Formigrams: Clustering Summaries of Dynamic Data*

Woojin Kim[†] Facundo Mémoli[‡]

Abstract

When studying flocking/swarming behaviors in animals one is interested in quantifying and comparing the dynamics of the clustering induced by the coalescence and disbanding of animals in different groups.

Motivated by this, we propose a summarization of time-dependent metric data which captures their time-dependent clustering features which we call *formigrams*. These set-valued functions generalize the notion of dendrogram, a prevalent object in the context of hierarchical clustering.

Also, we define a metric on formigrams for quantifying the degree of structural difference between any two given formigrams. In particular, the restriction of this metric to the collection of dendrograms recovers twice the Gromov-Hausdorff distance between the ultrametric spaces associated to the dendrograms. This fact enables us to show that constant factor approximations to the metric on formigrams cannot be obtained in polynomial time.

Finally, we investigate a sufficient condition for time-dependent metric spaces to be summarized into formigrams. In addition, we prove that this summarization process is stable under perturbations in the input time-dependent metric data.

1 Introduction

Given data represented as a static finite metric space (X, d_X), a hierarchical clustering method finds a hierarchical family of partitions that captures multi-scale features present in the dataset. These hierarchical families of partitions are called *dendrograms* and their visualization is straightforward (see figure on the left).

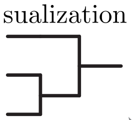

We now turn our attention to a problem of characterizing dynamic data. We model dynamic datasets as time varying finite metric spaces and study a simple generalization of the notion of dendrogram which we call *formigram* - a combination of the words formicarium[1] and diagram (see figure on the right).

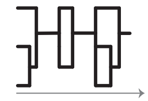

Whereas dendrograms are useful for modeling situations when data points aggregate along a certain scale parameter, formigrams are better suited for representing phenomena when data points may also separate or disband and then regroup at different parameter values. One motivation for considering this scenario comes from the study and characterization of *flocking/swarming/herding* behavior of animals [1, 10, 11, 12, 19, 21, 24, 28], convoys [14], moving clusters [15], or mobile groups [13, 29].

Related work. Let $\mathcal{X}$ be a set of points having piecewise linear trajectories with time-stamped vertices in Euclidean space $\mathbf{R}^d$. Buchin and et al. [3] provided explicit algorithms for studying the grouping structure of $\mathcal{X}$. This was subsequently enriched in [18, 25, 26, 27].

From the set $\mathcal{X}$, the authors of [3] construct a Reeb graph-like structure $\mathcal{R}_{\mathcal{X}}$ which is closely related to the *formigram* derived from $\mathcal{X}$ that we introduce (Section 3 and Theorem 4). The edges of $\mathcal{R}_{\mathcal{X}}$ are labeled by *maximal groups*, and they call $\mathcal{R}_{\mathcal{X}}$ together with these labels the *trajectory grouping structure* of $\mathcal{X}$, enabling the visualization of the life span of maximal groups.

Our contributions.

1. We generalize dendrograms to *formigrams* for the analysis of clustering features of dynamic data, such as dynamic metric spaces or dynamic graphs.

2. Any dendrogram over a finite set X induces an ultrametric on X [7]. Therefore, one can quantify the structural difference between any two dendrograms by computing the Gromov-Hausdorff distance between their two induced ultrametrics [7]. We propose a distance $d_{\mathrm{I}}^{\mathrm{F}}$ between formigrams which generalizes the method above for comparing two dendrograms (Theorems 1 and 2). The desire to obtain such a precise quantification of the difference between two dynamic clusterings was already made explicit in [3, Section 6]. Also, we show that constant factor approximations to $d_{\mathrm{I}}^{\mathrm{F}}$ cannot be obtained in polynomial time (Theorem 3).

*This work was partially supported by NSF grants IIS-1422400, CCF-1526513, and DMS-1723003.

†Department of Mathematics, The Ohio State University, kim.5235@osu.edu

‡Department of Mathematics, The Ohio State University, memoli@math.osu.edu

[1] A formicarium or ant farm is an enclosure for keeping ants under semi-natural conditions [30].

3. As an application, we propose a method for turning any (tame) dynamic metric data into a formigram. This method is closely related to the construction of trajectory grouping structures [3]. In particular, this method turns out to be stable under perturbations in the input dynamic metric data under a certain notion of distance between DMSs that we introduce (Theorem 5).

2 Background

2.1 Dendrograms and treegrams

Partitions and sub-partitions. Let X be a non-empty finite set. We will call any partition P of a subset X' of X a *sub-partition* of X (in particular, any partition of the empty set is defined as the empty set). In this case we call X' the *underlying set of* P.

1. By $\mathcal{P}^{\text{sub}}(X)$, we denote the set of *all sub-partitions of* X, i.e.

$$\mathcal{P}^{\text{sub}}(X) := \{P : \exists X' \subset X \text{ , } P \text{ is a partition of } X'\}.$$

2. By $\mathcal{P}(X)$, we denote the subcollection of $\mathcal{P}^{\text{sub}}(X)$ consisting solely of partitions of the *whole* X.

Given $P, Q \in \mathcal{P}^{\text{sub}}(X)$, by $P \leq Q$ we mean "P is finer than or equal to Q", i.e. for all $B \in P$, there exists $C \in Q$ such that $B \subset C$. For example, let $X = \{x_1, x_2, x_3\}$ and consider the sub-partitions $P := \{\{x_1, x_2\}\}$ and $Q := \{\{x_1, x_2\}, \{x_3\}\}$ of X. Then, it is easy to see that in this case $P \leq Q$.

Dendrograms. A *dendrogram* over a finite set X is any function $\theta_X : \mathbf{R}_+ \to \mathcal{P}(X)$ such that the following properties hold: (1) $\theta_X(0) = \{\{x\} : x \in X\}$, (2) if $t_1 \leq t_2$, then $\theta_X(t_1) \leq \theta_X(t_2)$, (3) there exists $T > 0$ such that and $\theta_X(t) = \{X\}$ for $t \geq T$, (4) for all t there exists $\epsilon > 0$ s.t. $\theta_X(s) = \theta_X(t)$ for $s \in [t, t+\epsilon]$ (right-continuity). See Figure 1 for an example.

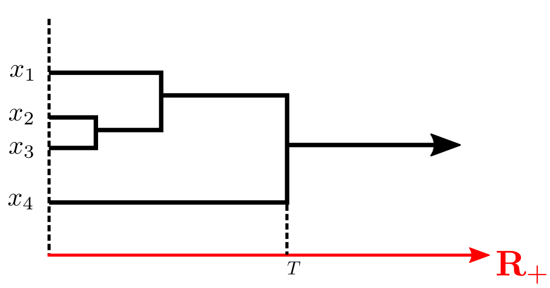

Figure 1: A dendrogram θ_X over the set $X = \{x_1, x_2, x_3, x_4\}$. Notice that $\theta_X(0) = \{\{x_1\}, \{x_2\}, \{x_3\}, \{x_4\}\}$ and $\theta_X(t) = \{X\}$ for all $t \in [T, \infty)$.

Treegrams. Dendrograms can be generalized to treegrams, a visual representation for hierarchical clustering of networks [23]. [2] A *treegram* over a finite set X is any function $\theta_X : \mathbf{R} \to \mathcal{P}^{\text{sub}}(X)$ such that the following properties hold: (1) if $t_1 \leq t_2$, then $\theta_X(t_1) \leq \theta_X(t_2)$, (2) (boundedness) there exists $T > 0$ such that $\theta_X(t) = \{X\}$ for $t \geq T$ and $\theta_X(t)$ is empty for $t \leq -T$. (3) for all t there exists $\epsilon > 0$ s.t. $\theta_X(s) = \theta_X(t)$ for $s \in [t, t+\epsilon]$ (right-continuity). See Figure 2 for an example.

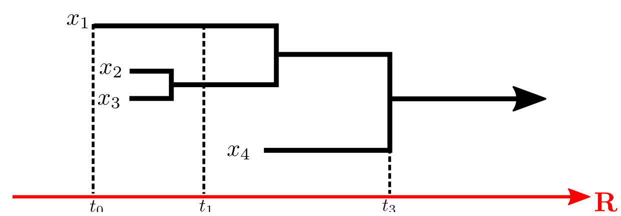

Figure 2: A treegram θ_X over the set $X = \{x_1, x_2, x_3, x_4\}$. Notice that $\theta_X(t) = \emptyset$ for $t \in (-\infty, t_0)$. Also, $\theta_X(t_0) = \{\{x_1\}\}$, $\theta_X(t_2) = \{\{x_1\}, \{x_2, x_3\}\}$, and $\theta_X(t) = \{X\}$ for all $t \in [t_3, \infty)$.

2.2 A distance between dendrograms

In this section we review the method of [7] for quantifying the structural difference between dendrograms. In short, we compare two dendrograms over sets X and Y by comparing their associated ultrametrics on X and Y, respectively.

Dendrograms and their associated ultrametrics. An *ultrametric space* (X, u_X) is a metric space satisfying the *strong triangle inequality*: for all $x, x', x'' \in X$, $u_X(x, x') \leq \max\{u_X(x, x''), u_X(x'', x')\}$.

Let X be a finite set and let $\theta_X : \mathbf{R}_+ \to \mathcal{P}(X)$ be a dendrogram over X. Recall from [7] that this θ_X induces a canonical ultrametric $u_{\theta_X} : X \times X \to \mathbf{R}_+$ on X defined by

$$u_{\theta_X}(x, x') := \inf\{\varepsilon \geq 0 : x, x' \text{ belong to the same block of } \theta_X(\varepsilon)\}.$$

For example, for the dendrogram θ_X depicted in Figure 1, it is easy to observe that $u_{\theta_X}(x_1, x_4) = T$.

Reciprocally, any ultrametric space (X, u_X) induces a dendrogram θ_X over X [7].

The Gromov-Hausdorff distance [4, Ch 7]. The Gromov-Hausdorff distance quantifies how far two compact metric spaces are from being isometric. This distance is widely used in applications such as shape comparison (for example, see [20]). In order to define the

[2] In order to regard a dendrogram $\theta_X : \mathbf{R}_+ \to \mathcal{P}(X)$ as a treegram, trivially extend θ_X to the whole $\mathbf{R}$: for $t \in (-\infty, 0)$, let $\theta_X(t) := \emptyset \in \mathcal{P}^{\text{sub}}(X)$ by definition.

Gromov-Hausdorff distance, one needs the notion of *correspondence.*

For sets X and Y, a subset $R \subset X \times Y$ is said to be a *correspondence* (between X and Y) if and and only if (1) for every $x \in X$, there exists $y \in Y$ such that $(x, y) \in R$, and (2) for every $y \in Y$, there exists $x \in X$ such that $(x, y) \in R$.

Let (X, d_X) and (Y, d_Y) be any two compact metric spaces. The Gromov-Hausdorff distance between (X, d_X) and (Y, d_Y) is defined by

$$d_{\mathrm{GH}}\left((X, d_X), (Y, d_Y)\right) := \frac{1}{2}\inf_R \sup_{\substack{(x,y)\in R\\(x',y')\in R}} \left|d_X(x, x') - d_Y(y, y')\right|,$$

where the infimum is taken over all correspondences between X and Y. Note that in the case where $(X, d_X), (Y, d_Y)$ are finite metric spaces, the infimum and the supremum above can be replaced with the minimum and the maximum, respectively.

A distance between dendrograms. Let θ_X and θ_Y be dendrograms over finite sets X and Y, respectively. One defines the *Gromov-Hausdorff distance [7] between the dendrograms* θ_X and θ_Y as

$$d_{\mathrm{GH}}(\theta_X, \theta_Y) := d_{\mathrm{GH}}\left((X, u_{\theta_X}), (Y, u_{\theta_Y})\right),$$

where u_{θ_X} and u_{θ_Y} are the ultrametrics associated to the dendrograms θ_X and θ_Y, respectively.

2.3 Finest common coarsening of (sub-)partitions

For a set X, we know that there exists a canonical one-to-one correspondence between the collection of all equivalence relations on X and the collection of all partitions $\mathcal{P}(X)$ of X. We will extend this correspondence in a certain way for defining the notion of *finest common coarsening* in the collection $\mathcal{P}^{\mathrm{sub}}(X)$ of all sub-partitions of X.

Sub-equivalence relations. Let X be a non-empty set. Let $\sim$ be any equivalence relation on any subset $X' \subset X$.[3] We call the relation $\sim$ a *sub-equivalence relation on* X. We also call X' the *underlying set* of $\sim$, which is identical to $\{x \in X : (x, x) \in \sim\}$.

Clearly, any equivalence relation on X is also a sub-equivalence relation with underlying set X.

There is the canonical one-to-one correspondence between the collection of all sub-equivalence relations on X and the collection $\mathcal{P}^{\mathrm{sub}}(X)$ of all sub-partitions of X: Any sub-equivalence relation $\sim$ on X corresponds to the sub-partition P with underlying set $X' = \{x \in X : (x, x) \in \sim\}$ such that $x \sim y$ iff x and y belong to the same block $B \in P$. Reciprocally, to any sub-partition P of X, one can associate the unique sub-equivalence relation $\sim_P$ on X defined by $x \sim_P y$ if and only if x and y belong to the same block $B \in P$.

[3]In particular, the unique equivalence relation on the empty set $\emptyset$ is $\emptyset$.

Sub-equivalence closure. Let X be a non-empty set. For an index set I, suppose that $\{\sim_i \subset X \times X : i \in I\}$ is a collection of sub-equivalence relations on X. The *sub-equivalence closure* of the collection $\{\sim_i \subset X \times X : i \in I\}$ is defined to be the transitive closure of the relation $\cup_{i\in I} \sim_i$ on X. In other words, by the *sub-equivalence closure* of the collection $\{\sim_i \subset X \times X : i \in I\}$, we mean the minimal sub-equivalence relation containing $\sim_i$ for all $i \in I$.

Finest common coarsening. Let $\{P_i\}_{i\in I}$ be any subcollection of $\mathcal{P}^{\mathrm{sub}}(X)$. For each $i \in I$, let $\sim_i$ be the sub-equivalence relation on X corresponding to P_i. By $\bigvee_{i\in I} P_i$, we mean the sub-partition of X corresponding to the sub-equivalence closure of the collection $\{\sim_i \subset X \times X : i \in I\}$. We will refer to $\bigvee_{i\in I} P_i$ as *the finest common coarsening of the collection* $\{P_i\}_{i\in I}$.

For example, let $X := \{x, y, z, w\}$. For $P_1 = \{\{x\}, \{y\}\}, P_2 = \{\{y, z\}\}$, and $P_3 = \{\{x, w\}\}$ in $\mathcal{P}^{\mathrm{sub}}(X)$, we have:

1. $\bigvee_{i=1}^2 P_i = \{\{x\}, \{y, z\}\} \in \mathcal{P}^{\mathrm{sub}}(X)$, and
2. $\bigvee_{i=1}^3 P_i = \{\{x, w\}, \{y, z\}\} \in \mathcal{P}(X)$.

3 Formigrams

Although the notions of dendrogram or treegram are useful when representing the output of a hierarchical clustering method (i.e. when partitions only become coarser with the increase of a parameter), in order to represent the diverse clustering behaviors of dynamic datasets we need a more flexible concept allowing for possible refinement of partitions. Here we suggest a "zigzag like" notion of dendrograms that we call *formigram.* We allow partitions to become finer sometimes, but require that partitions defined by a formigram change only finitely many times in any finite interval for visualization.

3.1 The definition of a formigram

Formigrams. A *formigram* over a finite set X is any function $\theta_X : \mathbf{R} \to \mathcal{P}^{\mathrm{sub}}(X)$ such that:

1. (Tameness) the set $\mathrm{crit}(\theta_X)$ of points of discontinuity of θ_X is locally finite.[4] We call the elements of $\mathrm{crit}(\theta_X)$ *the critical points of* θ_X.

[4]To say that $\mathrm{crit}(\theta_X)$ is locally finite means that for any bounded interval $I \subset \mathbf{R}$, the cardinality of $I \cap \mathrm{crit}(\theta_X)$ is finite. The purpose of this condition is twofold: on the one hand,

2. (Interval lifespan) for every $x \in X$, the set $I_x := \{t \in \mathbf{R} : x \in B \in \theta_X(t)\}$, said to be the *lifespan of* x, is a non-empty closed interval,
3. (Comparability) for every point $c \in \mathbf{R}$ it holds that $\theta_X(c-\varepsilon) \leq \theta_X(c) \geq \theta_X(c+\varepsilon)$ for all sufficiently small $\varepsilon > 0$.[5]

Note that the definition of formigrams generalizes those of dendrograms and treegrams.[6] In other words, every dendrogram and every treegram are formigrams. See Figure 3 for an example.

$$\theta_X(t) = \begin{cases} \{\{x_1,x_2,x_3\}\}, & t \in (-\infty, t_0) \\ \{\{x_1,x_2,x_3\},\{x_4\}\}, & t \in [t_0, t_1) \\ \{\{x_1,x_2,x_3,x_4\}\}, & t \in [t_1,t_2] \cup [t_5, \infty) \\ \{\{x_1,x_2\},\{x_3,x_4\}\}, & t \in (t_2,t_3] \cup [t_4,t_5) \\ \{\{x_1\},\{x_2\},\{x_3,x_4\}\}, & t \in (t_3,t_4). \end{cases}$$

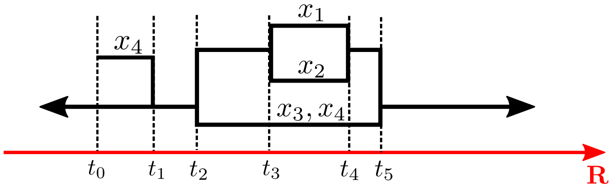

Figure 3: Top: The specification of a formigram θ_X over the set $X = \{x_1,x_2,x_3,x_4\}$. Bottom: A graphical representation of the formigram θ_X.

3.2 A distance between formigrams

In this section we introduce a (pseudo) metric on the collection of all formigrams. This metric quantifies the structural difference between two grouping/disbanding behaviors over time. In particular, when restricting this metric to the collection of dendrograms, (twice) the Gromov-Hausdorff distance between dendrograms is recovered (Theorem 2).

Partition morphisms. Before introducing a metric on formigrams, we first establish a method for interconnecting any two partitions with possibly different underlying sets. Recall that for any sets X and Y, a *multivalued map* $\varphi : X \rightrightarrows Y$ is a relation between X and Y such that for all $x \in X$, there exists (a not necessarily unique) $y \in Y$ with $(x,y) \in \varphi$.[7] For $x \in X$, the *image* $\varphi(x)$ of x is defined to be the set $\{y \in Y : (x,y) \in \varphi\}$.

For any two sets X and Y, let $P_X \in \mathcal{P}(X)$ and $P_Y \in \mathcal{P}(Y)$. Any multivalued map $\varphi : X \rightrightarrows Y$ (or map $\varphi : X \to Y$) is said to be a *partition morphism between* P_X *and* P_Y if for any $x, x' \in X$ belonging to the same block of P_X, their images $\varphi(x), \varphi(x')$ are included in the same block of P_Y (note that $\varphi(x), \varphi(x')$ can be sets containing more than one element). In this case, we write $P_X \leq_\varphi Q_Y$.

If $P_X \leq_\varphi Q_Y$, then there exists the canonical induced map $\varphi^* : P_X \to P_Y$ defined by sending each block $B \in P_X$ to the block $C \in P_Y$ such that $\varphi(B) \subset C$.

A distance between formigrams. Exploiting the fact that any formigram is a *"stack" of (sub-)partitions* of a specific set, we now introduce the *interleaving distance* $d_{\mathrm{I}}^{\mathrm{F}}$ on the collection of all formigrams. The construction of $d_{\mathrm{I}}^{\mathrm{F}}$ is inspired by the interleaving distance for Reeb graphs [9].

Let θ_X be a formigram over X and let $I \subset \mathbf{R}$ be an interval. We define $\bigvee_I \theta_X$ to be the finest common coarsening of the collection $\{\theta_X(t) : t \in I\}$ of sub-partitions of X. Also, for any $t \in \mathbf{R}$, define $[t]^\varepsilon := [t-\varepsilon, t+\varepsilon] \subset \mathbf{R}$.

Let θ_X and θ_Y be any two formigrams over X and Y, respectively. θ_X and θ_Y are said to be ε-*interleaved* if there exists a correspondence R between X and Y satisfying the following:

1. For any $(x,y) \in R$ and any $t \in \mathbf{R}$,
   (a) if x is in the underlying set of $\theta_X(t)$, then y is in the underlying set of $\bigvee_{[t]^\varepsilon} \theta_Y$.[8]
   (b) if y is in the underlying set of $\theta_Y(t)$, then x is in the underlying set of $\bigvee_{[t]^\varepsilon} \theta_X$.
2. For all $t \in \mathbf{R}$,
$$\theta_X(t) \leq_R \bigvee_{[t]^\varepsilon} \theta_Y \quad \text{and} \quad \theta_Y(t) \leq_{R^{-1}} \bigvee_{[t]^\varepsilon} \theta_X,$$
   where $R^{-1} = \{(y,x) \in Y \times X : (x,y) \in R\}$.

We call any such R an ε-*correspondence between* θ_X *and* θ_Y.[9] The interleaving distance $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y)$ between θ_X and θ_Y is defined by the infimum of $\varepsilon \geq 0$ for which there exists an ε-correspondence between θ_X and θ_Y. If there is no ε-correspondence between θ_X and θ_Y for any $\varepsilon \geq 0$, then we declare $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y) = +\infty$.

---

we want to guarantee easy visualization, on the other hand this condition is necessary for the simplification process of formigrams via zigzag persistence theory [5]. We refer the interested readers to [16].

[5] If θ_X is not continuous at c, then at least one of the relations of $\theta_X(c-\varepsilon) \leq \theta_X(c) \geq \theta_X(c+\varepsilon)$ would be strict for small $\varepsilon > 0$. But if c is a continuity point of θ_X, then $\theta_X(c-\varepsilon) = \theta_X(c) = \theta_X(c+\varepsilon)$ for small $\varepsilon > 0$.

[6] In order to regard a dendrogram $\theta_X : \mathbf{R}_+ \to \mathcal{P}(X)$ as a formigram, trivially extend θ_X to the whole $\mathbf{R}$: for $t \in (-\infty, 0)$, let $\theta_X(t) := \emptyset \in \mathcal{P}^{\mathrm{sub}}(X)$ by definition.

[7] In particular, any correspondence R between X and Y is a multivalued map.

[8] We remark that this condition is equivalent to saying that if x is in the underlying set of $\theta_X(t)$, then there exists $t_0 \in [t]^\varepsilon$ such that y is in the underlying set of $\theta_Y(t_0)$.

[9] Note that if R is an ε-correspondence between θ_X and θ_Y, then for any $\varepsilon' > \varepsilon$, R is also an ε'-correspondence between θ_X and θ_Y.

Theorem 1 *$d_{\mathrm{I}}^{\mathrm{F}}$ is an extended pseudo-metric on formigrams.*

See Appendix A for the proof of Theorem 1. For example, consider any formigram θ_X over a finite set X and let $\tau > 0$. Define another formigram θ_X^τ as $\theta_X^\tau(t) := \theta_X(t+\tau)$ for $t \in \mathbf{R}$. Then, it is not difficult to verify that $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_X^\tau) \leq \tau$ by checking that $R_X := \{(x,x) : x \in X\}$ is a τ-correspondence between θ_X and θ_X^τ.

Theorem 2 *$d_{\mathrm{I}}^{\mathrm{F}}$ generalizes the Gromov-Hausdorff distance between dendrograms. Namely, for any dendrograms θ_X and θ_Y over X and Y respectively,*

$$d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y) = 2\ d_{\mathrm{GH}}(\theta_X, \theta_Y).$$

Proof. Recall that by definition

$$d_{\mathrm{GH}}(\theta_X, \theta_Y) = d_{\mathrm{GH}}\left((X, u_{\theta_X}), (Y, u_{\theta_Y})\right)$$

where u_{θ_X} and u_{θ_Y} are the ultrametrics associated to the dendrograms θ_X and θ_Y, respectively. Therefore, we will show that $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y) = 2\ d_{\mathrm{GH}}\left((X, u_{\theta_X}), (Y, u_{\theta_Y})\right)$. First we show "$\geq$". If $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y) = \infty$, there is nothing to prove and hence we assume that $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y)$ is finite. Then, there exists an ε-correspondence $R \subset X \times Y$ between the two dendrograms θ_X and θ_Y for some $\varepsilon \geq 0$, implying that $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y) \leq \varepsilon$. Pick any $(x,y), (x',y') \in R$ and let $t := u_{\theta_X}(x,x')$. Then, x, x' belong to the same block of the partition $\theta_X(t)$. Since $\theta_X(t) \leq_R \bigvee_{[t]^\varepsilon} \theta_Y$, y, y' must belong to the same block of $\bigvee_{[t]^\varepsilon} \theta_Y$. Also, since θ_Y is a dendrogram, $\theta_Y(s_1) \leq \theta_Y(s_2)$ for any $s_1 \leq s_2$, and thus $\bigvee_{[t]^\varepsilon} \theta_Y = \theta_Y(t+\varepsilon)$. Therefore, y, y' belong to the same block of $\theta_Y(t+\varepsilon)$, and in turn $u_{\theta_Y}(y,y') \leq t + \varepsilon = u_{\theta_X}(x,x') + \varepsilon$. By symmetry, we also have $u_{\theta_X}(x,x') \leq u_{\theta_Y}(y,y') + \varepsilon$. Therefore, by the definition of $d_{\mathrm{GH}}((X, u_{\theta_X}), (Y, u_{\theta_Y}))$, we have $d_{\mathrm{GH}}((X, u_{\theta_X}), (Y, u_{\theta_Y})) \leq \varepsilon/2$.

Next, we prove "$\leq$". Let R be a correspondence between X and Y such that for all $(x,y), (x',y') \in R$, $|u_{\theta_X}(x,x') - u_{\theta_Y}(y,y')| \leq \varepsilon$, implying that $d_{\mathrm{GH}}((X, u_{\theta_X}), (Y, u_{\theta_Y})) \leq \varepsilon/2$. We wish to show that $\theta_X(t) \leq_R \theta_Y(t+\varepsilon)$ for all $t \in \mathbf{R}$. For $t < 0$, since $\theta_X(t) = \theta_Y(t) = \emptyset$, we trivially have $\theta_X(t) \leq_R \theta_Y(t+\varepsilon)$. Now pick any $t \geq 0$ and any $(x,y), (x',y') \in R$. Assume that x, x' belong to the same block of $\theta_X(t)$, implying that $u_{\theta_X}(x,x') \leq t$. Since $|u_{\theta_X}(x,x') - u_{\theta_Y}(y,y')| \leq \varepsilon$, we know $u_{\theta_Y}(y,y') \leq t + \varepsilon$, and hence y, y' belong to the same block of $\theta_Y(t+\varepsilon)$. Therefore, $\theta_X(t) \leq_R \theta_Y(t+\varepsilon)$ for all $t \in \mathbf{R}$. By symmetry, $\theta_X(t) \leq_R \theta_Y(t+\varepsilon)$ for all $t \in \mathbf{R}$ as well, completing the proof. □

Theorem 3 (Complexity of computing $d_{\mathrm{I}}^{\mathrm{F}}$) *Fix $\rho \in (1, 6)$. It is not possible to obtain a ρ approximation to the distance $d_{\mathrm{I}}^{\mathrm{F}}\left((X, \theta_X), (Y, \theta_Y)\right)$ between formigrams in time polynomial on $|X|, |Y|, |\mathrm{crit}(\theta_X)|$, $|\mathrm{crit}(\theta_Y)|$ unless $P = NP$.*

Proof. Pick any two dendrograms θ_X and θ_Y and invoke Theorem 2 to reduce the problem to the computation of the Gromov-Hausdorff distance

$$\Delta := d_{\mathrm{GH}}((X, u_{\theta_X}), (Y, u_{\theta_Y}))$$

between the ultrametric spaces $(X, u_{\theta_X}), (Y, u_{\theta_Y})$ associated to the dendrograms. However, according to [22, Corollary 3.8], Δ cannot be approximated within any factor less than 3 in polynomial time, unless $P = NP$. The author shows this by observing that any instance of the 3-partition problem can be reduced to an instance of the bottleneck ∞-Gromov-Hausdorff distance (∞-BGHD) problem between ultrametric spaces (see [22, p.865]). The proof follows. □

4 Application: Visualization of clustering features of dynamic metric data

In this section we explain how to extract scale dependent clustering features from time-dependent metric spaces in the form of formigrams. Furthermore, we will show that this summarization process is stable under perturbations in the input time-dependent metric spaces.

4.1 Dynamic metric spaces (DMSs)

Recall that a pseudo-metric space is a pair (X, d_X) where X is a (non-empty) set and $d_X : X \times X \to \mathbf{R}_+$ is a symmetric function which satisfies the triangle inequality, and such that $d_X(x,x) = 0$ for all $x \in X$. d_X is called the pseudo-metric. Note that one does not require that $d_X(x,x') = 0$ implies that $x = x'$ like in the case of standard metric spaces.

Dynamic metric spaces (DMSs). A *dynamic metric space* is a pair $\gamma_X = (X, d_X(\cdot))$ where X is a non-empty finite set and $d_X : \mathbf{R} \times X \times X \to \mathbf{R}_+$ satisfies:

1. For every $t \in \mathbf{R}$, $\gamma_X(t) = (X, d_X(t))$ is a pseudo-metric space.
2. There exists $t_0 \in \mathbf{R}$ such that $\gamma_X(t_0)$ is a (standard) metric space.
3. For fixed $x, x' \in X$, $d_X(\cdot)(x,x') : \mathbf{R} \to \mathbf{R}_+$ is continuous.

We refer to t as the *time* parameter. Condition 2 above is assumed since otherwise one could substitute the DMSs γ_X by another DMSs $\gamma_{X'}$ over a set X' which satisfies $|X'| < |X|$, and such that $\gamma_{X'}$ is point-wisely equivalent to γ_X.

A family of examples of DMSs is given by n particles/animals moving continuously inside an environment $\Omega \subset \mathbf{R}^d$ where particles are allowed to coalesce. If the n trajectories are $p_1(t), \ldots, p_n(t) \in \mathbf{R}^d$, then let

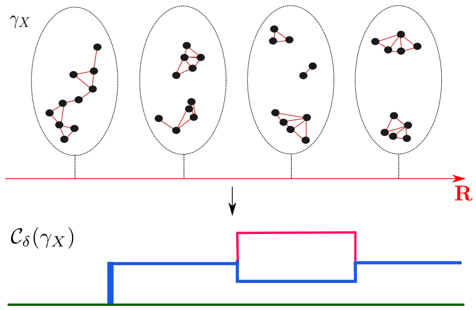

Figure 4: Top: A collection of moving particles (a DMS γ_X) is depicted over the time $\mathbf{R}$. In particular, any two points are connected by an edge if their distance does not exceed a certain $\delta > 0$. Bottom: The formigram $\mathcal{C}_\delta(\gamma_X)$ summarizes the clustering features of γ_X at the scale δ.

$P := \{1, \dots, n\}$ and define a DMS $\gamma_P := (P, d_P(\cdot))$ as follows: for $t \in \mathbf{R}$ and $i, j \in \{1, \dots, n\}$, let $d_P(t)(i,j) := \|p_i(t) - p_j(t)\|$, where $\|\cdot\|$ denotes the Euclidean norm.

Tame DMSs. We introduce a notion of *tameness* of DMS which will ultimately ensure that one can associate formigrams to tame DMSs. We first define *tame* functions $f : \mathbf{R} \to \mathbf{R}$: a continuous function $f : \mathbf{R} \to \mathbf{R}$ is *tame*, if for any $c \in \mathbf{R}$ and any finite interval $I \subset \mathbf{R}$, the set $f^{-1}(c) \cap I \subset \mathbf{R}$ is empty or has only finitely many connected components. For instance, polynomial functions (in particular, constant functions) and piecewise linear functions (with locally finitely many critical points) on $\mathbf{R}$ are tame. We say that a DMS $\gamma_X = (X, d_X(\cdot))$ is *tame* if for any $x, x' \in X$ the function $d_X(\cdot)(x, x') : \mathbf{R} \to \mathbf{R}_+$ is tame.

4.2 δ-clustering method for DMSs

δ-Clustering Method. Let $\delta \geq 0$. Recall (flat) single linkage clustering: Given any finite (pseudo-)metric space (X, d_X), define the partition $\mathcal{C}_\delta(X, d_X) := X/\sim_\delta$ where $\sim_\delta$ stands for the equivalence relation on X defined by $x \sim_\delta x'$ if and only if there exists a sequence $x = x_1, x_2, \dots, x_n = x'$ of points in X such that $d_X(x_i, x_{i+1}) \leq \delta$ for each $i \in \{1, \dots, n-1\}$.

From DMSs to Formigrams. We describe the process that, given a connectivity parameter $\delta \geq 0$, associates a formigram to any tame DMS:

Theorem 4 *Let γ_X be a tame DMS and fix $\delta \geq 0$. Then, the function $\mathcal{C}_\delta(\gamma_X) : \mathbf{R} \to \mathcal{P}(X)$ defined by $\mathcal{C}_\delta(\gamma_X)(t) = \mathcal{C}_\delta(\gamma_X(t))$ for $t \in \mathbf{R}$ is a formigram.*

See Figure 4 for an illustration. We prove Theorem 4 in Appendix A.

4.3 Stability of δ-clustering method for DMSs.

It turns out that the construction of formigrams from DMSs described in Theorem 4 is stable under perturbations in the input DMSs under a certain notion of distance between DMSs described below. Structurally, this distance is a hybrid between the Gromov-Hausdorff distance and the interleaving distance [2, 8] for Reeb graphs [9].

A distance between DMSs. Let γ_X, γ_Y be DMSs and $\varepsilon \geq 0$. We say that γ_X and γ_Y are *ε-interleaved* if there exists a correspondence R between X and Y such that $(*)$ $\forall (x,y), (x',y') \in R$, $\forall t \in \mathbf{R}$,

1. $\min_{s \in [t]^\varepsilon} d_Y(s)(y, y') \leq d_X(t)(x, x')$ and,
2. $\min_{s \in [t]^\varepsilon} d_X(s)(x, x') \leq d_Y(t)(y, y')$.

When γ_X and γ_Y are ε-interleaved we write $\gamma_X \approx_\varepsilon \gamma_Y$. The *interleaving distance* between γ_X and γ_Y is defined by $d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Y) := \inf\{\varepsilon \geq 0 : \gamma_X \approx_\varepsilon \gamma_Y\}$. If γ_X and γ_Y are not ε-interleaved for any $\varepsilon \geq 0$, declare $d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Y) = +\infty$. Also, any correspondence R satisfying $(*)$ is called an *ε-correspondence* between γ_X and γ_Y.

In Appendix B, we show that $d_{\mathrm{I}}^{\mathrm{dyn}}$ is indeed an extended metric on DMSs (Theorem 6).

Theorem 5 (Stability theorem) *For any tame DMSs γ_X, γ_Y and any $\delta \geq 0$, let $\theta_X := \mathcal{C}_\delta(\gamma_X)$ and $\theta_Y := \mathcal{C}_\delta(\gamma_Y)$ as in Theorem 4. Then,*

$$d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y) \leq d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Y).$$

See Appendix B for the proof of Theorem 5.

5 Conclusion and Discussion

We introduced formigrams: a generalization of the notion of dendrograms that is useful for characterizing and visualizing the clustering features of DMSs. We clarified a sufficient condition (tameness) for DMSs to admit a summarization as formigrams.

We also defined the distances $d_{\mathrm{I}}^{\mathrm{F}}$ and $d_{\mathrm{I}}^{\mathrm{dyn}}$ on formigrams and on DMSs, respectively, and showed that the δ-clustering method for DMSs is stable under perturbations in the input DMSs in terms of $d_{\mathrm{I}}^{\mathrm{F}}$ and $d_{\mathrm{I}}^{\mathrm{dyn}}$. Specifically, it is noteworthy that $d_{\mathrm{I}}^{\mathrm{F}}$ generalizes the Gromov-Hausdorff distance on dendrograms.

In [17], due to the high cost of computing $d_{\mathrm{I}}^{\mathrm{F}}$, we carry out a classification task for different flocking behaviors by making use of a tractable lower bound for $d_{\mathrm{I}}^{\mathrm{F}}$. The nature of this lower bound is related to zigzag persistence theory [5, 6]: One can find theoretical details in [16].

References

[1] M. Benkert, J. Gudmundsson, F. Hübner, and T. Wolle. Reporting flock patterns. *Computational Geometry*, 41(3):111–125, 2008.

[2] P. Bubenik and J. A. Scott. Categorification of persistent homology. *Discrete & Computational Geometry*, 51(3):600–627, 2014.

[3] K. Buchin, M. Buchin, M. J. van Kreveld, B. Speckmann, and F. Staals. Trajectory grouping structure. *JoCG*, 6(1):75–98, 2015.

[4] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry*, volume 33 of *AMS Graduate Studies in Math.* American Mathematical Society, 2001.

[5] G. Carlsson and V. De Silva. Zigzag persistence. *Foundations of computational mathematics*, 10(4):367–405, 2010.

[6] G. Carlsson, V. De Silva, and D. Morozov. Zigzag persistent homology and real-valued functions. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 247–256. ACM, 2009.

[7] G. Carlsson and F. Mémoli. Characterization, stability and convergence of hierarchical clustering methods. *Journal of Machine Learning Research*, 11:1425–1470, 2010.

[8] F. Chazal, D. Cohen-Steiner, M. Glisse, L. J. Guibas, and S. Oudot. Proximity of persistence modules and their diagrams. In *Proc. 25th ACM Sympos. on Comput. Geom.*, pages 237–246, 2009.

[9] V. De Silva, E. Munch, and A. Patel. Categorified reeb graphs. *Discrete & Computational Geometry*, 55(4):854–906, 2016.

[10] J. Gudmundsson and M. van Kreveld. Computing longest duration flocks in trajectory data. In *Proceedings of the 14th annual ACM international symposium on Advances in geographic information systems*, pages 35–42. ACM, 2006.

[11] J. Gudmundsson, M. van Kreveld, and B. Speckmann. Efficient detection of patterns in 2d trajectories of moving points. *Geoinformatica*, 11(2):195–215, 2007.

[12] Y. Huang, C. Chen, and P. Dong. Modeling herds and their evolvements from trajectory data. In *International Conference on Geographic Information Science*, pages 90–105. Springer, 2008.

[13] S.-Y. Hwang, Y.-H. Liu, J.-K. Chiu, and E.-P. Lim. Mining mobile group patterns: A trajectory-based approach. In *PAKDD*, volume 3518, pages 713–718. Springer, 2005.

[14] H. Jeung, M. L. Yiu, X. Zhou, C. S. Jensen, and H. T. Shen. Discovery of convoys in trajectory databases. *Proceedings of the VLDB Endowment*, 1(1):1068–1080, 2008.

[15] P. Kalnis, N. Mamoulis, and S. Bakiras. On discovering moving clusters in spatio-temporal data. In *SSTD*, volume 3633, pages 364–381. Springer, 2005.

[16] W. Kim and F. Memoli. Stable signatures for dynamic metric spaces via zigzag persistent homology. *arXiv preprint arXiv:1712.04064*, 2017.

[17] W. Kim, F. Mémoli, and Z. Smith. https://research.math.osu.edu/networks/formigrams.

[18] I. Kostitsyna, M. J. van Kreveld, M. Löffler, B. Speckmann, and F. Staals. Trajectory grouping structure under geodesic distance. In *31st International Symposium on Computational Geometry, SoCG 2015, June 22-25, 2015, Eindhoven, The Netherlands*, pages 674–688, 2015.

[19] Z. Li, B. Ding, J. Han, and R. Kays. Swarm: Mining relaxed temporal moving object clusters. *Proceedings of the VLDB Endowment*, 3(1-2):723–734, 2010.

[20] F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5(3):313–347, 2005.

[21] J. K. Parrish and W. M. Hamner. *Animal groups in three dimensions: how species aggregate*. Cambridge University Press, 1997.

[22] F. Schmiedl. Computational aspects of the Gromov–Hausdorff distance and its application in non-rigid shape matching. *Discrete & Computational Geometry*, 57(4):854–880, 2017.

[23] Z. Smith, S. Chowdhury, and F. Mémoli. Hierarchical representations of network data with optimal distortion bounds. In *Signals, Systems and Computers, 2016 50th Asilomar Conference on*, pages 1834–1838. IEEE, 2016.

[24] D. J. Sumpter. *Collective animal behavior*. Princeton University Press, 2010.

[25] A. van Goethem, M. J. van Kreveld, M. Löffler, B. Speckmann, and F. Staals. Grouping time-varying data for interactive exploration. In *32nd International Symposium on Computational Geometry, SoCG 2016, June 14-18, 2016, Boston, MA, USA*, pages 61:1–61:16, 2016.

[26] M. J. van Kreveld, M. Löffler, and F. Staals. Central trajectories. *Journal of Computational Geometry*, 8(1):366–386, 2017.

[27] M. J. van Kreveld, M. Löffler, F. Staals, and L. Wiratma. A refined definition for groups of moving entities and its computation. In *27th International Symposium on Algorithms and Computation, ISAAC 2016, December 12-14, 2016, Sydney, Australia*, pages 48:1–48:12, 2016.

[28] M. R. Vieira, P. Bakalov, and V. J. Tsotras. On-line discovery of flock patterns in spatio-temporal data. In *Proceedings of the 17th ACM SIGSPATIAL international conference on advances in geographic information systems*, pages 286–295. ACM, 2009.

[29] Y. Wang, E.-P. Lim, and S.-Y. Hwang. Efficient algorithms for mining maximal valid groups. *The VLDB JournalThe International Journal on Very Large Data Bases*, 17(3):515–535, 2008.

[30] Wikipedia. Formicarium — Wikipedia, the free encyclopedia, 2017. [Online; accessed 03-June-2017].

Appendix A

Proof of Theorem 1.

Proof. Symmetry of $d_{\mathrm{I}}^{\mathrm{F}}$ is clear and thus we only show reflexivity of $d_{\mathrm{I}}^{\mathrm{F}}$ and the triangle inequality. Let X be any finite set and let θ_X be a formigram over X. Then, one can easily check that $R_X := \{(x,x) : x \in X\}$ is a 0-correspondence between two copies of θ_X, implying that $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_X) = 0$.

Let Y and Z be some finite sets and let θ_Y and θ_Z be any formigrams over Y and Z, respectively. We wish to prove that $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Z) \leq d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y) + d_{\mathrm{I}}^{\mathrm{F}}(\theta_Y, \theta_Z)$. We assume that $d_{\mathrm{I}}^{\mathrm{F}}(\theta_X, \theta_Y)$ and $d_{\mathrm{I}}^{\mathrm{F}}(\theta_Y, \theta_Z)$ are finite because otherwise there is nothing to prove. By this assumption, for some $0 < \varepsilon_1, \varepsilon_2 < \infty$, there are an ε_1-correspondence $R_1 \subset X \times Y$ between θ_X and θ_Y and an ε_2-correspondence $R_2 \subset Y \times Z$ between θ_Y and θ_Z. Define the set $R_2 \circ R_1 \subset X \times Z$ by

$$R_2 \circ R_1 := \{(x,z) \in X \times Z : \exists y \in Y \text{ s.t.} (x,y) \in R_1 \text{ and } (y,z) \in R_2\}.$$

It is not difficult to check that $R_2 \circ R_1$ is a correspondence between X and Z. Therefore, it suffices to prove that $R_2 \circ R_1$ is an $(\varepsilon_1 + \varepsilon_2)$-correspondence between θ_X and θ_Z.

Fix any $(x,z) \in R_2 \circ R_1$ and $t \in \mathbf{R}$. Suppose that x belongs to the underlying set of the sub-partition $\theta_X(t)$ of X. By the definition of $R_2 \circ R_1$, there exists $y \in Y$ such that $(x,y) \in R_1$ and $(y,z) \in R_2$. Since R_1 is an ε_1-correspondence between θ_X and θ_Y, y must be in the underlying set of $\bigvee_{[t]^\varepsilon} \theta_Y$. This implies that there exists $t_0 \in [t]^\varepsilon = [t-\varepsilon, t+\varepsilon]$ such that y belongs to the underlying set of $\theta_Y(t_0)$. Then, invoking that R_2 is an ε_2-correspondence between θ_Y and θ_Z, there exists $t_1 \in [t_0]^{\varepsilon_2} \subset [t]^{\varepsilon_1+\varepsilon_2}$ such that z belongs to the underlying set of $\theta_Z(t_1)$. This implies that z belongs to the underlying set of $\bigvee_{[t]^{\varepsilon_1+\varepsilon_2}} \theta_Z$. Similarly, one can check that if z belongs to the underlying set of $\theta_Z(t)$, then x belongs to the underlying set of $\bigvee_{[t]^{\varepsilon_1+\varepsilon_2}} \theta_X$.

Now, we wish to show that $\theta_X(t) \leq_{R_2 \circ R_1} \bigvee_{[t]^{\varepsilon_1+\varepsilon_2}} \theta_Z$. To this end, it suffices to show that for any $(x,z), (x',z') \in R_2 \circ R_1$, if x, x' belong to the same block of $\theta_X(t)$, then z, z' belong to the same block of $\bigvee_{[t]^{\varepsilon_1+\varepsilon_2}} \theta_Z$. Pick any $(x,z), (x',z') \in R_2 \circ R_1$ and suppose that x, x' belong to the same block of $\theta_X(t)$. By the definition of $R_2 \circ R_1$, there exist $y, y' \in Y$ such that $(x,y), (x',y') \in R_1$ and $(y,z), (y',z') \in R_2$. Since $\theta_X(t) \leq_{R_1} \bigvee_{[t]^{\varepsilon_1}} \theta_Y$, y, y' must be in the same block of $\bigvee_{[t]^{\varepsilon_1}} \theta_Y$. Recall that the sub-equivalence relation corresponding to the sub-partition $\bigvee_{[t]^{\varepsilon_1}} \theta_Y$ is the transitive closure of the relation $\cup_{s \in [t]^{\varepsilon_1}} \sim_s \subset Y \times Y$, where $\sim_s$ is the sub-equivalence relation (on Y) corresponding to the sub-partition $\theta_Y(s)$ of Y. In particular, the set $\{\sim_s : s \in [t]^{\varepsilon_1}\}$ consists of finitely many relations on Y due to the tameness of θ_Y. Therefore, there exist (finite) sequences $y = y_0, y_1, \ldots, y_n = y'$ in Y and $s_0, s_1, \ldots, s_{n-1}$ in $[t]^{\varepsilon_1}$ such that y_i, y_{i+1} belong to the same block of $\theta_Y(s_i)$ for $i = 0, \ldots, n-1$. Since R_2 is a correspondence between Y and Z, there exists a sequence $z = z_0, \ldots, z_n = z'$ in Z such that $(y_i, z_i) \in R_2$ for $i = 0, \ldots n$. Also, since $\theta_Y(s_i) \leq_{R_2} \bigvee_{[s_i]^{\varepsilon_2}} \theta_Z$, z_i, z_{i+1} belong to the same block of $\bigvee_{[s_i]^{\varepsilon_2}} \theta_Z$ for each i. Since $s_i \in [t]^{\varepsilon_1}$, we have $[s_i]^{\varepsilon_2} \subset [t]^{\varepsilon_1+\varepsilon_2}$, and in turn $\bigvee_{[s_i]^{\varepsilon_2}} \theta_Z \leq \bigvee_{[t]^{\varepsilon_1+\varepsilon_2}} \theta_Z$. Therefore, z_i, z_{i+1} belong to the same block of $\bigvee_{[t]^{\varepsilon_1+\varepsilon_2}} \theta_Z$ for each i, and hence z, z' belong to the same block of $\bigvee_{[t]^{\varepsilon_1+\varepsilon_2}} \theta_Z$. Similarly, one can verify that $\theta_Z(t) \leq_{(R_2 \circ R_1)^{-1}} \bigvee_{[t]^{\varepsilon_1+\varepsilon_2}} \theta_X$. □

Proof of Theorem 4.

Proof. We show that $\theta_X := \mathcal{C}_\delta(\gamma_X)$ satisfies the three conditions (tameness, interval lifespan, and comparability) to be a formigram. First, by the definition of $\mathcal{C}_\delta$, $\mathcal{C}_\delta(\gamma_X)$ is a function from $\mathbf{R}$ to the set of all partitions $\mathcal{P}(X)(\subset \mathcal{P}^{\mathrm{sub}}(X))$ of X. Therefore, every element $x \in X$ has the full lifespan $I_x = (-\infty, \infty)$, in θ_X.

Next we show the comparability condition. For simplicity, assume that $X = \{1, 2, \ldots, n\}$ for some $n \in \mathbf{N}$. Fix $c \in \mathbf{R}$ and consider the following two subsets of $X \times X$:

$$A(c,\delta) := \{(i,j) : \ i < j \in X, \ d_X(c)(i,j) \leq \delta\},$$
$$B(c,\delta) := \{(i,j) : \ i < j \in X, \ d_X(c)(i,j) > \delta\}.$$

The continuity of $d_X(\cdot)(i,j)$ for each $(i,j) \in X \times X$ guarantees that there exists $\varepsilon > 0$ such that

$$B(t,\delta) \supset B(c,\delta) \quad \text{for all } t \in (c-\varepsilon, c+\varepsilon)$$

and in turn

$$A(t,\delta) \subset A(c,\delta) \quad \text{for all } t \in (c-\varepsilon, c+\varepsilon)$$

since $A(t,\delta) \cup B(t,\delta) = \{(i,j) : i < j \in X\}$ for all $t \in \mathbf{R}$. This implies that the partition $\mathcal{C}_\delta(\gamma_X(c))$ is coarser than or equal to $\mathcal{C}_\delta(\gamma_X(t))$ for each $t \in (c-\varepsilon, c+\varepsilon)$, which means that $\mathcal{C}_\delta(\gamma_X)$ satisfies the comparability condition.

It remains to prove that $\mathcal{C}_\delta(\gamma_X)$ is tame. For $i, j \in X$, let $f_{i,j} := d_X(\cdot)(i,j) : \mathbf{R} \to \mathbf{R}_+$ and let $I \subset \mathbf{R}$ be any finite interval. Note that discontinuity points of the function $\mathcal{C}_\delta(\gamma_X) : \mathbf{R} \to \mathcal{P}(X)$ can occur only at endpoints of connected components of the set ${f_{i,j}}^{-1}(\delta)$ for some $i, j \in X$. Fix any $i, j \in X$. Then, since γ_X is tame, the set ${f_{i,j}}^{-1}(\delta) \cap I$ has only finitely many connected components and thus there are only finitely many endpoints arising from those components. Since the set X is finite, this implies that $\mathcal{C}_\delta(\gamma_X)$ can have only finitely many critical points in I. □

Appendix B

Isomorphic DMSs. We now introduce a notion of *equality* between two DMSs. Let $\gamma_X = (X, d_X(\cdot))$ and $\gamma_Y = (Y, d_Y(\cdot))$ be DMSs. We say that γ_X and γ_Y are *isomorphic* if there exists a bijection $\varphi : X \to Y$ such that φ is an isometry between $\gamma_X(t)$ and $\gamma_Y(t)$ across all $t \in \mathbf{R}$.

Theorem 6 *$d_{\mathrm{I}}^{\mathrm{dyn}}$ is an extended metric modulo isomorphisms between DMSs.*

We will prove Theorem 6 after showing Theorem 5.

Proof of Theorem 5.

Proof. First, note that for all $t \in \mathbf{R}$, X and Y are the underlying sets of $\theta_X(t)$ and $\theta_Y(t)$, respectively.

Let $\varepsilon \geq 0$ and assume that $R \subset X \times Y$ is any ε-correspondence between γ_X and γ_Y. It suffices to prove that R is an ε-correspondence between the formigrams θ_X and θ_Y

as well. Let $(x, y), (x', y') \in R$ and fix any $t \in \mathbf{R}$. Assume that x, x' belong to the same block of $\theta_X(t)$, meaning that there is a sequence $x = x_0, x_1, \ldots, x_n = x'$ in X such that $d_X(t)(x_i, x_{i+1}) \leq \delta$ for $0 \leq i \leq n-1$. For each $0 \leq i \leq n-1$, pick $y_i \in Y$ such that $(x_i, y_i) \in R$ where $y = y_0$ and $y' = y_n$. Since R is an ε-correspondence between γ_X, γ_Y, we have $\min_{s \in [t]^\varepsilon} d_Y(s)(y_i, y_{i+1}) \leq d_X(t)(x_i, x_{i+1}) \leq \delta$. This implies that, for each i, there is $s_i \in [t]^\varepsilon$ such that $d_Y(s_i)(y_i, y_{i+1}) \leq \delta$ and in turn y_i, y_{i+1} are in the same block of $\theta_Y(s_i)$. Also for each i, since $s_i \in [t]^\varepsilon$, one has $\theta_Y(s_i) \leq \bigvee_{[t]^\varepsilon} \theta_Y$ and in turn y_i, y_{i+1} belong to the same block of $\bigvee_{[t]^\varepsilon} \theta_Y$. Therefore, we conclude that y, y' belong to the same block of $\bigvee_{[t]^\varepsilon} \theta_Y$. We have proved that $\theta_X(t) \leq_R \bigvee_{[t]^\varepsilon} \theta_Y$. Similarly, $\theta_Y(t) \leq_{R^{-1}} \bigvee_{[t]^\varepsilon} \theta_X$ can be shown, completing the proof. □

Proof of Theorem 6.

Proof. Reflexivity and symmetry of $d_{\mathrm{I}}^{\mathrm{dyn}}$ are clear so we shall show the triangle inequality only: that for all DMSs $\gamma_X, \gamma_Y, \gamma_Z$, one has $d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Z) \leq d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Y) + d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_Y, \gamma_Z)$. We assume that $d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Y)$ and $d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_Y, \gamma_Z)$ are finite because otherwise there is nothing to prove. Let $0 < \varepsilon_1, \varepsilon_2 < \infty$ ans suppose that there are an ε_1-correspondence $R_1 \subset X \times Y$ between γ_X and γ_Y and an ε_2-correspondence $R_2 \subset Y \times Z$ between γ_Y and γ_Z. Define the correspondence $R_2 \circ R_1$ between X and Z as follows:

$$R_2 \circ R_1 := \{(x, z) \in X \times Z : \exists y \in Y \text{ s.t.} (x, y) \in R_1 \text{ and } (y, z) \in R_2\}.$$

Pick any two pairs (x, z) and (x', z') in $R_2 \circ R_1$. Then, there are $y, y' \in Y$ such that $(x, y), (x', y') \in R_1$ and $(y, z), (y', z') \in R_2$. Then for all $t \in \mathbf{R}$, it holds that

$$\min_{s \in [t]^{\varepsilon_1 + \varepsilon_2}} d_Z(s)(z, z') \leq \min_{s \in [t]^{\varepsilon_1}} d_Y(s)(y, y') \leq d_X(t)(x, x'),$$
$$\min_{s \in [t]^{\varepsilon_2 + \varepsilon_1}} d_X(s)(x, x') \leq \min_{s \in [t]^{\varepsilon_2}} d_Y(s)(y, y') \leq d_Z(t)(z, z').$$

Therefore, $R_2 \circ R_1$ is an $(\varepsilon_1 + \varepsilon_2)$-correspondence between γ_X, γ_Z, implying that $d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Z) \leq d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Y) + d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_Y, \gamma_Z)$, as desired.

Now, we show that $d_{\mathrm{I}}^{\mathrm{dyn}}$ is not just an (extended) pseudo-metric but an (extended) metric. Assume that $d_{\mathrm{I}}^{\mathrm{dyn}}(\gamma_X, \gamma_Y) = 0$ for some DMSs γ_X, γ_Y. Since there exist only finitely many correspondences between X and Y, there must exist a correspondence $R \subset X \times Y$ such that for any $\varepsilon > 0$, R is an ε-correspondence between γ_X and γ_Y. We claim that this R is a 0-correspondence. To this end, we need the following:

Claim. Let $f : \mathbf{R} \to \mathbf{R}$ be a continuous map and $r, t \in \mathbf{R}$. Suppose that for every $\varepsilon > 0$, $\min_{s \in [t]^\varepsilon} f(s) \leq r$. Then $f(t) \leq r$.

Proof. [Proof of Claim] For each $k \in \mathbf{N}$, take any $s_k \in [t]^{1/k}$ such that $f(s_k) \leq r$. Then $(s_k)_{k \in \mathbf{N}}$ is a sequence in $f^{-1}(-\infty, r]$ converging to t. Since f is continuous, $f^{-1}(-\infty, r]$ is a closed set and thus t must belong to $f^{-1}(-\infty, r]$, i.e. $f(t) \leq r$, as desired. □

Remember that whenever $x, x' \in X$, $y, y' \in Y$ are fixed, the distance functions $d_X(\cdot)(x, x'), d_X(\cdot)(y, y') : \mathbf{R} \to \mathbf{R}_+$ are continuous. Since R is an ε-correspondence for any $\varepsilon > 0$, it follows that for any $(x, y), (x', y') \in R$, for any $\varepsilon > 0$, and for any $t \in \mathbf{R}$,

1. $\min_{s \in [t]^\varepsilon} d_X(s)(x, x') \leq d_Y(t)(y, y')$,
2. $\min_{s \in [t]^\varepsilon} d_Y(s)(y, y') \leq d_X(t)(x, x')$.

Thus by **Claim**, for all $(x, y), (x', y') \in R$ and all $t \in \mathbf{R}$ it holds that $d_Y(t)(y, y') = d_X(t)(x, x')$. In addition, invoking that there exist $t_0, t_0' \in \mathbf{R}$ such that $\gamma_X(t_0)$ and $\gamma_Y(t_0')$ are (standard) metric spaces by the definition of DMSs, the correspondence R must be the graph of a bijection between X and Y. This implies that γ_X and γ_Y are isomorphic DMSs, as desired. □