint union of compact, ω -path-connected components $\{Z_a : a \in A\}$. Let (A, v_A) be as in Definition 4.6.6, and let $(A, u_A^{\text{NR}}) = \mathcal{H}^{\text{NR}}(A, v_A)$. Also let $\mathcal{U} = \{U_1, \dots, U_m\}$ be a refined $\varepsilon/2$ -system on Z .*

Suppose that $S \subseteq Z$ is a finite subset equipped with the restriction $\omega_S := \omega_Z|_{S \times S}$ such that S has nonempty intersection with Z_a for each $a \in A$, and with U_i for each $1 \leq i \leq m$. Then,

$$d_{\mathcal{N}}((S, u_S^{\text{NR}}), (A, u_A^{\text{NR}})) < \varepsilon.$$

Proof of Lemma 4.8.3 For each $z \in Z$, let $a(z) \in A$ denote the index such that $z \in Z_{a(z)}$. Here the index is unique because Z is a disjoint union. Then define:

$$R := \{(s, a(s)) : s \in S\}.$$

Then $R \in \mathcal{R}(S, A)$. We wish to show $\text{dis}(R) < 2\varepsilon$, where the distortion is calculated with respect to u_S^{NR} and u_A^{NR} . Let $(s, a(s)), (s', a(s')) \in R$.

Claim 14 We have $u_A^{\text{NR}}(a(s), a(s')) \leq u_S^{\text{NR}}(s, s')$.

Proof Pick chains

$$\begin{aligned} c_1 &:= \{r_0 = s, r_1, r_2, \dots, r_k = s'\} \in C_S(s, s') \quad \text{and} \\ c_2 &:= \{t_0 = s', t_1, t_2, \dots, t_j = s\} \in C_S(s', s) \end{aligned}$$

such that $u_S^{\text{NR}}(s, s') = \max(\text{cost}(c_1), \text{cost}(c_2))$. Then for each $0 \leq i \leq k-1$, we have $\bar{\omega}_S(r_i, r_{i+1}) \leq u_S^{\text{NR}}(s, s')$. Similarly for each $0 \leq i \leq j-1$ we have $\bar{\omega}_S(t_i, t_{i+1}) \leq u_S^{\text{NR}}(s, s')$. Now observe that for each $z, z' \in Z$, we have:

$$v_A(a(z), a(z')) \leq \bar{\omega}_Z(z, z'). \quad (14)$$

Then for each $0 \leq i \leq k-1$, we have:

$$v_A(a(r_i), a(r_{i+1})) \leq \bar{\omega}_S(r_i, r_{i+1}) \leq u_S^{\text{NR}}(s, s').$$

Similarly for each $0 \leq i \leq j - 1$, we have:

$$\nu_A(a(t_i), a(t_{i+1})) \leq \bar{\omega}_S(t_i, t_{i+1}) \leq u_S^{\text{NR}}(s, s').$$

It follows that $u_A^{\text{NR}}(a(s), a(s')) \leq u_S^{\text{NR}}(s, s')$. \square

By Theorem 3.1.3, we obtain a correspondence $R' \in \mathcal{R}(S, Z)$ with $\text{dis}(R') < 2\varepsilon$ such that for each $(s, z) \in R'$, we have $\{s, z\} \subseteq U$ for some $U \in \mathcal{U}$. Here the distortion is measured with respect to ω_Z and ω_S . We will use this correspondence R' as follows: for each $z \in Z$, there exists $s \in S$ such that $(s, z) \in R'$. In other words, there exists $s \in S$ such that $\{s, z\} \subseteq U$ for some $U \in \mathcal{U}$. Since \mathcal{U} is a refined $\varepsilon/2$ -system, we know also that s, z belong to the same connected component Z_a , for some $a \in A$.

For each $z \in Z$, we will write $s(z)$ to denote the element of S obtained by the preceding construction.

Claim 15 Let $s, s' \in S$ be such that $a(s) = a(s')$, i.e. s, s' belong to the same ω -path-connected component of Z . Then $u_S^{\text{NR}}(s, s') < \nu_A(a(s), a(s')) + 2\varepsilon \leq u_A^{\text{NR}}(a(s), a(s')) + 2\varepsilon$.

Proof of Claim 15 Since $Z_{a(s)}$ is ω -path-connected, there exists a unique $r \in \mathbb{R}$ such that $\omega_Z(z, z) = r$ for all $z \in Z_{a(s)}$ by Lemma 4.6.4. By Definition 4.6.6, we have $r \leq \nu_A(a(s), a(s))$. Let $\eta > 0$, and let $z, z' \in Z$ be such that $(s, z), (s', z') \in R'$. Then by the definition of ω -path connectivity, we can take a chain $c = \{z_0 = z, z_1, z_2, \dots, z_n = z'\}$ joining z to z' such that $\text{cost}_Z(c) \leq r + \eta$. We can now convert this to a chain in S by using the correspondence R' . Define:

$$c_S := \{s, s(z_1), s(z_2), \dots, s(z_{n-1}), s'\}.$$

By construction, $(s(z_i), z_i) \in R'$ for each $1 \leq i \leq n - 1$. Furthermore we have $(s, z), (s', z') \in R'$ by our choice of z, z' . Now by using Remark 4.8.2 and the fact that $\text{dis}(R') < 2\varepsilon$, we have $\text{cost}_S(c_S) < r + \eta + 2\varepsilon \leq \nu_A(a(s), a(s)) + \eta + 2\varepsilon$.

By a similar process, we can obtain a chain $c'_S \in C_S(s', s)$ such that $\text{cost}_S(c'_S) < \nu_A(a(s), a(s)) + \eta + 2\varepsilon$. Thus $u_S^{\text{NR}}(s, s') < \nu_A(a(s), a(s)) + \eta + 2\varepsilon$. Since $\eta > 0$ was arbitrary, the result follows. \square

Claim 16 We have $u_S^{\text{NR}}(s, s') < u_A^{\text{NR}}(a(s), a(s')) + 2\varepsilon$.

Proof Let $\vec{c} := \{r_0, \dots, r_k\}$ be a chain in A such that $r_0 = a(s), r_k = a(s')$, and for each $0 \leq i \leq k - 1$, we have $\bar{\omega}_A(r_i, r_{i+1}) \leq u_A^{\text{NR}}(a(s), a(s'))$. Similarly let $\overleftarrow{c} := \{t_0, \dots, t_j\}$ be a chain in A such that $t_0 = a(s'), t_j = a(s)$, and for each $0 \leq i \leq j - 1$, we have $\bar{\omega}_A(t_i, t_{i+1}) \leq u_A^{\text{NR}}(a(s), a(s'))$.

By construction, we have $\nu_A(r_i, r_{i+1}) \leq u_A^{\text{NR}}(a(s), a(s'))$ for all $0 \leq i \leq k - 1$. Similarly we have $\nu_A(t_i, t_{i+1}) \leq u_A^{\text{NR}}(a(s), a(s'))$ for all $0 \leq i \leq j - 1$.

Next observe that by compactness of Z , for each $a, a' \in A$ we can obtain $z(a) \in Z_a, z(a') \in Z_{a'}$ such that $\bar{\omega}_Z(z(a), z(a')) = \nu_A(a, a')$. Applying this construction to consecutive elements in the chains \vec{c} and \overleftarrow{c} , we obtain the following chains in Z :

$$\begin{aligned} \vec{c}_Z &= \{z(r_0), \dots, z(r_k)\}, \text{ joining } z(a(s)) = z(r_0) \text{ to } z(a(s')) = z(r_k), \\ \overleftarrow{c}_Z &= \{z(t_0), \dots, z(t_j)\}, \text{ joining } z(a(s')) = z(t_0) \text{ to } z(a(s)) = z(t_j). \end{aligned}$$

In particular, for each $0 \leq i \leq k-1$, we have $\overline{\omega}_Z(z(r_i), z(r_{i+1})) = v_A(r_i, r_{i+1})$. Similarly for each $0 \leq i \leq j-1$, we have $\overline{\omega}_Z(z(t_i), z(t_{i+1})) = v_A(t_i, t_{i+1})$. Furthermore, we have $z(r_i) \in Z_{r_i}$ for each $0 \leq i \leq k$, and $z(t_i) \in Z_{t_i}$ for each $0 \leq i \leq j$.

Now we can use the correspondence $R' \in \mathcal{R}(S, Z)$ that we had fixed earlier. Recall the use of the notation $s(z) \in S$ for $z \in Z$ from the discussion preceding Claim 15. Now we obtain the following chains in S :

$$\begin{aligned} \overrightarrow{c}_S &= \{s(z(r_0)), \dots, s(z(r_k))\}, \text{ joining } s(z(a(s))) \in Z_{a(s)} \text{ to } s(z(a(s'))) \in Z_{a(s')}, \text{ and} \\ \overleftarrow{c}_S &= \{s(z(t_0)), \dots, s(z(t_j))\}, \text{ joining } s(z(a(s'))) \in Z_{a(s')} \text{ to } s(z(a(s))) \in Z_{a(s)}, \text{ such that} \end{aligned}$$

$$\begin{aligned} \overline{\omega}_S(s(z(r_i)), s(z(r_{i+1}))) &< \overline{\omega}_Z(z(r_i), z(r_{i+1})) + 2\varepsilon \quad \text{for all } 0 \leq i \leq k-1, \quad \text{and} \\ \overline{\omega}_S(s(z(t_i)), s(z(t_{i+1}))) &< \overline{\omega}_Z(z(t_i), z(t_{i+1})) + 2\varepsilon \quad \text{for all } 0 \leq i \leq j-1. \end{aligned}$$

Here we have applied Remark 4.8.2 on consecutive points in the chains to obtain the inequalities.

We know that s and $s(z(r_0)) = s(z(t_j))$ belong to the same ω -path-connected component $Z_{a(s)}$, and similarly s' and $s(z(r_k)) = s(z(t_0))$ belong to the same ω -path-connected component $Z_{a(s')}$. By Claim 15, we have:

$$\begin{aligned} u_S^{\text{NR}}(s, s(z(\rho_0))) &= u_S^{\text{NR}}(s, s(z(\tau_v))) < u_A^{\text{NR}}(a(s), a(s')) + 2\varepsilon, \\ u_S^{\text{NR}}(s', s(z(\rho_u))) &= u_S^{\text{NR}}(s', s(z(\tau_0))) < u_A^{\text{NR}}(a(s), a(s')) + 2\varepsilon. \end{aligned}$$

Finally it follows that:

$$u_S^{\text{NR}}(s, s') < u_A^{\text{NR}}(a(s), a(s')) + 2\varepsilon.$$

□

Thus we have $|u_S^{\text{NR}}(s, s') - u_A^{\text{NR}}(a(s), a(s'))| < 2\varepsilon$. Since $(s, a(s)), (s', a(s')) \in R$ were arbitrary, it now follows that $d_{\mathcal{N}}((S, u_S^{\text{NR}}), (A, u_A^{\text{NR}})) < \varepsilon$. □

Theorem 4.6.10 (Convergence of nonreciprocal clustering) *Let (Z, ω_Z, μ_Z) be a measure network. Suppose $\text{supp}(\mu_Z)$ is a finite, disjoint union of compact, ω -path-connected components $\{Z_a : a \in A\}$. Let (A, v_A) be as in Definition 4.6.6, and let $(A, u_A^{\text{NR}}) = \mathcal{H}^{\text{NR}}(A, v_A)$. For each $i \in \mathbb{N}$, let $z_i : \Omega \rightarrow Z$ be an independent random variable defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with distribution μ_Z . For each $n \in \mathbb{N}$, let $\mathcal{Z}_n = \{z_1, z_2, \dots, z_n\}$, and for each $\omega \in \Omega$, let $\mathcal{Z}_n(\omega)$ denote the subnetwork induced by $\{z_1(\omega), \dots, z_n(\omega)\}$. Let $\varepsilon > 0$. Then,*

$$\mathbb{P}(\{\omega \in \Omega : d_{\mathcal{N}}((A, u_A^{\text{NR}}), \mathcal{H}^{\text{NR}}(\mathcal{Z}_n(\omega))) \geq \varepsilon\}) \leq \frac{(1 - \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)))^n}{\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z))}.$$

In particular, the output of the nonreciprocal clustering method applied to the sampled network \mathcal{Z}_n converges almost surely to (A, u_A^{NR}) in the sense of $d_{\mathcal{N}}$ as the sample size increases.

Proof of Theorem 4.6.10 The proof is analogous to that of Theorem 4.1.4, but we repeat the argument here to facilitate the assessment of details. First observe that $\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)) \in (0, 1]$. Let $r \in (0, \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)))$, and let \mathcal{U}_r be a refined $\varepsilon/2$ -system on $\text{supp}(\mu_Z)$ such that $\mathfrak{m}(\mathcal{U}_r) \in (r, \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z))]$. For convenience, write $m := |\mathcal{U}_r|$, and also write $\mathcal{U}_r = \{U_1, \dots, U_m\}$.

For each $1 \leq i \leq m$, define A_i as in the statement of Lemma 4.1.6. Then by Lemma 4.1.6, the probability that at least one U_i has empty intersection with $\mathcal{Z}_n(\omega)$ is bounded as $\mathbb{P}(\bigcup_{k=1}^m A_k) \leq \frac{1}{\mathfrak{m}(\mathcal{U}_r)}(1 - \mathfrak{m}(\mathcal{U}_r))^n$.

On the other hand, if U_i has nonempty intersection with $\mathcal{Z}_n(\omega)$ for each $1 \leq i \leq m$, then by Lemma 4.8.3, we obtain $d_N((A, u_A^{\text{NR}}), \mathcal{H}^{\text{NR}}(\mathcal{Z}_n(\omega))) < \varepsilon$. Now define:

$$B_n := \{\omega \in \Omega : d_N((A, u_A^{\text{NR}}), \mathcal{H}^{\text{NR}}(\mathcal{Z}_n(\omega))) \geq \varepsilon\}.$$

Then we have:

$$\mathbb{P}(B_n) \leq \mathbb{P}\left(\bigcup_{k=1}^m A_k\right) \leq \frac{1}{\mathfrak{m}(\mathcal{U}_r)}(1 - \mathfrak{m}(\mathcal{U}_r))^n.$$

Since $r \in (0, \mathfrak{M}_{\varepsilon/2}(Z))$ was arbitrary, it follows that:

$$\mathbb{P}(B_n) \leq \frac{1}{\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z))}(1 - \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)))^n.$$

By an application of the Borel–Cantelli lemma as in Theorem 4.1.4, we have $\mathbb{P}(\limsup_{n \rightarrow \infty} B_n) = 0$. The result now follows. \square

4.8.3 Proofs related to Theorem 4.6.13

Lemma 4.8.4 (Reciprocal clustering on a ω -path-connected network) *Let (Z, ω_Z) be an ω -path-connected network with dissimilarity weights and finite reversibility ρ_Z . Then $(Z, u_Z^{\text{R}}) = \mathcal{H}^{\text{R}}(Z, \omega_Z)$ is given by writing $u_Z^{\text{R}}(z, z') = 0$ for all $z, z' \in Z$.*

Proof Let $z, z' \in Z$, and let $\varepsilon > 0$. By Definition 4.6.1 and the assumption that $\omega_Z(z, z) = 0$ for all $z \in Z$, there exists a continuous function $\gamma : [0, 1] \rightarrow Z$ such that $\gamma(0) = z$, $\gamma(1) = z'$, and there exist $t_0 = 0 \leq t_1 \leq t_2 \leq \dots \leq t_n = 1$ such that:

$$\omega_Z(\gamma(t_i), \gamma(t_{i+1})) \in \left[0, \frac{\varepsilon}{\rho_Z}\right) \text{ for each } 0 \leq i \leq n-1.$$

Since ω_Z is a dissimilarity, we have $\omega_Z(\gamma(t_{i+1}), \gamma(t_i)) \geq 0$. By finite reversibility, we also have

$$\omega_Z(\gamma(t_{i+1}), \gamma(t_i)) \leq \rho_Z \cdot \omega_Z(\gamma(t_i), \gamma(t_{i+1})) < \varepsilon.$$

Thus by using the chain $\{\gamma(t_0), \gamma(t_1), \dots, \gamma(t_n)\}$, we have $u_Z^{\text{R}}(z, z') < \varepsilon$. Since $z, z' \in Z$ and $\varepsilon > 0$ were arbitrary, the result now follows. \square

Lemma 4.8.5 (Reciprocal clustering collapses ω -path-connected subsets) *Let (Z, ω_Z) be a network with dissimilarity weights and finite reversibility such that Z is a disjoint collection $\{Z_a : a \in A\}$, where A is a finite indexing set and each Z_a is compact and ω -path-connected. Let (A, λ_A) be as in Definition 4.6.6, and let $(A, u_A^R) = \mathcal{H}^R(A, \lambda_A)$. Also let $\mathcal{U} = \{U_1, \dots, U_m\}$ be a refined $\varepsilon/2$ -system on Z .*

Suppose that $S \subseteq Z$ is a finite subset equipped with the restriction $\omega_S := \omega_Z|_{S \times S}$ such that S has nonempty intersection with U_i for each $1 \leq i \leq m$. Then we have:

$$d_{\mathcal{N}}((S, u_S^R), (A, u_A^R)) < \varepsilon.$$

Proof of Lemma 4.8.5 For each $z \in Z$, let $a(z) \in A$ denote the index such that $z \in Z_{a(z)}$. Then define:

$$R := \{(s, a(s)) : s \in S\}.$$

Then $R \in \mathcal{R}(S, A)$. We wish to show $\text{dis}(R) < 2\varepsilon$, where the distortion is calculated with respect to u_S^R and u_A^R . Let $(s, a(s)), (s', a(s')) \in R$.

Claim 17 We have $u_A^R(a(s), a(s')) \leq u_S^R(s, s')$.

Proof Pick a chain $c := \{r_0 = s, r_1, r_2, \dots, r_k = s'\} \in C_S(s, s')$ such that

$$u_S^R(s, s') = \max_{0 \leq i \leq k-1} (\max(\overline{\omega}_S(r_i, r_{i+1}), \overline{\omega}_S(r_{i+1}, r_i))).$$

Next consider the chain $c_A := \{a(r_0), \dots, a(r_k)\}$. By Definition 4.6.6, we have:

$$\max_{0 \leq i \leq k-1} \lambda_A(a(r_i), a(r_{i+1})) \leq u_S^R(s, s').$$

It follows that $u_A^R(a(s), a(s')) \leq u_S^R(s, s')$. □

By Theorem 3.1.3, we obtain a correspondence $R' \in \mathcal{R}(S, Z)$ with $\text{dis}(R') < 2\varepsilon$ such that for each $(s, z) \in R'$, we have $\{s, z\} \subseteq U$ for some $U \in \mathcal{U}$. Here the distortion is measured with respect to ω_Z and ω_S . We will use this correspondence R' as follows: for each $z \in Z$, there exists $s \in S$ such that $(s, z) \in R'$. In other words, there exists $s \in S$ such that $\{s, z\} \subseteq U$ for some $U \in \mathcal{U}$. Since \mathcal{U} is a refined $\varepsilon/2$ -system, we know also that s, z belong to the same connected component Z_a , for some $a \in A$.

For each $z \in Z$, we will write $s(z)$ to denote the element of S obtained by the preceding construction.

Claim 18 Let $s, s' \in S$ be such that $a(s) = a(s')$, i.e. s, s' belong to the same ω -path-connected component of Z . Then $u_S^R(s, s') < 2\varepsilon$.

Proof of Claim 18 Let $\eta > 0$, and let $z, z' \in Z$ be such that $(s, z), (s', z') \in R'$. By Lemma 4.8.4, we can take a chain $c = \{z_0 = z, z_1, z_2, \dots, z_n = z'\}$ from z to z' such that

$$\max_{0 \leq i \leq n-1} \max(\overline{\omega}_Z(z_i, z_{i+1}), \overline{\omega}_Z(z_{i+1}, z_i)) < \eta.$$

We can now convert this to a chain in S by using the correspondence R' . Define:

$$c_S := \{s, s(z_1), s(z_2), \dots, s(z_{n-1}), s'\}.$$

By construction, $(s(z_i), z_i) \in R'$ for each $1 \leq i \leq n-1$. Furthermore we have $(s, z), (s', z') \in R'$ by our choice of z, z' . Now by using Remark 4.8.2 and the fact that $\text{dis}(R') < 2\varepsilon$, we have $u_S^R(s, s') < \eta + 2\varepsilon$. Since $\eta > 0$ was arbitrary, the result follows. \square

Claim 19 We have $u_S^R(s, s') < u_A^R(a(s), a(s')) + 2\varepsilon$.

Proof Let $c := \{r_0, \dots, r_k\}$ be a chain in A such that $r_0 = a(s)$, $r_k = a(s')$, and

$$\max_{0 \leq i \leq k-1} \lambda_A(r_i, r_{i+1}) \leq u_A^R(a(s), a(s')).$$

Next observe that by compactness of Z , for each $a, a' \in A$ we can obtain $z(a) \in Z_a$, $z(a') \in Z_{a'}$ such that:

$$\max(\overline{\omega}_Z(z(a), z(a')), \overline{\omega}_Z(z(a'), z(a))) = \lambda_A(a, a').$$

Applying this construction to consecutive elements in the chain c , we obtain a chain in Z :

$$c_Z = \{z(r_0), \dots, z(r_k)\}, \text{ joining } z(a(s)) \text{ to } z(a(s')).$$

In particular, for each $0 \leq i \leq k-1$, we have

$$\max(\overline{\omega}_Z(z(r_i), z(r_{i+1})), \overline{\omega}_Z(z(r_{i+1}), z(r_i))) = \lambda_A(r_i, r_{i+1}).$$

Now we can use the correspondence $R' \in \mathcal{R}(S, Z)$ that we had fixed earlier. Recall the use of the notation $s(z) \in S$ for $z \in Z$ from the discussion preceding Claim 18. Now we obtain the following chain in S :

$$c_S = \{s(z(r_0)), \dots, s(z(r_k))\}, \text{ joining } s(z(a(s))) \in Z_{a(s)} \text{ to } s(z(a(s'))) \in Z_{a(s')}, \text{ such that}$$

$$\begin{aligned} \overline{\omega}_S(s(z(r_i)), s(z(r_{i+1}))) &< \overline{\omega}_Z(z(r_i), z(r_{i+1})) + 2\varepsilon \quad \text{for all } 0 \leq i \leq k-1, \quad \text{and} \\ \overline{\omega}_S(s(z(r_{i+1})), s(z(r_i))) &< \overline{\omega}_Z(z(r_{i+1}), z(r_i)) + 2\varepsilon \quad \text{for all } 0 \leq i \leq k-1. \end{aligned}$$

Here we have applied Remark 4.8.2 on consecutive points in the chains to obtain the inequalities.

We know that s and $s(z(r_0))$ belong to the same ω -path-connected component $Z_{a(s)}$, and similarly s' and $s(z(r_k))$ belong to the same ω -path-connected component $Z_{a(s')}$. By Claim 18, we have:

$$\begin{aligned} u_S^R(s, s(z(r_0))) &< 2\varepsilon, \\ u_S^R(s', s(z(r_k))) &< 2\varepsilon. \end{aligned}$$

Finally it follows that:

$$u_S^R(s, s') < u_A^R(a(s), a(s')) + 2\varepsilon.$$

□

Thus we have $|u_S^R(s, s') - u_A^R(a(s), a(s'))| < 2\varepsilon$. Since $(s, a(s)), (s', a(s')) \in R$ were arbitrary, it now follows that $d_{\mathcal{N}}((S, u_S^R), (A, u_A^R)) < \varepsilon$. □

Theorem 4.6.13 (Convergence of reciprocal clustering) *Let (Z, ω_Z, μ_Z) be a measure network with dissimilarity weights and finite reversibility. Suppose $\text{supp}(\mu_Z)$ is a finite, disjoint union of compact, ω -path-connected components $\{Z_a : a \in A\}$. Let (A, λ_A) be as in Definition 4.6.6, and let $(A, u_A^R) = \mathcal{H}^R(A, \lambda_A)$. For each $i \in \mathbb{N}$, let $z_i : \Omega \rightarrow Z$ be an independent random variable defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with distribution μ_Z . For each $n \in \mathbb{N}$, let $\mathcal{Z}_n = \{z_1, z_2, \dots, z_n\}$, and for each $\omega \in \Omega$, let $\mathcal{Z}_n(\omega)$ denote the subnetwork induced by $\{z_1(\omega), \dots, z_n(\omega)\}$. Let $\varepsilon > 0$. Then,*

$$\mathbb{P}(\{\omega \in \Omega : d_{\mathcal{N}}((A, u_A^R), \mathcal{H}^R(\mathcal{Z}_n(\omega))) \geq \varepsilon\}) \leq \frac{(1 - \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z), A))^n}{\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z), A)}.$$

In particular, the output of the reciprocal clustering method applied to the sampled network \mathcal{Z}_n converges almost surely to (A, u_A^R) in the sense of $d_{\mathcal{N}}$ as the sample size increases.

Proof of Theorem 4.6.13 The proof is analogous to that of Theorem 4.6.10, but we repeat the argument here to facilitate the assessment of details. First observe that $\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)) \in (0, 1]$. Let $r \in (0, \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)))$, and let \mathcal{U}_r be a refined $\varepsilon/2$ -system on $\text{supp}(\mu_Z)$ such that $\mathfrak{m}(\mathcal{U}_r) \in (r, \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z))]$. For convenience, write $m := |\mathcal{U}_r|$, and also write $\mathcal{U}_r = \{U_1, \dots, U_m\}$.

For each $1 \leq i \leq m$, define A_i as in the statement of Lemma 4.1.6. Then by Lemma 4.1.6, the probability that at least one U_i has empty intersection with \mathcal{Z}_n is bounded as $\mathbb{P}(\bigcup_{k=1}^m A_k) \leq \frac{1}{\mathfrak{m}(\mathcal{U}_r)} (1 - \mathfrak{m}(\mathcal{U}_r))^n$.

On the other hand, if U_i has nonempty intersection with $\mathcal{Z}_n(\omega)$ for each $1 \leq i \leq m$, then by Lemma 4.8.5, we obtain $d_{\mathcal{N}}((A, u_A^R), \mathcal{H}^R(\mathcal{Z}_n(\omega))) < \varepsilon$. Now define:

$$B_n := \{\omega \in \Omega : d_{\mathcal{N}}((A, u_A^R), \mathcal{H}^R(\mathcal{Z}_n(\omega))) \geq \varepsilon\}.$$

Then we have:

$$\mathbb{P}(B_n) \leq \mathbb{P}\left(\bigcup_{k=1}^m A_k\right) \leq \frac{1}{\mathfrak{m}(\mathcal{U}_r)} (1 - \mathfrak{m}(\mathcal{U}_r))^n.$$

Since $r \in (0, \mathfrak{M}_{\varepsilon/2}(X))$ was arbitrary, it follows that:

$$\mathbb{P}(B_n) \leq \frac{1}{\mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z))} (1 - \mathfrak{M}_{\varepsilon/2}(\text{supp}(\mu_Z)))^n.$$

By an application of the Borel–Cantelli lemma as in Theorem 4.1.4, we have $\mathbb{P}(\limsup_{n \rightarrow \infty} B_n) = 0$. The result now follows. \square

5 Discussion

The original motivation for this work came from a scientific problem: that of having a correspondence framework for graph-theoretic structures arising in systems neuroscience (Sporns 2012) in which undirected (*functional connectivity*) and directed (*effective connectivity*) graphs would be compatible with one another. The $d_{\mathcal{N}}$ metric that had recently been introduced in Carlsson et al. (2013, 2014) for extending hierarchical clustering to asymmetric networks was a natural proposal for addressing this limitation. After developing the $d_{\mathcal{N}}$ theory in the finite case (Chowdhury and Mémoli 2015, 2016a) and observing concurrent uses for developing stable persistent homology methods for finite networks (Chowdhury and Mémoli 2018b; Turner 2019), it was apparent that one needed to understand the convergence of sequences of finite networks to continuous objects, as had been done for the cut metric in Lovász (2012) and related earlier works Borgs et al. (2008, 2012). In particular, this question was made explicit in Turner (2019). Studying this convergence led us to realize that even though $d_{\mathcal{N}}$ metrizes the collection $\{(X, \omega_X)\}$ where X is a set and $\omega_X : X \times X \rightarrow \mathbb{R}$ is any function, such a collection was not amenable to meaningful convergence results or even a satisfactory characterization of isomorphism. This first set of obstacles led to the subsequent introduction of topologies and the intermediate collection \mathcal{CN} . This turned out to be a suitable ambient space which was sufficiently general while also having the required topological regularity.

Although our study of the consequences of relaxing metric assumptions was driven by empirical and applied considerations which stem from the data-driven scientific outlook of the recent decades, ultimately we were led to the literature developed by Fréchet and his contemporaries at the beginning of the 20th century (Fréchet 1906; Pitcher et al. 1918; Niemytzki 1927). These authors had been motivated to understand different constructions so they could settle on an appropriate notion of a metric space. Later works by Busemann (1950) and others on quasimetric spaces also turned out to be useful for understanding the subtleties that arise with asymmetry. Along these lines, a possible latent connection that we have not yet explored is to the field of directed topology (Fajstrup et al. 2016).

In this work, we provided a comprehensive theoretical study of $d_{\mathcal{N}}$, its isomorphism structure, and the metric structure that it induces on \mathcal{N} . In particular, we introduced \mathcal{CN} as a subfamily that is nested between \mathcal{FN} and \mathcal{N} , proved the sampling and convergence results on \mathcal{CN} that finally enabled one to obtain well-defined persistence diagrams on *infinite* asymmetric networks, and characterized the Dowker persistence diagrams of the directed circles with finite reversibility. These results now suggest numerous natural lines of continuation, some of which are as follows:

- *Developing new embeddings* As described in the introduction, graph structured data may be embedded into \mathcal{CN} in a multitude of ways. By virtue of this work, different embeddings $G \mapsto \mathcal{CN}$ are now seen to share a common ambient metric

space. Development of new and informative embeddings, perhaps by aggregating existing embeddings, is a broad-scale challenge. Some previous works Chowdhury and Mémoli (2018b); Chowdhury et al. (2017) have shown that building asymmetric Markov transition matrices can give better accuracy on certain classification tasks than their symmetric counterparts, so the systematic development of “asymmetric embeddings” seems to hold promise in extracting additional signal from data. With increasing availability of benchmark data sets for machine learning purposes, using data-driven supervision to generate “learned” embeddings is a natural next step (Litman and Bronstein 2013).

- *New inputs to persistent homology methods* Along these lines, it would be interesting to see which embeddings could serve as good preprocessing steps for the hierarchical clustering and persistent homology methods described earlier in this work. These would automatically benefit from $d_{\mathcal{N}}$ -stability with respect to the bottleneck distance. From a different perspective, methods such as the Dowker complex have natural connections to *hypergraphs* that are used to model multi-way relations (Zhou et al. 2006). Elaborating on this latent connection by extending our constructions to accept hypergraph data as input would open new directions in hypergraph data analysis.
- *More directed models with geometric regularity* Another interesting thread for future development would be to construct directed models of shapes beyond the circles with finite reversibility. This would likely provide an interesting connection to the theoretical development for Randers manifolds in Bao et al. (2004) and help towards characterizing how such spaces are embedded in the $(\mathcal{N}, d_{\mathcal{N}})$ hierarchy. Implicit in this construction should be some notion of geometric regularity such as the finite reversibility assumption.

Finally we remark that embeddings of the form $G \mapsto \mathcal{N}$ can be further extended via maps $\mathcal{N} \rightarrow \mathcal{N}^m$ where the latter is the collection of measure networks equipped with a generalized Gromov–Wasserstein distance (Chowdhury and Mémoli 2019). This approach follows the formulation of Mémoli (2007) in taking a relaxation of the $d_{\mathcal{N}}$ problem, and benefits from the deep study of Riemannian structure developed by Sturm (2012). This particular context has seen extensive growth in the machine learning community in recent years (Peyré et al. 2016; Solomon et al. 2016; Chowdhury and Needham 2020b), and as has already been shown in Xu et al. (2019a), Chowdhury and Needham (2020a), in various cases the results demonstrated in these works arise from utilizing different embeddings $G \mapsto \mathcal{N}$ and studying the particular structure that each embedding imposes upon the problem.

Acknowledgements This work is partially supported by the National Science Foundation under grants DMS-1723003, CCF-1740761, CCF-1526513, IIS-1422400, and DMS-1547357. Facundo Mémoli acknowledges support from the Mathematical Biosciences Institute at The Ohio State University.

Declarations

Conflict of interest The authors declare no conflict of interest.

A Experiments

A.1 Computational aspects and an algorithm for computing minimum matchings

In this section we first discuss some algorithmic details on how to compute the lower bounds for d_N involving local spectra, and then present computational examples. All networks in this section are assumed to be finite. Our software and datasets are available on <https://github.com/fmemoli/PersNet> as part of the PersNet software package.

Lower bounds for d_N involving the comparison of local spectra of two networks such as those in Proposition 4.3.4 require computing the minimum of a functional $J(R) := \max_{(z,y) \in R} C(z, y)$ where $C : Z \times Y \rightarrow \mathbb{R}_+$ is a given *cost* function and R ranges in $\mathcal{R}(Z, Y)$. This is an instance of a bottleneck linear assignment problem (or LBAP) (Burkard et al. 2009). We remark that the current instance differs from the standard formulation in that one is now optimizing over correspondences and not over permutations. Hence the standard algorithms need to be modified.

Assume $n = |Z|$ and $m = |Y|$. In this section we adopt matrix notation and regard R as a matrix $((r_{i,j})) \in \{0, 1\}^{n \times m}$. The condition $R \in \mathcal{R}(Z, Y)$ then requires that $\sum_i r_{i,j} \geq 1$ for all j and $\sum_j r_{i,j} \geq 1$ for all i . We denote by $C = ((c_{i,j})) \in \mathbb{R}_+^{n \times m}$ the matrix representation of the cost function C described above. With the goal of identifying a suitable algorithm, the key observation is that the optimal value $\min_{R \in \mathcal{R}} J(R)$ must coincide with a value realized in the matrix C .

An algorithm with complexity $O(n^2 \times m^2)$ is the one in Algorithm 1 (we give it in Matlab pseudo-code). The algorithm belongs to the family of *thresholding algorithms* for solving matching problems over permutations, see Burkard et al. (2009). Notice that R is a binary matrix and that procedure **TestCorrespondence** has complexity $O(n \times m)$. In the worst case, the matrix C has $n \times m$ distinct entries, and the **while** loop will need to exhaustively test them all, hence the claimed complexity of $O(n^2 \times m^2)$. Even though a more efficient version (with complexity $O((n \times m) \log(n \times m))$) can be obtained by using a bisection strategy on the range of possible values contained in the matrix C (in a manner similar to what is described for the case of permutations in Burkard et al. 2009), here for clarity we limit our presentation to the version detailed above.

Algorithm 1 MinMax matching

```

1: procedure MINMAXMATCH(C)
2:    $v = \text{sort}(\text{unique}(C(:)));$ 
3:    $k = 1;$ 
4:   while  $\sim \text{done}$  do
5:      $c = v(k);$ 
6:      $R = (C \leq c);$ 
7:      $\text{done} = \text{TESTCORRESPONDENCE}(R);$ 
8:      $k = k + 1;$ 
9:   return  $c$ 
10: procedure TESTCORRESPONDENCE(R)
11:    $\text{done} = \text{prod}(\text{sum}(R)) * \text{prod}(\text{sum}(R')) > 0;$ 
12:   return  $\text{done}$ 

```

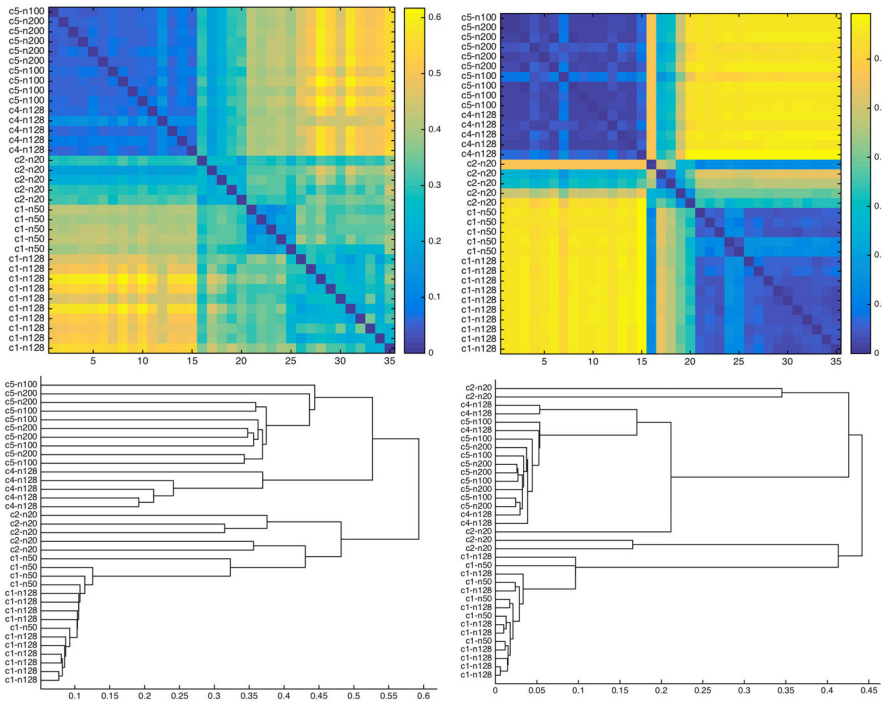


Fig. 17 Performance of local (*left*) and global (*right*) spectrum lower bounds on the random community graph classification task. The global spectrum can reasonably tell apart classes of graphs with 4–5 communities from those with a single community, but is confounded by graphs with two communities. In comparison, the local spectrum is able to better distinguish the two-community graphs

A.2 Computational example: randomly generated networks

As a first application of our ideas we generated a database of weighted directed networks with different numbers of “communities” and different total cardinalities using the software provided by [For]. Using this software, we generated 35 random networks as follows: 5 networks with 5 communities and 200 nodes each (class $c5-n200$), 5 networks with 5 communities and 100 nodes each (class $c5-n100$), 5 networks with 4 communities and 128 nodes each (class $c4-n128$), 5 networks with 2 communities and 20 nodes each (class $c2-n20$), 5 networks with 1 community and 50 nodes each (class $c1-n50$), and 10 networks with 1 community and 128 nodes each (class $c1-n128$). In order to make the comparison more realistic, as a preprocessing step we divided all the weights in each network by the diameter of the network. In this manner, discriminating between networks requires differentiating their structure and not just the scale of the weights. Note that the (random) weights produced by the software [For] are all non-negative.

Using a Matlab implementation of Algorithm 1 we computed a 35×35 matrix of values corresponding to a lower bound based simultaneously on both the in and out local spectra. This strengthening of Proposition 4.3.4 is stated below.

Proposition A.2.1 For all $X, Y \in \mathcal{FN}$,

$$d_{\mathcal{N}}(X, Y) \geq \frac{1}{2} \min_{R \in \mathcal{R}} \max_{(x, y) \in R} C(x, y), \quad \text{where}$$

$$C(x, y) = \max \left(d_{\mathcal{H}}^{\mathbb{R}}(\text{spec}_X^{\text{in}}(x), \text{spec}_Y^{\text{in}}(y)), d_{\mathcal{H}}^{\mathbb{R}}(\text{spec}_X^{\text{out}}(x), \text{spec}_Y^{\text{out}}(y)) \right).$$

This bound follows from Proposition 4.3.4 by the discussion at the beginning of Sect. 4.3.

The results are shown in the form of the lower bound matrix and its single linkage dendrogram in Fig. 17. Notice that the labels in the dendrogram permit ascertaining the quality of the classification provided by the local spectra bound. With only very few exceptions, networks with similar structure (same number of communities) were clustered together regardless of their cardinality. Notice furthermore how networks with 4 and 5 communities merge together before merging with networks with 1 and 2 communities, and vice versa. For comparison, we also provide details about the performance of the global spectra lower bound on the same database in Fig. 17. The results are clearly inferior to those produced by the local version, as predicted by the inequality in Proposition 4.3.4.

A.3 Computational example: simulated hippocampal networks

A natural observation about humans is that as they navigate an environment, they produce “mental maps” which enable them to recall the features of the environment at a later time. This is also true for other animals with higher cognitive function. In the neuroscience literature, it is accepted that the hippocampus in an animal’s brain is responsible for producing a mental map of its physical environment (Best et al. 2001; Brown et al. 1998). More specifically, it has been shown that as a rat explores a given environment, specific physical regions (“place fields”) become linked to specific “place cells” in the hippocampus (O’Keefe and Dostrovsky 1971; O’Keefe and Nadel 1978). Each place cell shows a spike in neuronal activity when the rat enters its place field, accompanied by a drop in activity as the rat goes elsewhere. In order to understand how the brain processes this data, a natural question to ask is the following: Is the time series data of the place cell activity, often referred to as “spike trains”, enough to recover the structure of the environment?

Approaches based on homology (Curto and Itskov 2008) and persistent homology (Dabaghian et al. 2012) have shown that the preceding question admits a positive answer. We were interested in determining if, instead of computing homology groups, we could represent the time series data as networks, and then apply our invariants to distinguish between different environments. Our preliminary results on simulated hippocampal data indicate that such may be the case.

In our experiment, there were two environments: (1) a square of side length L , and (2) a square of side length L , with a disk of radius $0.33L$ removed from the center. In what follows, we refer to the environments of the second type as *1-hole environments*, and those of the first type as *0-hole environments*. For each environment, a random-walk trajectory of 5000 steps was generated, where the agent could move above, below,

left, or right with equal probability. If one or more of these moves took the agent outside the environment (a disallowed move), then the probabilities were redistributed uniformly among the allowed moves. The length of each step in the trajectory was $0.1L$.

In the first set of 20 trials for each environment, 200 place fields of radius $0.1L$ were scattered uniformly at random. In the next two sets, the place field radii were changed to $0.2L$ and $0.05L$. This produced a total of 60 trials for each environment. For each trial, the corresponding network (X, ω_X) was constructed as follows: X consisted of 200 place cells, and for each $1 \leq i, j \leq 200$, the weight $\omega_X(x_i, x_j)$ was given by:

$$\omega_X(x_i, x_j) = 1 - \frac{\# \text{ times cell } x_j \text{ spiked in a window of five time units after cell } x_i \text{ spiked}}{\# \text{ times cell } x_j \text{ spiked}}.$$

The results of applying the local spectra lower bound are shown in the inset of Fig. 18. The labels *env-0*, *env-1* correspond to 0 and 1-hole environments, respectively. Note that with some exceptions, networks corresponding to the same environment are clustered together, regardless of place field radius. In Fig. 18 we also present the single linkage dendrogram obtained from comparing all 120 networks together. In light of these results, we are interested in seeing how these methods can be applied to other time series data arising from biology.

As a final remark, we note that it is possible to obtain better clustering on the hippocampal network dataset by using d_N -invariants that arise from *persistent homology*. We refer the reader to Chowdhury and Mémoli (2018b) for details.

References

- Adamaszek, M., Adams, H.: The Vietoris–Rips complexes of a circle. *Pac. J. Math.* **290**(1), 1–40 (2017)
- Adamaszek, M., Adams, H., Frick, F., Peterson, C., Previte-Johnson, C.: Nerve complexes of circular arcs. *Discret. Comput. Geom.* **56**(2), 251–273 (2016)
- Bao, D., Robles, C., Shen, Z.: Zermelo navigation on Riemannian manifolds. *Journal of Differential Geometry* **66**(3), 377–435 (2004)
- Bao, D., Chern, S.-S., Shen, Z.: An introduction to Riemann–Finsler geometry, vol. 200. Springer, Berlin (2012)
- Bauer, U., Landi, C., Mémoli, F.: The Reeb graph edit distance is universal. In: Cabello, S., Chen, D.Z. (eds.) 36th International Symposium on Computational Geometry, SoCG 2020, June 23–26, 2020, Zürich, Switzerland, vol. 164 of LIPIcs, pp. 15:1–15:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
- Best, P.J., White, A.M., Minai, A.: Spatial processing in the brain: the activity of hippocampal place cells. *Annu. Rev. Neurosci.* **24**(1), 459–486 (2001)
- Blumberg, A.J., Lesnick, M.: Universality of the homotopy interleaving distance. *arXiv preprint arXiv:1705.01690* (2017)
- Borgs, C., Chayes, J.T., Lovász, L., Sós, V.T., Vesztegombi, K.: Convergent sequences of dense graphs I: subgraph frequencies, metric properties and testing. *Adv. Math.* **219**(6), 1801–1851 (2008)
- Borgs, C., Chayes, J.T., Lovász, L., Sós, V.T., Vesztegombi, K.: Convergent sequences of dense graphs II. Multiway cuts and statistical physics. *Ann. Math.* 151–219 (2012)
- Boutin, M., Kemper, G.: Lossless representation of graphs using distributions. *arXiv preprint arXiv:0710.1870* (2007)
- Bridson, M.R., Haefliger, A.: Metric Spaces of Non-positive Curvature, vol. 319. Springer, Berlin (2011)
- Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C., Wilson, M.A.: A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J. Neurosci.* **18**(18), 7411–7425 (1998)
- Burago, D., Burago, Y., Ivanov, S.: A Course in Metric Geometry. AMS Graduate Studies in Math, vol. 33. American Mathematical Society (2001)
- Burkard, R.E., Dell’Amico, M., Martello, S.: Assignment Problems. SIAM (2009)
- Busemann, H.: The geometry of Finsler spaces. *Bull. Am. Math. Soc.* **56**(1), 5–16 (1950)

- Calissano, A., Feragen, A., Vantini, S.: Graph space geometry and geodesic principal components. *MOX Report, Populations of unlabeled networks* (2020)
- Carlsson, G.: Topology and data. *Bull. Am. Math. Soc.* **46**(2), 255–308 (2009)
- Carlsson, G., Mémoli, F.: Characterization, stability and convergence of hierarchical clustering methods. *J. Mach. Learn. Res.* **11**, 1425–1470 (2010)
- Carlsson, G., Mémoli, F.: Classifying clustering schemes. *Found. Comput. Math.* **13**(2), 221–252 (2013)
- Carlsson, G., Zomorodian, A., Collins, A., Guibas, L.J.: Persistence barcodes for shapes. *Int. J. Shape Model.* **11**(02), 149–187 (2005)
- Carlsson, G., Mémoli, F., Ribeiro, A., Segarra, S.: Axiomatic construction of hierarchical clustering in asymmetric networks. In: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 5219–5223. IEEE (2013)
- Carlsson, G., Mémoli, F., Ribeiro, A., Segarra, S.: Hierarchical quasi-clustering methods for asymmetric networks. In: Proceedings of the 31th International Conference on Machine Learning, ICML 2014 (2014)
- Carlsson, G., Mémoli, F., Ribeiro, A., Segarra, S.: Hierarchical clustering of asymmetric networks. *Adv. Data Anal. Classif.* **12**(1), 65–105 (2018)
- Carlsson, G., Mémoli, F., Segarra, S.: Robust hierarchical clustering for directed networks: an axiomatic approach. *SIAM J. Appl. Algebra Geom.* **5**(4), 675–700 (2021)
- Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L.J., Oudot, S.Y.: Proximity of persistence modules and their diagrams. In: Proceedings of the Twenty-Fifth Annual Symposium on Computational Geometry, pp. 237–246. ACM (2009a)
- Chazal, F., Cohen-Steiner, D., Guibas, L.J., Mémoli, F., Oudot, S.Y.: Gromov–Hausdorff stable signatures for shapes using persistence. In: Computer Graphics Forum, vol. 28, pp. 1393–1403. Wiley Online Library (2009b)
- Chazal, F., De Silva, V., Oudot, S.: Persistence stability for geometric complexes. *Geom. Dedicata.* **173**(1), 193–214 (2014)
- Chazal, F., De Silva, V., Glisse, M., Oudot, S.: The Structure and Stability of Persistence Modules. Springer, Berlin (2016)
- Chowdhury, S., Mémoli, F.: Metric structures on networks and applications. In: 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 1470–1472 (2015)
- Chowdhury, S., Mémoli, F.: Distances between directed networks and applications. In: 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6420–6424. IEEE (2016a)
- Chowdhury, S., Mémoli, F.: Persistent homology of directed networks. In: 2016 50th Asilomar Conference on Signals, Systems and Computers, pp. 77–81. IEEE (2016b)
- Chowdhury, S., Mémoli, F.: Distances and isomorphism between networks and the stability of network invariants. *arXiv preprint [arXiv:1708.04727](https://arxiv.org/abs/1708.04727)* (2017)
- Chowdhury, S., Mémoli, F.: Explicit geodesics in Gromov–Hausdorff space. *Electron. Res. Announc.* **25**, 48–59 (2018a)
- Chowdhury, S., Mémoli, F.: A functorial Dowker theorem and persistent homology of asymmetric networks. *J. Appl. Comput. Topol.* **2**(1–2), 115–175 (2018b)
- Chowdhury, S., Mémoli, F.: Persistent path homology of directed networks. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1152–1169. SIAM (2018c)
- Chowdhury, S., Mémoli, F.: The Gromov–Wasserstein distance between networks and stable network invariants. *Inf. Inference J. IMA* **8**(4), 757–787 (2019)
- Chowdhury, S., Needham, T.: Generalized spectral clustering via Gromov–Wasserstein learning. *arXiv preprint [arXiv:2006.04163](https://arxiv.org/abs/2006.04163)* (2020a)
- Chowdhury, S., Needham, T.: Gromov–Wasserstein averaging in a Riemannian framework. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 842–843 (2020b)
- Chowdhury, S., Dai, B., Mémoli, F.: Topology of stimulus space via directed network persistent homology. *Cosyne Abstracts* (2017)
- Chowdhury, S., Gebhart, T., Huntsman, S., Yutin, M.: Path homologies of deep feedforward networks. In: 2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA), pp. 1077–1082. IEEE (2019)
- Chowdhury, S., Clause, N., Mémoli, F., Sánchez, J., Wellner, Z.: New families of stable simplicial filtration functors. *Topol. Appl.* (2020)

- Chung, F.: Laplacians and the Cheeger inequality for directed graphs. *Ann. Comb.* **9**(1), 1–19 (2005)
- Cobzas, S.: *Functional Analysis in Asymmetric Normed Spaces*. Springer, Berlin (2012)
- Curto, C., Itskov, V.: Cell groups reveal structure of stimulus space. *PLoS Comput. Biol.* **4**(10) (2008)
- Dabaghian, Y., Mémoli, F., Frank, L., Carlsson, G.: A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS Comput. Biol.* **8**(8) (2012)
- Dey, T.K., Mémoli, F., Wang, Y.: Multiscale mapper: topological summarization via codomain covers. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 997–1013. Society for Industrial and Applied Mathematics (2016)
- Dey, T.K., Li, T., Wang, Y.: An efficient algorithm for 1-dimensional (persistent) path homology. In: *36th International Symposium on Computational Geometry (SoCG 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2020)
- Dowker, C.H.: Homology groups of relations. *Ann. Math.* 84–95 (1952)
- Dudley, R.M.: *Real Analysis and Probability*, vol. 74. Cambridge University Press, Cambridge (2002)
- Edelsbrunner, H., Harer, J.: *Computational Topology: An Introduction*. American Mathematical Society (2010)
- Edelsbrunner, H., Morozov, D.: *Persistent Homology: Theory and Practice* (2014)
- Edgar, G.A.: *Classics on Fractals*. (1993)
- Fagin, R., Stockmeyer, L.: Relaxing the triangle inequality in pattern matching. *Int. J. Comput. Vis.* **30**(3), 219–231 (1998)
- Fajstrup, L., Goubault, E., Haucourt, E., Mimram, S., Raussen, M.: *Directed Algebraic Topology and Concurrency*. Springer, Berlin (2016)
- Fortunato, S.: Benchmark graphs to test community detection algorithms. <https://sites.google.com/site/santofortunato/inthepress2> (2017)
- Franklin, S.P.: Spaces in which sequences suffice. *Fundam. Math.* **57**(1), 107–115 (1965)
- Fréchet, M.: Sur quelques points du calcul fonctionnel. *Rendiconti del Circolo Matematico di Palermo* (1884–1940) **22**(1), 1–72 (1906)
- Friston, K.J.: Functional and effective connectivity in neuroimaging: a synthesis. *Hum. Brain Mapp.* **2**(1–2), 56–78 (1994)
- Frosini, P.: Measuring shapes by size functions. In: *Intelligent Robots and Computer Vision X: Algorithms and Techniques*, pp. 122–133. International Society for Optics and Photonics (1992)
- Galvin, F., Shore, S.: Completeness in semimetric spaces. *Pac. J. Math.* **113**(1), 67–75 (1984)
- Galvin, F., Shore, S.: Distance functions and topologies. *Am. Math. Mon.* **98**(7), 620–623 (1991)
- Grigor'yan, A., Lin, Y., Muranov, Y., Yau, S.-T.: Homotopy theory for digraphs. *Pure Appl. Math. Q.* **10**(4), 619–674 (2014)
- Gromov, M.: Structures métriques pour les variétés Riemanniennes. *Textes Mathématiques [Mathematical Texts]* **1** (1981)
- Gromov, M.: *Metric Structures for Riemannian and Non-Riemannian Spaces*. Progress in Mathematics, vol. 152. Birkhäuser Boston Inc., Boston (1999)
- Gruenhage, G.: Generalized metric spaces. *Handbook of Set-Theoretic Topology*, pp. 423–501 (1984)
- Hartigan, J.A.: *Clustering Algorithms*, vol. 209. Wiley, New York (1975)
- Hartigan, J.A.: Consistency of single linkage for high-density clusters. *J. Am. Stat. Assoc.* **76**(374), 388–394 (1981)
- Hartigan, J.A.: Statistical theory in clustering. *J. Classif.* **2**(1), 63–76 (1985)
- Heinonen, J.: *Lectures on Analysis on Metric Spaces*. Springer, Berlin (2012)
- Hendrikson, R.: Using Gromov–Wasserstein distance to explore sets of networks. Master's thesis, University of Tartu (2016)
- Hill, T.L.: *Free Energy Transduction and Biochemical Cycle Kinetics*. Courier Corporation (2005)
- Hongtengl, X.: Gromov–Wasserstein factorization models for graph clustering. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 6478–6485 (2020)
- Hörster, E., Lienhart, R., Slaney, M.: Image retrieval on large-scale image databases. In: *Proceedings of the 6th ACM International Conference on Image and Video Retrieval*, pp. 17–24 (2007)
- Huntsman, S.: Generalizing cyclomatic complexity via path homology. *arXiv preprint arXiv:2003.00944* (2020)
- Ivanov, A.O., Nikolaeva, N.K., Tuzhilin, A.A.: The Gromov–Hausdorff metric on the space of compact metric spaces is strictly intrinsic. *Math. Notes* **5**(100), 883–885 (2016)
- Jain, B.J., Obermayer, K.: Structure spaces. *J. Mach. Learn. Res.* **10**(11) (2009)

- Jardine, N., Sibson, R.: *Mathematical Taxonomy*. Wiley Series in Probability and Mathematical Statistics, Wiley, New York (1971)
- Jegou, H., Schmid, C., Harzallah, H., Verbeek, J.: Accurate image search using the contextual dissimilarity measure. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(1), 2–11 (2008)
- Kalton, N.J., Ostrovskii, M.I.: Distances between Banach spaces. In: *Forum Mathematicum*, vol. 11, pp. 17–48. Walter de Gruyter (1999)
- Kelly, J.C.: Bitopological spaces. *Proc. Lond. Math. Soc.* **3**(1), 71–89 (1963)
- Le Gouic, T., Loubes, J.-M.: Existence and consistency of Wasserstein barycenters. *Probab. Theory Relat. Fields* **168**(3–4), 901–917 (2017)
- Leustean, L., Nicolae, A., Zaharescu, A.: Barycenters in uniformly convex geodesic spaces. *arXiv preprint arXiv:1609.02589* (2016)
- Litman, R., Bronstein, A.M.: Learning spectral descriptors for deformable shape correspondence. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(1), 171–180 (2013)
- Lovász, L.: *Large Networks and Graph Limits*, vol. 60. American Mathematical Society (2012)
- Lyu, H., Memoli, F., Sivakoff, D.: Sampling random graph homomorphisms and applications to network data analysis. *arXiv preprint arXiv:1910.09483* (2019)
- Matsumoto, M., et al.: A slope of a mountain is a Finsler surface with respect to a time measure. *J. Math. Kyoto Univ.* **29**(1), 17–25 (1989)
- Melleray, J.: On the geometry of Urysohn’s universal metric space. *Topol. Appl.* **154**(2), 384–403 (2007)
- Mémoli, F.: On the Use of Gromov–Hausdorff Distances for Shape Comparison. *The Eurographics Association* (2007)
- Mémoli, F.: Gromov–Wasserstein distances and the metric approach to object matching. *Found. Comput. Math.* **11**(4), 417–487 (2011)
- Mémoli, F.: Some properties of Gromov–Hausdorff distances. *Discret. Comput. Geom.* (2012). <https://doi.org/10.1007/s00454-012-9406-8>
- Mémoli, F.: A distance between filtered spaces via tripods. *arXiv preprint arXiv:1704.03965* (2017)
- Mémoli, F., Sapiro, G.: Comparing point clouds. In: *SGP ’04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry processing*, pp. 32–40, New York, NY, USA, (2004). ACM
- Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.* **5**(3), 313–347 (2005)
- Mennucci, A.C.G.: On asymmetric distances. *Anal. Geom. Metr. Spaces* **1**, 200–231 (2013)
- Munkres, J.R.: *Elements of algebraic topology*, vol. 7. Addison-Wesley Reading (1984)
- Munkres, J.R.: *Topology*. Prentice Hall (2000)
- Newman, M.: *Networks: An Introduction*. Oxford University Press, Oxford (2010)
- Newman, M.E.J.: Estimating network structure from unreliable measurements. *Phys. Rev. E* **98**(6), 062321 (2018)
- Newman, M.E.J.: Network structure from rich but noisy data. *Nat. Phys.* **14**(6), 542–545 (2018)
- Niemytzki, V.W.: On the “third axiom of metric space”. *Trans. Am. Math. Soc.* **29**(3), 507–513 (1927)
- Ohta, S.: Barycenters in alexandrov spaces of curvature bounded below. *Adv. Geom.* **12**(4), 571–587 (2012)
- O’Keefe, J., Dostrovsky, J.: The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**(1), 171–175 (1971)
- O’Keefe, J., Nadel, L.: *The Hippocampus as a Cognitive Map*, vol. 3. Clarendon Press, Oxford (1978)
- Oles, V., Lemons, N., Panchenko, A.: Efficient estimation of a Gromov–Hausdorff distance between unweighted graphs. *arXiv preprint arXiv:1909.09772* (2019)
- Pardalos, P.M., Wolkowicz, H. (eds.): *Quadratic Assignment and Related Problems*. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 16. American Mathematical Society, Providence (1994)
- Pennec, X., Fillard, P., Ayache, N.: A Riemannian framework for tensor computing. *Int. J. Comput. Vis.* **66**(1), 41–66 (2006)
- Perez, L., Husain, K., Chowdhury, S., Schweinhart, B., Galstyan, V., Mehta, P., Fakhri, N., Murugan, A.: Quantifying scale-dependent irreversibility using persistent homology. *Bull. Am. Phys. Soc.* **65** (2020)
- Perrault-Joncas, D.C., Meila, M.: Directed graph embedding: an algorithm based on continuous limits of Laplacian-type operators. In: *Advances in Neural Information Processing Systems*, pp. 990–998 (2011)
- Petersen, P.: *Riemannian Geometry*, vol. 171. Springer, Berlin (2006)
- Peyré, G., Cuturi, M., Solomon, J.: Gromov–Wasserstein averaging of kernel and distance matrices. In: *International Conference on Machine Learning*, pp. 2664–2672 (2016)

- Pinto, G.V.F.: Motivic constructions on graphs and networks with stability results (2020)
- Pitcher, A.D., Chittenden, E.W.: On the foundations of the calcul fonctionnel de Fréchet. *Trans. Am. Math. Soc.* **19**(1), 66–78 (1918)
- Robins, V.: Towards computing homology from finite approximations. In: *Topology Proceedings*, vol. 24, pp. 503–532 (1999)
- Sabau, S.V., Shibuya, K., Shimada, H.: Metric structures associated to Finsler metrics. *arXiv preprint [arXiv:1305.5880](https://arxiv.org/abs/1305.5880)* (2013)
- Schmiedl, F.: Shape matching and mesh segmentation: mathematical analysis, algorithms and an application in automated manufacturing. PhD thesis, München, Technische Universität München, Diss., 2015 (2015)
- Schmiedl, F.: Computational aspects of the Gromov–Hausdorff distance and its application in non-rigid shape matching. *Discret. Comput. Geom.* **57**(4), 854–880 (2017)
- Semple, C., Steel, M.A.: *Phylogenetics*, vol. 24. Oxford University Press on Demand, Oxford (2003)
- Shen, Y.-B., Zhao, W.: Gromov pre-compactness theorems for nonreversible Finsler manifolds. *Differ. Geom. Appl.* **28**(5), 565–581 (2010)
- Singh, G., Mémoli, F., Carlsson, G.: Topological methods for the analysis of high dimensional data sets and 3d object recognition. In: *Symposium on Point-Based Graphics*, pp. 91–100 (2007)
- Smith, Z., Chowdhury, S., Mémoli, F.: Hierarchical representations of network data with optimal distortion bounds. In: *2016 50th Asilomar Conference on Signals, Systems and Computers*, pp. 1834–1838. IEEE (2016)
- Solomon, J., Peyré, G., Kim, V.G., Sra, S.: Entropic metric alignment for correspondence problems. *ACM Trans. Graph. (TOG)* **35**(4), 72 (2016)
- Sporns, O.: *Networks of the Brain*. MIT press, Cambridge (2011)
- Sporns, O.: *Discovering the Human Connectome*. MIT press, Cambridge (2012)
- Steen, L.A., Seebach, J.A.: Counterexamples in Topology, vol. 18. Springer, Berlin (1978)
- Stojmirović, A., Yi-Kuo, Yu.: Geometric aspects of biological sequence comparison. *J. Comput. Biol.* **16**(4), 579–610 (2009)
- Sturm, K.-T.: The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces. *arXiv preprint [arXiv:1208.0434](https://arxiv.org/abs/1208.0434)* (2012)
- Turner, K.: Rips filtrations for quasimetric spaces and asymmetric functions with stability results. *Algebraic Geom. Topol.* **19**(3), 1135–1170 (2019)
- Tversky, A., Gati, I.: Similarity, separability, and the triangle inequality. *Psychol. Rev.* **89**(2), 123 (1982)
- Venkatesh, M., Jaja, J., Pessoa, L.: Comparing functional connectivity matrices: a geometry-aware approach applied to participant identification. *Neuroimage* **207**, 116398 (2020)
- Waszkiewicz, P.: The local triangle axiom in topology and domain theory. *Appl. Gen. Topol.* **4**(1), 47–70 (2013)
- Xu, H., Luo, D., Carin, L.: Scalable Gromov–Wasserstein learning for graph partitioning and matching. In: *Advances in Neural Information Processing Systems*, pp. 3052–3062 (2019)
- Xu, H., Luo, D., Zha, H., Duke, L.C.: Gromov–Wasserstein learning for graph matching and node embedding. In: *International Conference on Machine Learning*, pp. 6932–6941. PMLR (2019)
- Zaustinsky, E.M.: *Spaces with Non-symmetric Distance*, vol. 34. American Mathematical Society (1959)
- Zhou, D., Huang, J., Schölkopf, B.: Learning with hypergraphs: clustering, classification, and embedding. *Adv. Neural. Inf. Process. Syst.* **19**, 1601–1608 (2006)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.