REVIEW

Distances Between Immersed Graphs: Metric Properties

Maike Buchin[1] · **Erin Chambers[2]** · **Pan Fang[3]** · **Brittany Terese Fasy[4]** ·
Ellen Gasparovic[5] · **Elizabeth Munch[6,7]** · **Carola Wenk[3]**

Abstract

Graphs in metric spaces appear in a wide range of data sets, and there is a large body of work focused on comparing, matching, or analyzing collections of graphs in different ambient spaces. In this survey, we provide an overview of a diverse collection of distance measures that can be defined on the set of finite graphs immersed (and in some cases, embedded) in a metric space. For each of the distance measures, we recall their definitions and investigate which of the properties of a metric they satisfy. Furthermore we compare the distance measures based on these properties and discuss their computational complexity.

1 Introduction

In this survey, we provide a methodical overview of a collection of distance measures aimed at comparing graphs immersed and embedded in an ambient metric space, such as Euclidean space. Data in this form arises in a wide range of areas, including GIS, trajectory analysis, protein alignment, plant morphology, and commodity networks such as electrical grids. Intuitively comparing two networks might require a mapping or correspondence between the networks. For example, if one has a ground truth of a road network and a simplification or reconstruction of the same network, one may wish to measure the error of the latter. In this case, a mapping between the two networks would identify the parts of the ground truth that are successfully reconstructed/simplified and would enable one to study the local error of the reconstruction. In another example, two embedded graphs could serve as representations of a geographic network (e.g., rivers) in two different years, and mappings between them allow one to measure where and how much the networks have changed. On the other hand, many networks are not isomorphic, nor is one necessarily interested in true subgraph isomorphism.

✉ Maike Buchin
  maike.buchin@rub.de

Extended author information available on the last page of the article

Given the prevalence of immersed and embedded graphs, as well as the wide range of potential domain questions, we need mathematical foundations for comparing and measuring the resemblance of such structures. Moreover, such data is often collected from a noisy and error-prone process, leading to a need for a wide variety of distances that are robust to different types of issues in the data. Such measures are useful in machine learning and statistical approaches, where an understanding of the mathematical structure of the options gives stronger guarantees and theoretical analyses.

More formally, in this paper, we consider the set $\mathcal{G}_M$ of finite graphs immersed in a topological metric space (M, δ) and the set $\widetilde{\mathcal{G}}_M$ of finite graphs embedded in M. The most common setting is when we compare planar graphs immersed or embedded in Euclidean space, so that $M = \mathbb{R}^2$ and δ is the Euclidean distance d_2. We note that immersions (where edges are allowed to cross) and embeddings (where the image in the space M may not have crossings) are both of interest; for example, a road network is often considered to be embedded, but, in fact, given overpasses and bridges, in $M = \mathbb{R}^2$ an immersion is the correct representation. Our core question is the following: given $G_1, G_2 \in \mathcal{G}_M$ (or $\in \widetilde{\mathcal{G}}_M$), how can we measure the distance between G_1 and G_2? Moreover, what metric properties does a given distance satisfy?

Goal.   The goal of this survey is to enumerate and to understand various metric properties of $\mathcal{G}_M$ and $\widetilde{\mathcal{G}}_M$ under different graph distance measures. We note that different distance measures capture different aspects of the given graphs, and we conjecture that a formal study of the mathematical properties strengthens our understanding of these tradeoffs. For example, some prioritize global similarity, while others focus on local matching; some are theoretically sound but difficult to compute. After briefly summarizing relevant definitions and background in Section 2, we discuss distances between immersed graphs in Section 3. Finally, we conclude in Section 4 by mentioning several natural open problems and possible future work in this area.

2 Preliminaries

We are motivated by the study of the spaces of graphs embedded in Euclidean space induced by different distances. In this work, we consider a more general setting: spaces of graphs immersed in a metric space (M, δ). (In fact, we consider spaces of pseudographs, allowing self-loops and multiple edges between pairs of vertices.) We begin by briefly defining some of the key concepts used in the remainder of this paper, but refer the reader to a topology text for full definitions [17, 24].

2.1 Distances and Metrics

A number of key properties are considered desirable for comparisons between sets. Here, we focus on the properties most relevant for graph comparison; see for example [9, 10] for more thorough surveys and discussion of distances in a range of more general settings. We note that there are additional studies of the geometry of graph spaces, e.g., [20].

Table 1 Types of distances

	Finiteness	Identity	Sym.	Sep.	Subadd.
Distance	?	?	?	?	?
Metric	Y	Y	Y	Y	Y
Pseudo-metric (PM)	Y	Y	Y	?	Y
Semi-metric (SM)	Y	Y	Y	Y	?
Quasi-metric (QM)	Y	Y	?	Y	Y
Extended metric	N	Y	Y	Y	Y
Directed PM	Y	Y	N	?	Y
Extended directed PM	N	Y	N	?	Y
Extended directed QM	N	Y	N	Y	Y

For each distance discussed in this paper, we note whether or not the distance satisfies Finiteness, Identity, Symmetry, Separability, and Subadditivity. A question mark indicates that it does not need to satisfy that property, but may

Definition 1 (Key Properties of Dissimilarity Functions) Let $\mathbb{X}$ be a set. Consider a function $d \colon \mathbb{X} \times \mathbb{X} \to \overline{\mathbb{R}}_{\geq 0}$, where $\overline{\mathbb{R}}_{\geq 0}$ denotes the non-negative extended real line, i.e., $\overline{\mathbb{R}}_{\geq 0} := \mathbb{R}_{\geq 0} \cup \infty$. We define the following properties:

1. Finiteness: for all $x, y \in \mathbb{X}$, $d(x, y) < \infty$.
2. Identity: for all $x \in \mathbb{X}$, $d(x, x) = 0$.
3. Symmetry: for all $x, y \in \mathbb{X}$, $d(x, y) = d(y, x)$.
4. Separability: for all $x, y \in \mathbb{X}$, $d(x, y) = 0$ implies $x = y$.
5. Subadditivity (Triangle Ineq.): for all $x, y, z \in \mathbb{X}$, $d(x, y) \leq d(x, z) + d(z, y)$.

We say that d is a **distance metric** (or simply a **metric**) if it satisfies all five properties. However, less strict notions of dissimilarity can be defined. We say that d is a **pseudo-metric** if it satisfies finiteness, identity, symmetry, and subadditivity, but not necessarily separability.[1]

We say that d is a **semi-metric** if it satisfies finiteness, identity, symmetry, and separability, but not necessarily subadditivity. We say that d is a **quasi-metric** if it satisfies finiteness, identity, separability, and subadditivity, but not necessarily symmetry; see Table 1. If we do not know which, if any, of these properties d satisfies, then we simply refer to d as a **distance**.

When we know that certain properties are not satisfied, we can add certain adjectives. We refer to distance functions that allow infinite distances as **extended** distances.

We say that d is **asymmetric** or **directed**[2] if it does not satisfy symmetry (and may or may not satisfy identity, separability, and subadditivity). For example, an extended directed quasi-metric is a distance that satisfies identity, separability, and subadditivity, but does not satisfy finiteness and symmetry.

---

[1] Note that the terminology for different comparison functions is not standard across fields. We choose to use the terms that are used in topology and computational geometry (e.g., [9]).

[2] The term *directed distance* is sometimes used to refer to a distance function that allows for the range to include negative numbers. This is not the sense in which we are using this term. Throughout this paper, our distances are always nonnegative.

Remark 1 (Distances Between Sets) Let $\mathbb{X}$ be a set and let $d\colon \mathbb{X} \times \mathbb{X} \to \overline{\mathbb{R}}_{\geq 0}$ be a distance. Let $A, B \subseteq \mathbb{X}$. Then, we define $d(A, B) := \inf\limits_{(x,y) \in A \times B} d(x, y)$.

Remark 2 (Metrizing Distances) If d is not a metric, we can often metrize it. For example, if d is asymmetric, then we can symmetrize it to define a symmetric distance $D\colon \mathbb{X} \times \mathbb{X} \to \overline{\mathbb{R}}_{\geq 0}$, where $D(x, y) = \max\{d(x, y), d(y, x)\}$. Note that: if d satisfies separability, then D is also separable. Likewise, finiteness, identity, and subadditivity of D follow from finiteness, identity, and subaditivity of d, respectively. Throughout this paper, this is how we symmetrize asymmetric distances, but we note that there are other options as well.

As a more complicated example, if d is a pseudo-metric, we can make it separable by defining equivalence classes. For each $x \in \mathbb{X}$, we define its equivalence class, denoted $[x]$, as the set of points $[x] := \{y \in \mathbb{X} \mid d(x, y) = 0\}$ and let $\mathbb{Y}$ be the set of all equivalence classes. Then, we define $D\colon \mathbb{Y} \times \mathbb{Y} \to \overline{\mathbb{R}}_{\geq 0}$ by $D([x], [y]) = \max\limits_{x' \in [x]} \max\limits_{y' \in [y]} d(x', y')$.

2.2 Correspondences and Matchings

Let (M, δ) be a metric space, and $A, B \subset M$. One way to define the distance between A and B is to "line up" the points in the two sets and to use the distance δ to measure "how well" the two sets are aligned. For this, we define correspondences and matchings.

A **correspondence** τ between A and B is simply a subset of $A \times B$. Often, it is convenient to think of τ both as a subset and as a function that can take subsets of A to subsets of B and vice versa. Specifically, for $A' \subseteq A$, we define $\tau(A') := \{b \in B \mid \exists a \in A'$ s.t. $(a, b) \in \tau\}$, and define $\tau^{-1}(B')$ similarly to be $\{a \in A \mid (a, b) \in \tau\}$. We call τ a **matching** if each element of $A \sqcup B$ appears in τ at most once. We call τ a **perfect matching** if τ induces a bijection between the two sets, $b\colon A \to B$ with $b(a) = \tau(\{a\})$.

2.3 Open Balls and the Metric Topology

Let $\mathbb{X}$ be a set and $d\colon \mathbb{X} \times \mathbb{X} \to \overline{\mathbb{R}}_{\geq 0}$ a distance. For each $r \geq 0$ and $x \in \mathbb{X}$, we define the **open ball** of radius r centered at x by taking all points in $\mathbb{X}$ whose distance to x is less than r:

$$\mathbb{B}_d(x, r) := \{y \in \mathbb{X} \mid d(x, y) < r\}.$$

In the case that d is a metric, this is called a metric ball; if d is a pseudo-metric, it is a pseudo-metric ball, and so on. The closure of an open ball is a **closed ball** and can be written: $\overline{\mathbb{B}}_d(x, r) = \{y \in \mathbb{X} \mid d(x, y) \leq r\}$.

We can use the set of all open balls in $\mathbb{X}$ under the distance d in order to define a topology on $\mathbb{X}$. This is called the **open ball topology** in general, and denoted $(\mathbb{X}, d)$. If d is a metric, the resulting topological space is called a **metric space**. If d is a pseudo-metric, then it is a **pseudo-metric space**, and so on. We say that $(\mathbb{X}, d)$ is a **metric measure space** if there exists a Borel measure μ on $(\mathbb{X}, d)$ such that all open

balls have positive measure (i.e., for all $x \in \mathbb{X}$ and $r > 0$, $\mu(\mathbb{B}_d(x, r)) > 0$); see [16, Ch. $3\frac{1}{2}$] for more details on metric measure spaces.

At times, we may find it convenient to assume that a topological space is compact. We say that a topological space T is **compact** if every open cover has a finite subcover. On compact sets, minimum and maximum are defined:

Lemma 1 (Minimum and Maximum Over a Compact Space) *If T is a compact topological space and $f \colon T \to \mathbb{R}$ is a continuous function, then $\sup f$ and $\inf f$ are attained by elements in T.*

See, e.g., [27, Corollary 13.18] for a proof. In particular, the previous lemma allows us to write $\max f = \sup f$ and $\min f = \inf f$.

2.4 The Space of Paths

Let $\mathbb{X}$ be a set and $d \colon \mathbb{X} \times \mathbb{X} \to \overline{\mathbb{R}}_{\geq 0}$ be a distance function. Using the open ball topology in $\mathbb{X}$ and the subspace topology[3] on the unit interval $I = [0, 1] \subset \mathbb{R}$, a **path** in $\mathbb{X}$ is a continuous map $\gamma \colon I \to \mathbb{X}$, and we shall refer to γ as a path from $\gamma(0)$ to $\gamma(1)$. If $\mathbb{X}$ is a metric space, then the length of γ can be defined, and we call γ **rectifiable** if its length is finite; see [26] for the formal definition of path lengths. A path is a **geodesic** if it is locally shortest, that is, if no local perturbation of γ results in a shorter path.

A **reparameterization** φ of the unit interval I is a continuous, non-decreasing, and bijective map $\varphi \colon I \to I$. Thus, we may also call φ an orientation-preserving homeomorphism on I (here, the homeomorphism is orientation-preserving since $\phi(0) = 0$ and $\phi(1) = 1$). Thus, two paths $\gamma_1, \gamma_2 \colon I \to \mathbb{X}$ are equivalent, up to orientation-preserving homeomorphism, if there exist two reparameterizations φ_1 and φ_2 such that $\gamma_1 \circ \varphi_1 = \gamma_2 \circ \varphi_2$. (Note: we can always force one of the reparameterizations to be the identity map). For $a, b \in \mathbb{X}$, we use $\mathrm{Paths}_{\mathbb{X}}(a, b)$ to denote the collection of all paths from a to b in $\mathbb{X}$, up to reparameterization. Then, for $A, B \subseteq \mathbb{X}$, we define

$$\mathrm{Paths}_{\mathbb{X}}(A, B) := \bigcup_{(a,b) \in A \times B} \mathrm{Paths}_{\mathbb{X}}(a, b).$$

Let (M, δ) be a metric space, and let $\Pi := \mathrm{Paths}_M(M, M)$. If we topologize Π with the compact-open topology, then Π is compact (and, by Lemma 1, we can take minima and maxima over Π).

2.4.1 The Hausdorff Distance (General Form)

Let (M, δ) be a metric space and let 2^M denote the collection of all subsets of M. Then, the **directed Hausdorff distance** $\overrightarrow{\delta}_H \colon 2^M \times 2^M \to \overline{\overline{\mathbb{R}}}_{\geq 0}$ in (M, δ) is defined by

$$\overrightarrow{\delta}_H(A, B) := \sup_{a \in A} \delta(a, B) = \sup_{a \in A} \inf_{b \in B} \delta(a, b), \tag{1}$$

---

[3] Here, the unit interval $I = [0, 1]$ is a (closed) subspace of $\mathbb{R}$ under the standard topology.
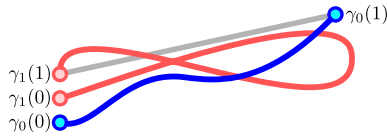
Fig. 1 Two curves γ_1 and γ_2 with small Hausdorff distance, as each point on one curve is quite near some point on the other. However, these two curves have large Fréchet distance, realized by the gray leash connecting $\gamma_1(1)$ to $\gamma_0(1)$

where the supremum over an empty set is defined to be zero and the infimum over an empty set is defined to be ∞. [4] See Fig. 1 We symmetrize this distance by taking the maximum (as in Remark 2) in order to obtain the **Hausdorff distance** $\delta_H : 2^M \times 2^M \to \overline{\mathbb{R}}_{\geq 0}$, which is defined by

$$\delta_H(A, B) := \max\{\overrightarrow{\delta}_H(A, B), \overrightarrow{\delta}_H(B, A)\}.$$

The directed Hausdorff distance is not a metric, as it does not satisfy finiteness, symmetry, nor separability (e.g., consider $M = \mathbb{R}$, $A = \{0\}$ and $B = \mathbb{Z}$ and observe that $\overrightarrow{\delta}_H(A, B) = \infty$ and $\overrightarrow{\delta}_H(B, A) = 0$). However, when restricted to the set of non-empty compact subspaces of M, the directed Hausdorff distance and the Hausdorff distance both satisfy finiteness.

Lemma 2 (The Hausdorff Distances) *Let (M, δ) be a metric space. The directed Hausdorff distance on M satisfies identity and subadditivity, but not finiteness, symmetry, and separability (i.e., it is an extended directed pseudo-metric), and the Hausdorff distance is an extended metric on M.*

Let $\mathbb{Y} \subset 2^M$ be a collection of non-empty, compact subspaces of M. Then, the directed Hausdorff distance when restricted to $\mathbb{Y}$ satisfies finiteness, identity, and subadditivity, but neither symmetry nor separability (i.e., it is a directed psuedo-metric on $\mathbb{Y}$), and the Hausdorff distance is a metric when restricted to $\mathbb{Y}$.

The proof of this lemma follows from the metric properties of δ; see [19] for details. For a detailed study of the Hausdorff distance, see [22].

2.4.2 Fréchet Distance Between Paths

We define the **Fréchet distance** $\delta_F : \Pi \times \Pi \to \overline{\mathbb{R}}_{\geq 0}$ by

$$\delta_F(\gamma_1, \gamma_2) = \inf_{\varphi} \max_{t \in [0,1]} \delta\left(\gamma_1(t), \gamma_2(\varphi(t))\right),$$

where φ ranges over all reparameterizations of I; see Fig. 2. The Fréchet distance is sometimes called the dog-walking distance, with the following intuition: if a woman is walking on one path, and her dog is walking on the other, find the shortest possible

---

[4] Felix Hausdorff introduced this distance in his 1914 book [18], where he restricted his attention to the space of closed, nonempty subsets of 2^M.
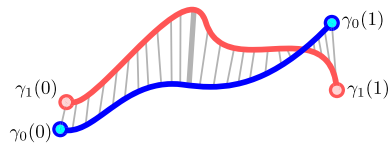
Fig. 2 Two curves γ_0 and γ_1, with the corresponding points on the curves shown connected in light grey. The Fréchet distance asks for the correspondence which minimizes the longest grey connection, shown here as a thicker line

leash length that allows them to walk on their curves from start to end. In a sense, the Fréchet distance is really a continuous perfect matching between the curves, where the cost of the matching is the worst case distance between points that are paired under the correspondence.

Similarly, the **weak Fréchet distance** $\delta_{wF} \colon \Pi \times \Pi \to \overline{\mathbb{R}}_{\geq 0}$ is defined using the same formula, but where ϕ is allowed to range over all continuous surjections (as opposed to bijections) such that $\phi(0) = 0$ and $\phi(1) = 1$. In other words, these maps do not need to be monotonic and are allowed to decrease.

While introduced originally by Fréchet in his thesis as one of the first examples of a metric [14], this distance was first investigated from a computational perspective by Alt and Godau [4], who also established the following lemma.

Lemma 3 (The Fréchet Metrics) *Let (M, δ) be a metric space, and let $\Pi =$ Paths$_M(M, M)$. Both the Fréchet distance δ_F and the weak Fréchet distance δ_{wF} are metrics on M.*

2.5 Immersions and Embeddings into Metric Spaces

Let (M, δ) be a metric space and $G = (V, E)$ be an abstract graph, which we view as a stratified one-dimensional simplicial complex. We restrict our attention to nontrivial graphs, so that each graph must contain at least one vertex. In this setting, note that when we write $x \in G$, the point x can either be a vertex in V_G or a point interior to an edge in E_G. We topologize G using the quotient space topology, where each (closed) edge is homeomorphic to the unit interval $I = [0, 1]$ in $\mathbb{R}^1$, and we topologize M using the open ball topology.

An **immersion** of G into M is a continuous map $\phi \colon G \to M$ such that for each point $x \in G$ and for all small enough open neighborhoods n_x of x in G, the map ϕ restricted to n_x (denoted $\phi|_{n_x}$) is a homeomorphism onto its image.[5] We call the pair (G, ϕ) an **immersed graph** in M. We say that (G, ϕ) is **rectifiable** if V and E are finite sets, and, for every edge $e \in E_G$, the length of $\phi(e)$ is finite. To simplify notation, we sometimes use G in place of the pair (G, ϕ).

---

[5] In differential geometry, immersions are required to be smooth maps. We do not require that in this paper.

(a) Immersion, Not Embedding   (b) Immersion and Embedding

Fig. 3 Two graphs with the same image in (M, δ). In (a), a graph with two vertices and one edge is immersed in M. Note that the edge self-intersects. In (b), a graph with three vertices and three edges is embedded in M with the same image as the graph in (a). In the embedding, the graph and its image in M are homeomorphic

Given two graphs (G_1, ϕ_1) and (G_2, ϕ_2) immersed in M, we say that they are equivalent **up to homeomorphism**, which we denote $(G_1, \phi_1) \cong (G_2, \phi_2)$, if there exists a homeomorphism $\alpha \colon G_1 \to G_2$ such that the following diagram commutes:

$$
\begin{array}{ccc}
G_1 & \xrightarrow{\quad \alpha \quad} & G_2 \\
& \searrow{\phi_1} \quad \phi_2 \swarrow & \\
& M &
\end{array}
$$

In other words, for all $x \in G$, $\phi_1(x) = (\phi_2 \circ \alpha)(x)$. The set of all rectifiable immersions of nontrivial graphs (up to homeomorphism) into M is denoted $\mathcal{G}_M$.

An immersion of G is said to be an **embedding** if G is homeomorphic onto the image $\phi(G)$. In other words, immersions allow edge crossings and embeddings do not; see Fig. 3 and [12, 21, 23]. The set of all rectifiable embeddings of nontrivial graphs (up to homeomorphism) into M is denoted $\widetilde{\mathcal{G}}_M$.

3 Comparisons on Immersed Graphs

Given a metric space (M, δ), we consider the collection $\mathcal{G}_M$ of finite graphs immersed in this metric space. Each $(G, \phi) \in \mathcal{G}_M$ is a graph $G = (V, E)$, where V is the vertex set and E is the edge set, together with an immersion $\phi \colon G \to M$. In this section, we investigate the metric properties of different distance functions on $\mathcal{G}_M$.

By default, we work on immersed graphs, where edges may cross each other. As every embedding is an immersion, the metric properties hold automatically when restricted to embedded graphs. However, when not all metric properties hold for immersed graphs, some additional properties may hold for embedded graphs that do not hold for the immersed graphs. We explicitly discuss this restriction and why it holds in the relevant sections.

3.1 Hausdorff Distances

In Sect. 2.4.1, we introduced the directed and undirected Hausdorff distances in the general setting, $\overrightarrow{\delta_H}$ and δ_H, respectively. We now consider $\overrightarrow{\delta_H}$ restricted to $\mathcal{G}_M$. We call this restriction the **directed Hausdorff distance between immersed graphs** $\overrightarrow{d_H} \colon \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$ and define it by

$$\overrightarrow{d_H}(G_1, G_2) = \overrightarrow{\delta_H}(\phi_1(G_1), \phi_2(G_2)),$$

and the **(undirected) Hausdorff distance** $d_H : \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$ is

$$d_H(G_1, G_2) = \delta_H(\phi_1(G_1), \phi_2(G_2))$$
$$= \max\{\overrightarrow{d_H}(G_1, G_2), \overrightarrow{d_H}(G_2, G_1)\}.$$

Since each graph in $\mathcal{G}_M$ is a compact subset of M, by Lemma 1, we know that the Hausdorff distance between two graphs in $\mathcal{G}_M$ is realized by points in the graphs (and hence the inf and sup of Equation (1) are actually min and max).

Theorem 1 *(Metric Properties of Hausdorff Distances Between Immersed Graphs) Let (M, δ) be a metric space. The following statements hold:*

1. *The directed Hausdorff distance $\overrightarrow{d_H}$ on $\mathcal{G}_M$ satisfies finiteness, identity, and sub-additivity, but not symmetry and separability (i.e., it is a directed pseudo-metric).*

2. *The Hausdorff distance d_H on $\mathcal{G}_M$ satisfies finiteness, identity, symmetry, and subadditivity, but does not satisfy separability (i.e., the Hausdorff distance on immersed graphs is a pseudo-metric).*

3. *The Hausdorff distance d_H restricted to embedded graphs in $\widetilde{\mathcal{G}}_M$ is a metric.*

Proof First, note that Statement 3 follows from Lemma 2 and the fact that each $G \in \widetilde{\mathcal{G}}_M$ is a compact subspace of M. Now, for Statements 1 and 1, $\overrightarrow{d_H}$ and d_H on $\mathcal{G}_M$ also satisfy nonnegativity, identity, and subadditivity. Neither $\overrightarrow{d_H}$ nor d_H satisfy separability, for two different immersed graphs may have the same image (e.g., a single edge mapped to a line segment and a Y-shaped graph mapped onto the same line segment; see Fig. 3), where two different graphs are mapped to the same subset of $\mathbb{R}^2$.

By construction, just as in the general case, $\overrightarrow{d_H}$ is asymmetric and d_H is symmetric; for instance, if the immersion of G_1 is a true subset of the immersion of G_2, then $\overrightarrow{d_H}(G_1, G_2) = 0 > \overrightarrow{d_H}(G_2, G_1)$.

3.2 Fréchet Distance

The Fréchet distance is a metric originally defined on paths in $\mathbb{R}^n$ (see Sect. 2.4.2), but it can also be defined for more general objects, such as oriented manifolds. Let $A, B \subseteq \mathbb{R}^d$ be two oriented manifolds and let $f : A \to M$ and $g : B \to M$ be two immersions.[6] Then the Fréchet distance between A and B is given by

$$\inf_{\alpha} \max_{t \in A} \delta(f(t), g(\alpha(t))),$$

where $\alpha : A \to B$ ranges over all orientation-preserving homeomorphisms. In fact, this definition can be generalized to any two homeomorphic spaces, such as graphs. However, care must be taken to either define an appropriate notion of orientation, or define the distance without requiring the homeomorphisms to be orientation-preserving.

---

[6] This definition can be generalized to continuous maps.

Fig. 4 Graphs with large Fréchet
but small Hausdorff distance

Here, the **Fréchet distance between immersed graphs** $d_F \colon \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$
can be defined by restricting our attention to the Fréchet distance between paths, and
avoiding any mention of orientation-preserving homeomorphims. In particular, for
$(G_1, \phi_1), (G_2, \phi_2) \in \mathcal{G}_M$, the Fréchet distance is defined as

$$d_F(G_1, G_2) = \min_{\alpha} \max_{e \in E_1} \delta_F(\phi_1(e), \phi_2(\alpha(e)))$$

if G_1, G_2 are homeomorphic, and ∞ if they are not.[7] Here, α ranges over all edge
mappings corresponding to the isomorphisms of G_1 and G_2, see [5], and we assume
the graphs have no degree two vertices. The latter can be assumed since we consider
graphs in $\mathcal{G}_M$ equivalent up to homeomorphism. Note that we are abusing notation
slightly here and viewing $\phi_G(e)$ as a parameterized curve, rather than just an immersion
of an edge; since the Fréchet distance considers all reparameterizations of the curve
regardless, any parameterization of the image of the edge is sufficient. Since graphs
are compared using an isomorphism, the Fréchet distance can be arbitrarily larger than
the Hausdorff distance, see Fig 4.

Theorem 2 (Metric Properties of Fréchet Distance Between Immersed Graphs) *The
Fréchet distance d_F is a metric.*

Proof It is well-known that the Fréchet distance is a pseudo-metric: Identity and sym-
metry follow directly from the definition.

Separability is also fulfilled: consider two graphs $(G_1, \phi_1), (G_2, \phi_2)$ with Fréchet
distance 0. These graphs must be isomorphic if their distance is 0, and hence G_1 is
the same as G_2.

For completeness, we provide a proof that the subadditivity is satisfied: Consider
three graphs $(G_1, \phi_1), (G_2, \phi_2)(G_3, \phi_3) \in \mathcal{G}_M$. Then

$$
\begin{aligned}
&d_F(G_1, G_2) + d_F(G_2, G_3) \\
&= \min_{\alpha_1 : G_1 \to G_2} \max_{e_1 \in E_1} d_F(\phi_1(e_1), \phi_2(\alpha_1(e_1))) \\
&\quad + \min_{\alpha_2 : G_2 \to G_3} \max_{e_2 \in E_2} d_F(\phi_2(e_2), \phi_3(\alpha_2(e_2))) \\
&\geq \min_{\alpha_1 : G_1 \to G_2} \min_{\alpha_2 : G_2 \to G_3} \Big(\max_{e_1 \in E_1} d_F(\phi_1(e_1), \phi_2(\alpha_1(e_1))) \\
&\quad + \max_{e_2 \in \alpha_1(e_1)} d_F(\phi_2(e_2), \phi_3(\alpha_2(e_2))) \Big) \\
&\geq \min_{\alpha_1 : G_1 \to G_2} \min_{\alpha_2 : G_2 \to G_3} \max_{e_1 \in E_1} d_F(\phi_1(e_1), \alpha_2(\phi_2(\alpha_1(e_1)))) \\
&= \min_{\alpha : G_1 \to G_3} \max_{e_1 \in E_1} d_F(\phi_1(e_1), \phi_3(\alpha(e_1))) = d_F(G_1, G_3)
\end{aligned}
$$

---

[7] Typically, this would be an infimum rather than a minimum, but since we are simply reducing to the
Fréchet distance between paths in Euclidean space, this infimum is realized.

The first inequality follows simply by combining the two terms, and the second inequality follows from the triangle inequality that is fulfilled by d_F for curves (which in turn follows from it being fulfilled by δ_F).

Note that if G_1 and G_2 are simple graphs, each isomorphism has a unique edge mapping. Hence computing this distance is at least as hard as determining graph isomorphism. For trees [5], it can be computed in polynomial time and for graphs of bounded tree width [7] it is fixed-parameter tractable. For planar embedded graphs, it is desirable that isomorphisms are "orientation-preserving" in the sense that they preserve orderings of edges around each vertex. This property can be used to enumerate all such planar orientation-preserving isomorphisms, and thus compute the distance for embedded graphs in polynomial time [13].

3.3 Path-Based Distance

The path-based distance was originally presented in [1]. This distance uses the Fréchet distance between paths in graphs (see Sect. 2.4.2) to define a distance between graphs.

Let (M, δ) be a metric space. For each $(G, \phi) \in \mathcal{G}_M$, let Π_G denote the set of all paths in G up to reparameterization; that is, $\Pi_G = \text{Paths}_G(G, G)$. See Sect. 2.4. The **directed path-based distance** $\overrightarrow{d_{path}} \colon \mathcal{G}_M \times \mathcal{G}_M \to \mathbb{R}_{\geq 0}$ is the directed Hausdorff distance between (G_1, ϕ_1) and (G_2, ϕ_2):

$$\overrightarrow{d_{path}}(G_1, G_2) = \max_{P \in \Pi_1} \min_{Q \in \Pi_2} \delta_F(\phi_1(P), \phi_2(Q)).$$

As noted in Remark 2, we can symmetrize this asymmetric distance in order to define the **path-based distance** $d_{path} \colon \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$ as the Hausdorff distance between path sets Π_1 and Π_2. Specifically,

$$d_{path}(G_1, G_2) = \max \left\{ \overrightarrow{d_{path}}(G_1, G_2), \overrightarrow{d_{path}}(G_2, G_1) \right\}.$$

Due to compactness of Π_1 and Π_2 and by Lemma 1, we can always find paths where this distance is realized. In other words, if $\overrightarrow{d_{path}}(G_1, G_2) = \varepsilon$, then there exist paths $P_1 \in \Pi_1$ and $P_2 \in \Pi_2$ such that the Fréchet distance between P_1 and P_2 is ε. See Fig. 5 for an example; note that the two graphs shown have infinite Fréchet distance, as they are not homeomorphic.

Theorem 3 *(Metric Properties of Path-Based Distances) The directed path-based distance satisfies finiteness, identity, and subadditivity, but not separability nor symmetry. The path-based distance is a metric.*

Proof For each $(G, \phi) \in \mathcal{G}_M$, define the set of paths $\Pi_G = \text{Paths}_G(G, G)$ and the set of immersed paths $\phi_G(\Pi_G) := \{\phi_G \circ f \mid f \in \Pi_G\}$. Since ϕ_G is an immersion and Π_G is compact, we know that $\phi_G(\Pi_G)$ is also compact. Hence, the theorem follows from setting $\mathbb{Y} = \{\phi_G(\Pi_G) \mid (G, \phi_G) \in \mathcal{G}_M\}$ in Lemma 2, where $\mathbb{Y}$ is a collection of compact subspaces of $\text{Paths}_M(M, M)$.
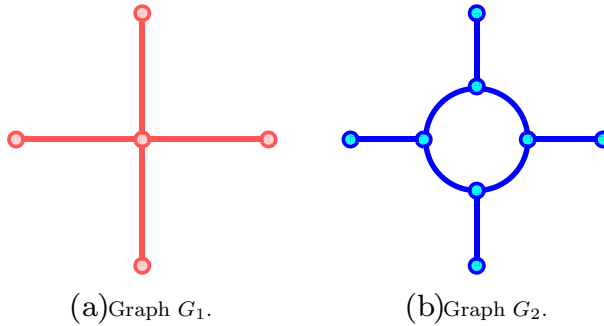
Fig. 5 Two graphs, G_1 and G_2, shown separately for ease of visibility. When they are embedded to be overlapping (so that the outer four vertices are in the same location in the plane), these graphs have small path based distance: each edge in G_1 will can map to two edges in G_2, and the path based distance will be less than the radius of G_2's inner loop

The exact complexity of the path-based distance is still an open and potentially challenging question, as the measure depends up Fréchet mappings between curves in the graph, of which there could be exponentially many to consider. However, in the paper that introduces it [1], the authors present a polynomial time approximation algorithm which is based on the maximum Fréchet distance and demonstrate its efficacy on real-world map datasets.

3.4 Traversal Distance

The *traversal distance*, introduced by Alt et al. in [3], for a pair of immersed graphs $(G_1, \phi_1), (G_2, \phi_2) \in \mathcal{G}_M$ is defined as follows. Let $f \colon [0, 1] \to G_1$ be a continuous and surjective function, called a (full) *traversal* of G_1, and let $g \colon [0, 1] \to G_2$ be a continuous—but not necessarily surjective—function, called a *partial traversal* of G_2. We compare functions f and g by their L_∞ norm, $\max_{t \in [0,1]} \delta(f(t), g(t))$. Taking the infimum over all possible f and g, we arrive at the **traversal distance** $\overrightarrow{d_T} \colon \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$. Formally, the traversal distance from (G_1, ϕ_1) to (G_2, ϕ_2) is defined by

$$\overrightarrow{d_T}(G_1, G_2) = \inf_{f,g} \max_{t \in [0,1]} \delta(f(t), g(t)),$$

where f ranges over all full traversals of G_1 and g ranges over all partial traversals of G_2. Noticing that reparametrizations of each traversal f and each partial traversal g are included in that infimum, we observe the following equivalence:

$$\overrightarrow{d_T}(G_1, G_2) = \inf_{f,g} \delta_F(f, g).$$

Compared to the Fréchet distance, the traversal distance can also be applied to non-homeomorphic graphs. On homeomorphic graphs, the Fréchet distance is an upper bound for the traversal distance, as the Fréchet correspondence yields a candidate
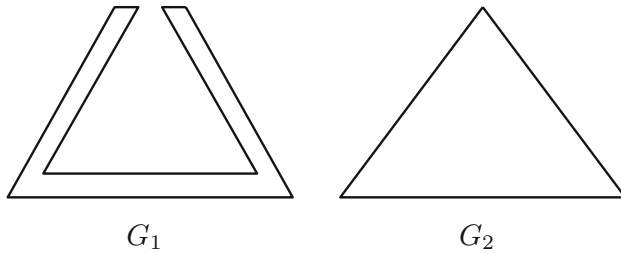
Fig. 6 Homeomorphic graphs with small traversal but large Fréchet and path-based distances
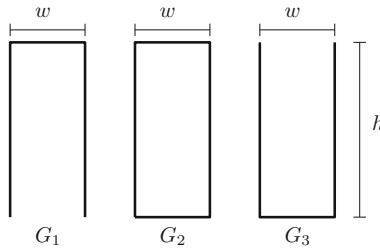
Fig. 7 Example showing that the traversal distance violates both separability and subadditivity. Assume that the four vertices of these graphs are the same, i.e., G_1 and G_3 are subgraphs of G_2. Since G_1 is a subgraph of G_2, the traversal distance from G_1 to G_2 is zero. To compute $\overrightarrow{d_T}(G_1, G_3)$, conside the traversal of G_1 that starts at the bottom left, goes up to the next vertex, right to the third vertex, and finally down to the last vertex. The best partial traversal of G_3 would be one that goes up and down one of the vertical edges

traversal to consider. However, the traversal distance can be arbitrarily smaller than the Fréchet distance, see Fig. 6.

We take the symmetric version of the traversal distance by maximizing the two directed distances. In other words, we define $d_T : \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$ by

$$d_T(G_1, G_2) = \max\{\overrightarrow{d_T}(G_1, G_2), \overrightarrow{d_T}(G_2, G_1)\}.$$

Theorem 4 (Metric Properties of Traversal Distance) *The directed traversal distance $\overrightarrow{d_T}$ satisfies finiteness and identity but does not satisfy symmetry, separability, and subadditivity.*

The symmetric traversal distance d_T satisfies finitesness, identity, and symmetry, but it does not satisfy separability nor subadditivity.

Proof Since the distance between f and g is finite and since the traversal distance is taking the infimum over all possible traversals f and partial traversals g, we know that $\overrightarrow{d_T}$ is finite. When $(G_1, \phi_1) = (G_2, \phi_2)$, taking the same traversal yields a traversal distance of zero. Hence, the traversal distance satisfies the identity property. Since any proper subgraph has distance zero to its supergraph, but not vice versa, we also have that it is not symmetric.

To see that the traversal distance does not fulfill separability and subadditivity, consider the following graphs: G_2 is comprised of four vertices and four edges forming an axis-aligned rectangle in $\mathbb{R}^2$, G_1 is the subgraph of G_2 obtained by
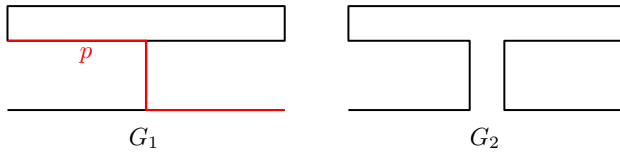
Fig. 8 The path p highlighted red in (G_1, ϕ_1) has large distance to any path in (G_2, ϕ_2), but the traversal distance from (G_1, ϕ_1) to (G_2, ϕ_2) is small

removing the bottom edge, and G_3 is the subgraph of G_2 obtained by removing the top edge. See Fig. 7. Then, $\overrightarrow{d_T}(G_1, G_2) = 0$, but (G_1, ϕ_1) and (G_2, ϕ_2) are not equivalent (up to homeomorphism). Thus, separability is not satisfied. In addition, we have $\overrightarrow{d_T}(G_2, G_3) = w/2$ and $\overrightarrow{d_T}(G_1, G_3) = w$. Since $w > w/2 + 0$, subadditivity is not satisfied.

Since $\overrightarrow{d_T}$ satisfies finiteness and symmetry and identity, so does the symmetrized version d_T. By construction, d_T is also symmetric. The two graphs of Fig. 3 have distance zero, which means that d_T does not satisfy separability. However, for d_T, subadditivity does not hold. For this, consider again the graphs in Fig. 7. We have $d_T(G_1, G_3) = w$. If we add small outward left and right spikes (say, they are length ϵ with $\epsilon < \frac{w}{2}$) on all graphs at height $h/2$ (calling the resulting graphs G_1^*, G_2^*, and G_3^*, respectively), then we have $d_T(G_1, G_2) = d_T(G_2, G_3) = w/2$ and $d_T(G_1^*, G_3^*) = w + \epsilon$. Hence, subadditivity does not hold.

Since a (partial) traversal is also a path in the graph, the traversal distance is related to the path-based distance. Both take the Fréchet distance from a traversal/path in (G_1, ϕ_1) to a closest partial traversal/path in (G_2, ϕ_2). However, the traversal distance minimizes over all traversals, whereas the path-based distance maximizes over all paths. Hence, the traversal distance is a lower bound to the path-based distance, i.e.,

$$d_T(G_2, G_1) \leq \overrightarrow{d_{path}}(G_2, G_1).$$

Fig. 8 gives an example where the path-based distance is strictly larger than the traversal distance.

3.5 Strong and Weak Graph Distance

The strong and weak graph distances were first introduced in [2], with the goal of combining topology and geometry. They are based on the strong/normal and weak Fréchet distance between graphs.

Let $(G_1, \phi_1), (G_2, \phi_2) \in \mathcal{G}_M$. A **graph mapping** $s \colon G_1 \to G_2$ is a continuous map. An alternative characterization of a graph mapping for planar embedded graphs is as follows: each vertex $v \in V_1$ is sent to a point $s(u) \in G_2$, and each edge $\{u, v\} \in E_1$ is sent to a path[8] from $s(u)$ to $s(v)$ in G_2. Note that $s(u)$ can be a vertex or any point internal to an edge.

---

[8] In [10] a graph mapping is required to map each edge to a *simple* path, and it is shown that for computing the distances it suffices to check such simple graph mappings.
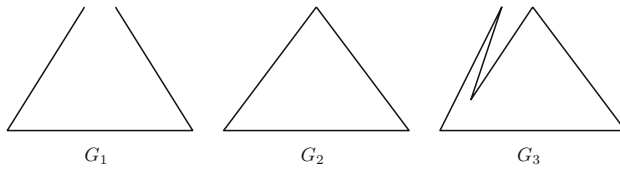
Fig. 9 Graph G_2 has strictly larger (strong or weak) graph distance than traversal distance to G_1. And graph G_2 has strictly larger strong than weak graph distance to G_3

Letting s range over all graph mappings from G_1 to G_2, the **directed strong graph distance** $\overrightarrow{d_S} : \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$ between G_1 and G_2 is given by

$$\overrightarrow{d_S}(G_1, G_2) = \inf_{s:G_1 \to G_2} \max_{e \in E_1} \delta_F(\phi_1(e), \phi_2(s(e))),$$

and the **directed weak graph distance** $\overrightarrow{d_W} : \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$ is given by

$$\overrightarrow{d_W}(G_1, G_2) = \inf_{s:G_1 \to G_2} \max_{e \in E_1} \delta_{wF}(\phi_1(e), (\phi_2(s(e))).$$

Note that these distances are not symmetric. However, we may define their undirected versions by taking the maximum of the directed distances, i.e.,

$$d_S(G_1, G_2) = \max\{\overrightarrow{d_S}(G_1, G_2), \overrightarrow{d_S}(G_2, G_1)\}$$

and

$$d_W(G_1, G_2) = \max\{\overrightarrow{d_W}(G_1, G_2), \overrightarrow{d_W}(G_2, G_1)\}.$$

In [2] it was shown that these distances are metrics for planar embedded graphs and pseudo-metrics for non-planar graphs. However, for immersed graphs, separability also holds for graphs when defining graph mappings as continuous maps between the (abstract) graphs. Hence we obtain:

Theorem 5 (*Metric Properties of Strong and Weak Graph Distances*) *The directed strong and weak graph distances* ($\overrightarrow{d_S}$ and $\overrightarrow{d_W}$) *satisfy finiteness, identity, separability, and subadditivity, but not symmetry (i.e., they are quasi-metrics). The undirected strong and weak graph distances* (d_S and d_W, respectively) *are metrics.*

Furthermore, in [2] the authors point out that

$$d_T(G_1, G_2) \leq \overrightarrow{d_W}(G_1, G_2) \leq \overrightarrow{d_S}(G_1, G_2),$$

which connects this distance to the traversal distance discussed in Sect. 3.4. See also Fig. 9.

Both the strong and weak graph distances are NP-hard to decide for general graphs. However, for trees we can compute them in cubic time for the strong graph distance

and quadratic time for the weak graph distance. For planar embedded graphs under a geometric assumption (which requires cycles to have a nice shape), the weak distance can be computed in quadratic time, but the strong distance remains NP-hard. An open question is whether a "stronger/symmetric" version of the graph distance, which requires the same graph mapping for both directions, would be equivalent to a generalization of the contour tree distance; see Sect. 3.6 for a discussion of that distance.

3.6 Contour Tree Distance

In [6], motivated by computing the Fréchet distance between two surfaces, the "contour tree distance" is defined between the contour trees of two surfaces. We naturally extend this distance to a distance on the subset of $\mathcal{G}_M$ consisting of connected (immersed) graphs; let $\mathcal{C}_M$ denote this space. The **contour tree distance** $d_C : \mathcal{C}_M \times \mathcal{C}_M \to \mathbb{R}$ is defined to be

$$d_C(G_1, G_2) = \inf_{\tau} \sup_{(x,y)\in\tau} \delta(\phi_1(x), \phi_2(y)),$$

where τ ranges over the set of all correspondences between (G_1, ϕ_1) and (G_2, ϕ_2) such that:

1. τ is a connected subset of $G_1 \times G_2$.
2. For each $x \in G_1$, $\tau \cap (\{x\} \times G_2)$ is a non-empty, connected subset of G_2.
3. For each $y \in G_2$, $\tau \cap (G_1 \times \{y\})$ is a non-empty, connected subset of G_1.

The connectedness of the correspondence τ requires the graphs to be connected, too. This is the reason that we are restricting our attention to $\mathcal{C}_M$ as opposed to $\mathcal{G}_M$.

The contour tree distance resembles the Fréchet distance in that it establishes a correspondence between portions of the graphs. However, unlike standard Fréchet distance, the contour tree distance allows a comparison of non-homeomorphic graphs using $G_1 \times G_2$, in a manner similar to the Fréchet distance. It allows for "stretching" a region of the graph, as any vertex in one can correspond to a connected subregion in the other. See Fig. 10 for an illustration.

Theorem 6 (d_C *is a Metric) On the space* $\mathcal{C}_M$ *of connected graphs,* d_C *is a metric.*

Proof The distance fulfills identity via the trivial correspondence, and symmetry because the same correspondence works for $G_1 \times G_2$ and $G_2 \times G_1$.

For separability, assume that $d_C(G_1, G_2) = 0$. Then either a correspondence τ exists such that $d(\phi(x), \phi(y)) = 0$ for all $(x, y) \in \tau$ or there is a limit of correspondences such that $d(\phi(x), \phi(y)) < \varepsilon$ for arbitrary small ε and $(x, y) \in \tau$. In both cases it follows that G_1 and G_2 have the same immersion.

To see that it also satisfies subadditivity, we concatenate the correspondences and use subadditivity in (M, δ).

The contour tree distance is NP-complete to compute, even for trees [6]. It seems that the contour tree distance can be considered as a symmetric version of the strong
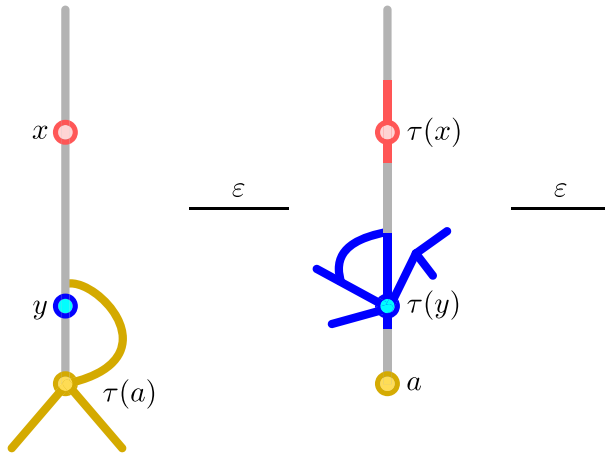
Fig. 10 Two graphs with small contour tree distance showing corresponding parts of the graphs

graph distance; see Sect. 3.5. Both align portions of the graphs, but where the strong graph distance uses two separate mappings between the two graphs, the contour tree distance uses symmetric correspondences.

3.7 Local Persistent Homology Distance

The next distance we investigate is the local persistent homology (LPH) distance, originally presented in [1] as a metric on plane graphs. To define this distance, we compare the graphs at a local level using *persistent homology* (see, e.g., [11]). Briefly, persistent homology is a multiscale version of the fundamental topological notion of homology, which measures the "features" in a space (i. e. connected components, holes, and higher-dimensional voids).

More formally: given a nested sequence of topological spaces (called a filtration), persistent homology tracks the appearance ("births") and disappearance ("deaths") of topological features within the filtration. The results are then encoded in a *persistence diagram* as pairs of (birth, death) points in the first quadrant of the plane. A standard distance between persistence diagrams is as follows: given persistence diagrams D_1 and D_2, their **bottleneck distance** d_b is defined to be

$$d_b(D_1, D_2) := \inf_{f : D_1 \to D_2} \sup_{x \in D_1} ||x - f(x)||$$

where f ranges over all bijections between D_1 and D_2.

The local persistent homology distance compares graphs at a local level and requires the following additional preliminary definition. Let $\mathbb{Y} \subseteq \mathbb{X}$ a set. We define an ε-**thickening** of Y, denoted $\mathbb{Y}^\varepsilon$, to be

$$\mathbb{Y}^\varepsilon := \bigcup_{x \in \mathbb{Y}} \mathbb{B}_d(x, \varepsilon)$$
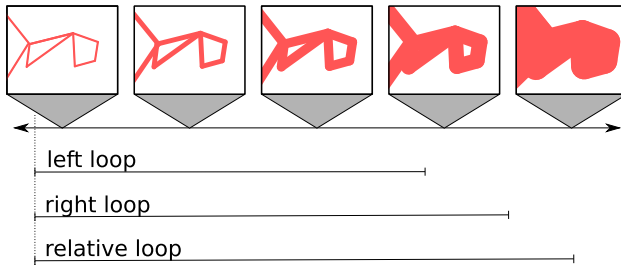
Fig. 11 A picture of a geometric graph embedded in $\mathbb{R}^2$, shown in the leftmost box, along with ε-thickenings for increasing values of ε. The graph contains 3 loops: two that are entirely in red graph, as well as one "relative" loop which is formed by the leftmost partial cycle of the graph along with its boundary. The birth and death times of each are indicated below the filtration; each such pair will form a point in the persistence diagram, with longer lifetime loops appearing further from the diagonal that has slope 1

$$= \{x \in \mathbb{X} \mid d(\mathbb{Y}, x) \le \varepsilon\}.$$

When $\mathbb{X}$ is a real vector space, $\mathbb{Y}^\varepsilon$ is referred to as the Minkowski sum of $\mathbb{Y}$ with a closed ball of radius ε in $(\mathbb{X}, d)$.

Let $(G, \phi_G) \in \mathcal{G}_M$. Then G^ε is the ε-thickening of G in (M, δ). We define the function $\delta_G : M \to \mathbb{R}$ as the distance function to the set $\phi_G(G)$; namely, for all $x \in M$, we define $\delta_G(x) := \min_{y \in \phi_G(G)} \delta(x, y)$. Equivalently, $\delta_G(x)$ is the smallest non-negative ϵ such that x is in the ϵ-thickening G^ϵ. Let U be a closed subset of M, and let ∂U denote the boundary of U. We consider the distance function restricted to the quotient space $U_* := U/\partial U$; that is, we define $\delta_{G,U} : U_* \to \mathbb{R}$ by

$$\delta_{G,U}([x]) = \begin{cases} \delta_G(x), & [x] \ne \partial U \\ \min_{y \in \partial U} \delta_G(y), & [x] = \partial U. \end{cases}$$

See Fig. 11 for an example.

Given graphs (G_1, ϕ_1) and (G_2, ϕ_2) in $\mathcal{G}_M$, let D_1 and D_2 be the persistence diagrams of $\delta_{G_1,U}$ and $\delta_{G_2,U}$, respectively. A *local distance signature* between G_1 and G_2 is the assignment of a distance to the neighborhood U; in our case, we assign the bottleneck distance between D_1 and D_2 to the set U, and denote it $B_U(G_1, G_2)$. The local distance signature is valuable in its own right and can be used in heatmaps to visualize the locations of large differences between the two graphs, as shown in Fig. 12.

We use the construction described above to define an overall distance between the two graphs.

Let $\mathcal{B}$ denote the set of all metric balls in (M, δ); that is,

$$\mathcal{B} = \{\mathbb{B}_\delta(x, r) \mid x \in M, r \in \mathbb{R}_{\ge 0}\}.$$

As a set, $\mathcal{B}$ is equivalent to the product $M \times \mathbb{R}_{\ge 0}$. And so, we use the product topology on $M \times \mathbb{R}_{\ge 0}$ to define a topology on $\mathcal{B}$.

Fig. 12 Comparison of two constructed road networks G_1 in color and G_2 in gray. The color-coded network visualizes the heatmap of the local distance signature for $x \in G_1$. Figure rights: ©2014 Ahmed, Fasy, and Wenk. First appeared in [1]
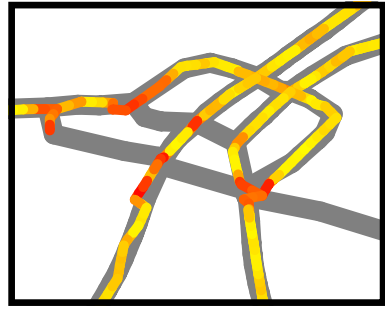
We call $\omega \colon \mathcal{B} \to \mathbb{R}_{\geq 0}$ a **weight function** if it: (i) vanishes for large enough balls (i.e., there is an $R \in \mathbb{R}$ such that $\omega(\mathbb{B}_\delta(x, r)) = 0$ for all $r > R$), (ii) has a finite integral ($\int_{U \in \mathcal{B}} \omega(U)\, dU < \infty$), and (iii) is continuous. The **local persistent homology distance with respect to the weight function** ω is the function $d_{LH} \colon \mathcal{G}_M \times \mathcal{G}_M \to \mathbb{R}$ defined by

$$d_{LH}(G_1, G_2) = \int_{U \in \mathcal{B}} \omega(U) B_U(G_1, G_2)\, dU. \tag{2}$$

If M is a metric measure space, this integral is well-defined. In particular, we can expand it to a double integral as follows:

$$d_{LH}(G_1, G_2) = \int_{r=0}^{\infty} \int_{x \in M} \omega(U) B_U(G_1, G_2)\, dx\, dr, \tag{3}$$

where $U = \mathbb{B}_\delta(x, r)$. Moreover, as we see in the next theorem, this distance is a pseudo-metric:

Theorem 7 *(Metric Properties of d_{LH}) For a given metric measure space M and weight function ω, the local homology distance d_{LH} satisfies finiteness, identity, symmetry, and subadditivity but does not satisfy separability (i.e., it is a pseudo-metric). The LPH distance restricted to comparing embedded graphs is a metric.*

Proof Let $\omega \colon \mathcal{B} \to \mathbb{R}_{\geq 0}$ be the weight function defining d_{LH}.

We prove that d_{LH} satisfies nonnegativity, finiteness, identity, symmetry, and subadditivity.

Nonnegativity. Let $G_1, G_2 \in \mathcal{G}_M$. The distance d_{LH} is defined in Equation (2) using an integral over $\mathcal{U}$ with integrand $\omega(U) B_U(G_1, G_2)$. Since ω is a weight function, we know that ω is nonnegative. As $B_U(G_1, G_2)$ is the bottleneck distance between two persistence diagrams, we know this term is also nonnegative. Thus, $d_{LH}(G_1, G_2)$, being the integral of the product of two nonnegative functions, is nonnegative.

Finiteness. Let $X \subset M$ be the support of ω. Since ω is a weight function, we know that there exists an R such that ω vanishes on all $U = \mathbb{B}_\delta(x, r) \in \mathcal{B}$ satisfying $R \leq r$, i.e., we have $\omega(U) = 0$. On the other hand, for all balls $U = \mathbb{B}_\delta(x, r) \in \mathcal{B}$ such that $r \leq R$, we know that $B_U(G_1, G_2) \leq 2R$ (since thickening by $2R$ or more results in the empty topological space). Then, letting $c = \int_{U \in \mathcal{B}} \omega(U)\, dU$ and $\mathcal{B}'$ denote the

support of ω, we bound the expression of Equation (2):

$$d_{LH}(G_1, G_2) \le \int_{U \in \mathcal{B}'} \omega(U) R \, dU + \int_{U \in \mathcal{B} \setminus \mathcal{B}'} 0 \cdot B_U(G_1, G_2) \, dU = cR + 0 < \infty.$$

Identity. If $G_1 = G_2 \in \mathcal{G}_M$, then $B_U(G_1, G_2) = 0$ for all U, and so $d_{LH}(G_1, G_2) = 0$.

Symmetry. The symmetry of d_{LH} follows immediately from the symmetry of the bottleneck distance between persistence diagrams.

Subadditivity. Subadditivity follows, again, from the subadditivity of the bottleneck distance between persistence diagrams.

Inseparability. Let (G_1, ϕ_1) be the immersion of the graph with two vertices and one edge from Fig. 3(a) and let G_2 be the graph with three vertices and three edges from Fig. 3(b). Note that G_1 and G_2 are not homeomorphic, yet have the same image in M. Then, for all metric balls U in (M, δ), we have $B_U(G_1, G_2) = 0$. Thus, d_{LH} is not separable.

Remark 3 (Variants of the LPH Distance) When restricted to the setting where M is compact, then d_{LH} also satisfies separability if ω is nonvanishing on small enough balls. This was observed in [1] for M a compact subspace of $\mathbb{R}^n$.

The definition of the LPH distance given in Equation (2) integrates the bottleneck distance over the space of all metric balls $\mathcal{B}$. We can aggregate over $\mathcal{B}$ in other ways, such as using the L_p-norm $\left(\int_{U \in \mathcal{B}} \omega(U) B_U(G_1, G_2)^p \, dU \right)^{1/p}$ in place of the bottleneck distance. We can also replace the bottleneck distance B_U with the Wasserstein metric or erosion distance between persistence diagrams. Changing the distance in these ways, the same metric properties hold. In addition, there are many different descriptors that one could consider in place of the persistence diagram, such as the Euler characteristic curve of $\delta_{G,U}$ or a simpler descriptor such as the number of connected components in $G \cap U$. However, if we use a weaker invariant, we may lose additional metric properties.

3.8 Graph Edit Distances

We next consider graph edit distances, which take a completely different and more combinatorial view when comparing graphs. Introduced in the 1980s, edit distances for general graphs are a well-studied way to compare graphs [25]. Given associated weights or costs for graph operations (e.g., vertices or edges to be inserted or deleted), the edit distance is the infimum of the sum of edit costs over all sequences of edit operations needed to transform one graph into the other. Edit distances are drastically different than other notions covered so far, in that they do not minimize some maximum correspondence, but rather sum all costs, leading to much larger distances. We give a brief overview of two versions of geometric edit distances and their metric properties here, as studied in [8]; see [15] for a survey of the many variants and heuristics on broader classes of graphs.
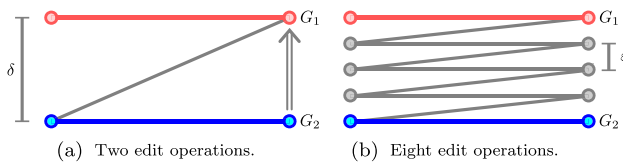
Fig. 13 Two examples of edit sequences between graphs G_1 and G_2. Both the edit distance and the geometric graph distance are $\geq 2\delta$, which Fréchet, Hausdorff, and path-based distances are all equal to δ on this example

A natural notion of edit distance for geometric graphs, which to the best of our knowledge was first discussed in [8], allows deleting and inserting vertices and edges for some cost related to the distance in the ambient space M, as well as moving vertices for cost propotional to the distance moved in M and change in edge lengths. More formally, the **cost** is then defined for each edit operation as follows:

1. *Edge Deletion and Insertion*. The cost of removing or inserting an edge is the length of the edge.
2. *Vertex Deletion and Insertion*. The cost of inserting or deleting an isolated vertex is 1. The cost of inserting a vertex in the middle of an edge to create two edges is free, as is the reverse operation. Otherwise, the cost of deleting a vertex is 1, plus the cost of deleting all incident edges.
3. *Vertex Moving*. The cost of moving a vertex is the distance that the vertex is moved, plus the sum of the changes in edge lengths of all incident edges.

The **graph edit distance** $d_{edit} : \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$ is defined by

$$d_{edit}(G_1, G_2) = \inf_{\tau} \sum_{\tau_i \in \tau} c(\tau_i), \tag{4}$$

where τ ranges over all finite sequences of edit operations that start at G_1 and end at G_2 and $c(\tau_i)$ denotes the cost of the edit operation τ_i.

In Fig. 13a, we see two graphs, G_1 and G_2, which each have two vertices and a single edge. The edit sequence between them is moving the right vertex first (shown as a grey double arrow), for a cost of δ to move the vertex itself plus the difference in length between the blue edge and the grey diagonal; the next edit move will shift the left vertex up, which again costs δ plus the difference in length between the diagonal and the red upper edge. This simple example already demonstrates the distinction of the edit distance and previous distances introduced in this survey, as the Fréchet, Hausdorff, and the path-based distance are all equal to δ on these graphs, while the edit distance must always be $\geq 2\delta$, given that each vertex must move δ.

We note that variations of edit distance could be examined, particularly with regards to different costs for vertex insertion or deletion costs; however, we begin with this definition as it is the one proposed and studied in the literature [8]. We begin by exploring metric proerties of graph edit distance, before discussing some of its limitations. The following properties hold for this distance:[9]

---

[9] While we do not believe these lemmas and theorems to be new results, finding exact references to these lemmas and theorems was not straightforward. So, for completeness, we provide proofs here.

Theorem 8 (Metric Properties of Graph Edit Distance) *The graph edit distance d_{edit} is a semi-metric that satisfies sub-additivity.*

Proof To see that d_{edit} is finite, consider the following weighted graph $\mathbb{G} = (\mathcal{G}_M, \mathcal{E}, \omega)$, where the vertex sets are all immersed graphs. A pair of immersed graphs $((G_1, \phi_1), (G_2, \phi_2))$ corresponds to an edge in $\mathcal{E}$ if and only if they are connected through a single edit operation. The weight of the edge, $\omega((G_1, \phi_1), (G_2, \phi_2))$, is the cost of that edit. Then, $d_{edit}((G_1, \phi_1), (G_2, \phi_2))$ is the length of the shortest path from (G_1, ϕ_1) to (G_2, ϕ_2) in $\mathbb{G}$. Since any graph is connected to the empty graph by a finite sequence of edit operations by removing all edges then removing all vertices, we know that the edit distance between two graphs is finite.

Identity follows because an empty sequence of edits transforms a graph to itself. Separability follows because, given $(G_1, \phi_1) \not\cong (G_2, \phi_2) \in \mathcal{G}_M$, we know that there must be some edit operation of positive cost to convert (G_1, ϕ_1) to (G_2, ϕ_2).

To prove symmetry, consider a fixed finite sequence τ in Equation (4). Let $-\tau$ denote the reverse sequence. Since any finite sequence of edits can be reversed to transform one graph to another, we know that $-\tau$ is one of the sequences considered by the infimum when defining $d_{edit}(G_2, G_1)$. Moreover, since costs of edit operations is symmetric, the cost of τ is the same as the cost of $-\tau$. Hence, we obtain:

$$d_{edit}(G_1, G_2) = \inf_{\tau} \sum_{\tau_i \in \tau} c(\tau_i) = \inf_{-\tau} \sum_{-\tau_i \in \tau} c(\tau_i) = d_{edit}(G_2, G_1).$$

Sub-additivity follows a similar argument. □

However, one major limitation, especially when considering algorithms to calculate this distance, is that the edit distance may never be attained by a finite sequence of edit operations. For instance, consider again our simple example of two graphs consisting of single straight edges, shown in Fig. 13. As pointed out by [8], an optimal edit sequence would be to alternate moving the vertices by an infinitesimal amount so as to minimize the change in edge length incurred; see Fig. 13b. Taking the limit as that infinitesmal amount decreases to zero, we get that the edit distance is exactly 2δ, but that value can never be realized by any finite sequence of edits.

In order to get around this issue, [8] introduces an edit-like distance, which they call the geometric graph distance, that makes a few notable changes to the edit distance above. Rather than charging a unit for a vertex addition or deletion, vertex additions and deletions are free. Also, they introduce two fixed parameters—an edge weight and a vertex weight—which factor multiplicatively into the relevant edit operations.[10] To avoid pathological examples of infinite edit sequences as discussed above, instead of charging costs for individual moves, they instead look at the total change in length at the end of all edits. They then observe that edits can be required to be done in the following order:

1. *Edge Deletion Phase*. The cost of removing an edge is the length of the edge times the edge weight.

---

[10] note that choosing both parameters as 1 results in the usual cost for deleting and inserting edges, and moving vertices

2. *Vertex Deletion Phase*. Deleting an isolated vertex is free. Non-isolated vertices may not be deleted.
3. *Vertex Moving Phase*. The cost of moving a vertex is the distance that the vertex is moved times the vertex weight, plus for all incident edges the change in edge lengths times the edge weight.
4. *Vertex Insertions Phase*. The cost of adding vertices is free. In contrast to the edit distance defined above, adding a vertex in the middle of an edge is not an allowable edit operation. (However, it can be attained by removing the edge in the Edge Deletion phase, then adding a vertex in this phase, and adding the two edges in the next phase).
5. *Edge Insertion Phase*. The cost of inserting an edge is the length of the edge times the edge weight.

Then, the *geometric graph distance* $d_{ggd} \colon \mathcal{G}_M \times \mathcal{G}_M \to \overline{\mathbb{R}}_{\geq 0}$ is defined by

$$d_{ggd}(G_1, G_2) = \inf_{\kappa} \sum_{\kappa_i \in \kappa} c(\kappa_i),$$

where κ ranges over all finite sequences of edit operations that start at G_1 and end at G_2 and preserve the phases described above. As proven in [8], this distance is a metric:

Theorem 9 (Metric Properties of Geometric Graph Distance) *The geometric graph distance d_{ggd} is a metric on the set of geometric graphs without isolated vertices for positive edge and vertex weight.*

However, as also shown in [8] the geometric graph distance is NP-complete to compute when considering non-planar graphs or choosing the edge weight much larger than the vertex weight. It is unknown if the problem remains NP-hard for planar graphs, or if the original edit distance formulation is also NP-hard.

4 Conclusion

In this paper, we survey commonly used distance measures on graphs immersed or embedded in a metric space. We also explored the metric properties induced on this space of graphs. For this, we studied the immersed graphs up to homeomorphism. If we did not consider graphs up to homeomorphism, then none of these distances would satisfy the separability property, since adding a vertex in the middle of an edge would result in a combinatorially different, yet homeomorphic, graph that has the same collection of immersions. In circumstances where the combinatorics are important, the proofs above can be adapted for studying immersions up to graph isomorphism instead of up to homeomorphism, in which case separability would no longer hold.

Many future areas of study remain in this area, so we conclude by mentioning one or two particularly natural questions. First, we are unsure if any of these metrics are equivalent to each other, where by "equivalent" we mean bounded within constant factors of each other. Another natural research path is to explore topological equivalences of the spaces of graphs as well as other topological properties of these spaces,

such as stability under certain types of perturbations. Finally, while we have briefly mentioned computational complexity of a few of these where results exist, the exact complexity or hardness of most remains open.

Declarations

Conflict of Interest The authors declare that they have no conflicts of interest.

Additional declarations for articles in life science journals that report the results of studies involving humans and/or animals Not applicable.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication If accepted, the publication of this manuscript is accepted by all coauthors.

Code availability Not applicable.

References

1. Ahmed, M., Fasy, B.T., Wenk, C.: Local persistent homology based distance between maps. In: Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, pp. 43–52. ACM (2014). https://doi.org/10.1145/2666310.2666390
2. Akitaya, H., Buchin, M., Kilgus, B., Sijben, S., Wenk, C.: Distance measures for embedded graphs. Comput. Geom. **95**, 101743 (2021)
3. Alt, H., Efrat, A., Rote, G., Wenk, C.: Matching planar maps. J. Algorithms **49**(2), 262–283 (2003). https://doi.org/10.1016/s0196-6774(03)00085-3

4. Alt, H., Godau, M.: Computing the Fréchet distance between two polygonal curves. IJCGA **5**(1–2), 75–91 (1995)
5. Buchin, K., Buchin, M., Schulz, A.: Fréchet distance of surfaces: some simple hard cases. In: European Symposium on Algorithms, pp. 63–74. Springer, New York (2010)
6. Buchin, K., Ophelders, T., Speckmann, B.: Computing the Fréchet distance between real-valued surfaces. In: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 2443—2455. ACM (2017)
7. Buchin, M., Krivosija, A., Neuhaus, A.: Computing the Fréchet distance of trees and graphs of bounded tree width. In: Proceedings of the 36th European Workshop on Computational Geometry (2020)
8. Cheong, O., Gudmundsson, J., Kim, H.S., Schymura, D., Stehn, F.: Measuring the similarity of geometric graphs. In: International Symposium on Experimental Algorithms, pp. 101–112. Springer (2009)
9. Conci, A., Kubrusly, C.: Distance between sets—a survey. Adv. Math. Sci. Appl. **26**, 1–18 (2017)
10. Deza, M.M., Deza, E.: Encyclopedia of Distances. Springer, Berlin (2009)
11. Edelsbrunner, H., Harer, J.: Computational Topology: An Introduction. American Mathematical Society, Washington, DC (2010)
12. Eppstein, D., Mumford, E.: Self-overlapping curves revisited. In: Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 160–169. SIAM (2009)
13. Fang, P., Wenk, C.: The Fréchet distance for plane graphs. In: Proceedings of the 37th European Workshop on Computational Geometry (2021)
14. Frechet, M.: Sur quelques points du calcul fonctionnel. Rendiconti del Circolo Matematico di Palermo (1884 - 1940) **22**(1), 1–72 (1906). https://doi.org/10.1007/BF03018603
15. Gao, X., Xiao, B., Tao, D., Li, X.: A survey of graph edit distance. Pattern Anal. Appl. **13**(1), 113–129 (2009). https://doi.org/10.1007/s10044-008-0141-y
16. Gromov, M.: Metric Structures for Riemannian and Non-Riemannian Spaces. Birkhäuser Boston (2007). https://doi.org/10.1007/978-0-8176-4583-0.Translation by S.M. Bates
17. Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge (2000). https://cds.cern.ch/record/478079
18. Hausdorff, F.: Grundzüge der Mengenlehre, vol. 7. Verlag von Veit & Comp, Leipzig (1914)
19. Hausdorff, F.: Set Theory. Chelsea Publishing Company, New York (1962)
20. Jain, B.J.: On the geometry of graph spaces. Discrete Appl. Math. **214**, 126–144 (2016). https://doi.org/10.1016/j.dam.2016.06.027
21. Li, Y., Barbič, J.: Immersion of self-intersecting solids and surfaces. ACM Trans. Graph. **37**(4), 1–14 (2018)
22. Michael, E.: Topologies on spaces of subsets. Trans. Am. Math. Soc. **71**, 152–182 (1951)
23. Mukherjee, U., Gopi, M., Rossignac, J.: Immersion and embedding of self-crossing loops. In: Proceedings of the Eighth Eurographics Symposium on Sketch-Based Interfaces and Modeling, p. 31-38. Association for Computing Machinery, New York, NY, USA (2011). https://doi.org/10.1145/2021164.2021170
24. Munkres, J.R.: Algebraic Topology. Prentice Hall, Upper Saddle River, NJ (1964)
25. Sanfeliu, A., Fu, K.S.: A distance measure between attributed relational graphs for pattern recognition. IEEE Trans. Syst. Man Cybernetics **13**(3), 353–362 (1983). https://doi.org/10.1109/TSMC.1983.6313167
26. Sullivan, J.M.: Curves of finite total curvature. In: Discrete Differential Geometry, pp. 137–161. Springer (2008)
27. Sutherland, W.A.: Introduction to Metric and Topological Spaces. Oxford University Press, Oxford (2009)

Authors and Affiliations

Maike Buchin[1] ⬥ · **Erin Chambers**[2] ⬥ · **Pan Fang**[3] · **Brittany Terese Fasy**[4] ⬥ ·
Ellen Gasparovic[5] ⬥ · **Elizabeth Munch**[6,7] ⬥ · **Carola Wenk**[3] ⬥

Erin Chambers
erin.chambers@slu.edu

Pan Fang
pfang@tulane.edu

Brittany Terese Fasy
brittany.fasy@montana.edu

Ellen Gasparovic
gasparoe@union.edu

Elizabeth Munch
muncheli@msu.edu

Carola Wenk
cwenk@tulane.edu

[1]   Department of Computer Science, Ruhr University Bochum, 44780 Bochum, Germany

[2]   Department of Computer Science, Saint Louis University, Saint Louis, MO, USA

[3]   Department of Computer Science, Tulane University, New Orleans, LA, USA

[4]   School of Computing & Department of Mathematical Sciences, Montana State University,
       Bozeman, MT, USA

[5]   Department of Mathematics, Union College, Schenectady, NY, USA

[6]   Department of Computational Mathematics, Science, and Engineering, Michigan State
       University, East Lansing, MI, USA

[7]   Department of Mathematics, Michigan State University, East Lansing, MI, USA