7-15-2020 2:00 PM

Locally Persistent Categories And Metric Properties Of Interleaving Distances

Luis N. Scoccola, *The University of Western Ontario*

Supervisor: Christensen, Daniel J., *The University of Western Ontario*
A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Mathematics

Abstract

This thesis presents a uniform treatment of different distances used in the applied topology literature. We introduce the notion of a locally persistent category, which is a category with a notion of approximate morphism that lets one define an interleaving distance on its collection of objects. The framework is based on a combination of enriched category theory and homotopy theory, and encompasses many well-known examples of interleaving distances, as well as weaker notions of distance, such as the homotopy interleaving distance and the Gromov–Hausdorff distance.

We show that the approach is not only an organizational tool, but a useful theoretical tool that allows one to formulate simple conditions under which a certain construction is stable, or under which an interleaving distance is, e.g., complete and geodesic. Being based on the well-developed theory of enriched categories, constructions in the theory of interleavings can be conveniently cast as enriched universal constructions.

We give several applications. We generalize Blumberg and Lesnick's homotopy interleaving distance to categories of persistent objects of a model category and prove that this distance is intrinsic and complete. We identify a universal property for the Gromov–Hausdorff distance that gives simple conditions under which an invariant of metric spaces is stable. We define a distance for persistent metric spaces, a generalization of filtered metric spaces, that specializes to known distances on filtered metric spaces and dynamic metric spaces, and use it to lift stability results for invariants of metric spaces to invariants of persistent metric spaces. We present a new stable invariant of metric measure spaces, the kernel density filtration, that encodes the information of a kernel density estimate for all choices of bandwidth. We study the interleaving distance in the category of persistent sets and show that, when restricted to a well-behaved subcategory that in particular contains all dendrograms and merge trees, one gets a complete and geodesic distance.

We relate our approach to previous categorical approaches by showing that categories of generalized persistence modules and categories with a flow give rise to locally persistent categories in a way that preserves both metric and categorical structure.

Keywords: Persistence, enriched category, extended pseudo metric, interleaving distance, quotient metric, weak equivalence.

Summary for lay audience

Algorithms in data science often require an input as well as a choice of parameters. In order to avoid arbitrary choices, one can study the evolution of the output of the algorithms as the parameters range over all possible choices.

In the context of applied topology, many algorithms first construct a representation of a topological space and then compute an invariant of this space. For example, many clustering algorithms work by computing the connected components of a graph that encodes some of the topology of the data set. When letting the parameters range over all possible choices, instead of constructing a single topological space, the algorithm constructs a persistent topological space, that is, a topological space parametrized by the poset of real numbers, and then computes an invariant of this persistent space, yielding a parametrized invariant. For example, the connected components of a topological space give a clustering of the space, while the connected components of a persistent topological space give a hierarchical clustering. Parametrized invariants are often stable, meaning that they are robust to small perturbations of the input dataset, making them a convenient practical tool. Parametrized invariants are studied by Topological Persistence.

It was observed in the work of Chazal, Cohen-Steiner, Glisse, Guibas, Oudot, Bubenik, Scott, Lesnick, and others that category theory can be used to organize and strengthen stability and consistency results about topological persistence methods. Categories are used to group mathematical objects with comparable structures together, such as the collection of all topological spaces.

This thesis studies a notion of category whose objects can be treated as persistent or parametrized objects. We show the benefits of this approach by recovering and generalizing previous results in the persistence literature in a uniform way, as well as giving new applications.

Co-authorship statement

Sections 6.5 and 6.6 contain joint work with Alex Rolle that appears in [RS20].

Acknowledgements

I thank my supervisor Dan Christensen for his encouragement and for many helpful and enlightening conversations about the topics of this thesis and about mathematics in general.

I thank Rick Jardine for many exchanges about the interactions between persistence and homotopy theory.

I thank Mike Lesnick for sharing his invaluable knowledge about multi-parameter persistence and applied topology.

I thank Alex Rolle for countless discussions about mathematics.

I thank Emily Riehl and Benedikt Fluhr for discussions about enriched category theory and persistence.

I thank Mark Daley, Rick Jardine, Chris Kapulkin, and Mike Lesnick for being part of the thesis examining committee and for helpful comments about the thesis.

I thank my family, my friends, and my partner for their unconditional support and their continuous encouragement during my PhD.

Contents

Chapter 1

Introduction

In this chapter, we motivate the use of persistence when doing topological inference and the use of category theory when studying persistence. We then describe the main contributions of the thesis and give an overview of related work.

1.1 Persistence: motivation and context

We start this section by giving two examples of how persistence can be used to consistently estimate topological features of continuous objects from finite samples. The first example is based on the persistence-based clustering algorithm introduced in [CGOS13]. The second example is about estimating the homology of a manifold from a sample, a problem for which many solutions have been proposed (see, e.g., [CL05; CSEH05; NSW08]).

We then show that category theory helps in formalizing and proving the consistency of the workflow of the examples, by providing us with distances between suitable relaxations of the topological invariants we wanted to estimate. We conclude by explaining what kind of metric properties are desirable for the distances category theory has provided us with.

Example: density-based clustering. Let $X \subseteq \mathbb{R}^d$ be a finite set of points that we want to cluster into disjoint groups. Many density-based clustering techniques assume that X was sampled from an unknown distribution given by a well-behaved probability density function $f : \mathbb{R}^d \to \mathbb{R}$. This assumption gives us something to work with: we can use f to specify a well-defined "true clustering" of the support of f. Given a density function f such as the one of Fig. 1.1, one could decide that there should be as many true clusters as there are modes (local maxima) of the distribution.

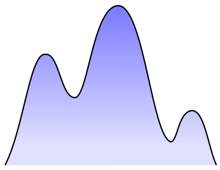

Figure 1.1: A density function $f : \mathbb{R} \to \mathbb{R}$.

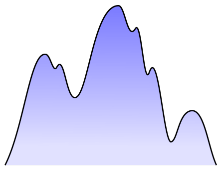

Figure 1.2: An estimate $\hat{f} : \mathbb{R} \to \mathbb{R}$ of f.

A problem one is faced with when estimating the modes of f from an estimate $\hat{f}$ is that one is often only guaranteed that $||f - \hat{f}||_\infty \leq \varepsilon$ for some small ε, so, even if $\hat{f}$ is well-behaved, it could have small, spurious local maxima, as in Fig. 1.2. To address this problem, many practical approaches to mode estimation construct a tree that tracks the evolution of the connected components of the superlevel sets of $\hat{f}$, then prune this tree and let the leaves of the pruned tree be the estimated modes ([SN10; CGOS13; KCBRW16; MH17]).

To formally define such a tree T for a function f, consider the superlevel sets of f, which we interpret as a functor $F : \mathbf{R}^{\mathrm{op}} \to \mathbf{Top}$ indexed by the poset $\mathbf{R}^{\mathrm{op}} = (\mathbb{R}, \geq)$:

$$F(r) = \left\{ x \in \mathbb{R}^d \,:\, f(x) \geq r \right\} \in \mathbf{Top}.$$

Composing F with the path components functor $\pi_0 : \mathbf{Top} \to \mathbf{Set}$, one obtains a persistent set $T : \mathbf{R}^{\mathrm{op}} \to \mathbf{Set}$, which can be represented as a merge tree, as in Fig. 1.3.

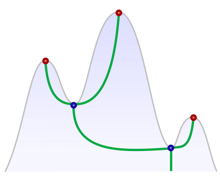

Figure 1.3: The tree T given by the connected components of $\{f \geq r\}$.

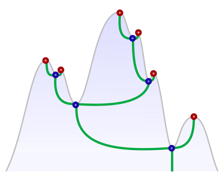

Figure 1.4: The tree $\hat{T}$ given by the connected components of $\{\hat{f} \geq r\}$.

Note that T has as many leaves as there are modes of f. The goal is to prune the estimated tree $\hat{T}$ to make it look like the true tree T. When pruning $\hat{T}$, it is useful to have a measure of prominence of modes. Since one is going to use this measure of prominence to prune all the modes that are not prominent enough, the measure must be *stable*, so that, in particular, the number of sufficiently prominent modes of an estimate $\hat{f}$ is the same as the number of sufficiently prominent modes of f, as long as $||f - \hat{f}||_\infty$ is sufficiently small.

Persistence theory provides us with a stable measure of prominence: the 0-th *persistence diagram* of a (sufficiently tame) function f, denoted by $\mathsf{PD}_0(f)$. The persistence diagram $\mathsf{PD}_0(f)$ is a collection of points above the diagonal of $\mathbb{R}^2$ that has exactly as many points as there are modes of the function f. One can use the vertical distance from a point in $\mathsf{PD}_0(\hat{f})$, which corresponds to a mode of $\hat{f}$, to the diagonal of $\mathbb{R}^2$ as a measure of the prominence of the corresponding mode of $\hat{f}$. This measure of prominence is stable in that there is a distance between persistence diagrams, the *bottleneck distance* d_B, such that the following holds

$$d_B\left(\mathsf{PD}_0\left(f\right), \mathsf{PD}_0\left(\hat{f}\right)\right) \leq ||f - \hat{f}||_\infty.$$

This is the 0th case of the celebrated stability theorem for persistence diagrams of tame functions, originally proven in [CSEH05] and [AFL03].

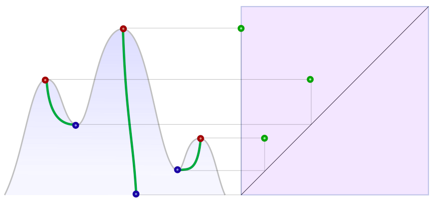

Figure 1.5: Constructing $\mathsf{PD}_0(f)$.

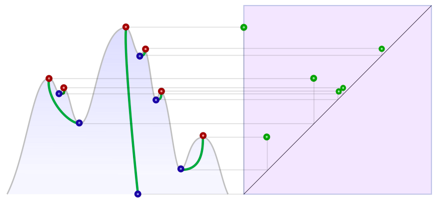

Figure 1.6: Constructing $\mathsf{PD}_0(\hat{f})$.

We see in Fig. 1.6 that, although $\hat{f}$ has more modes than f, it has exactly three prominent modes and three significantly less prominent ones, the ones close to the diagonal, that will have to pruned. We refer the interested reader to [CGOS13] where a clustering algorithm based on a more sophisticated version of these principles is introduced and proven consistent.

We haven't described how we constructed the persistence diagrams of Fig. 1.5 and Fig. 1.6. For the purposes of this introduction, it is enough to know that one can associate a persistence diagram to every (sufficiently tame) persistent vector space, that is, to a sufficiently tame (covariant or contravariant) functor $\mathbf{R} \to \mathbf{Vec}_k$. This perspective was first taken in [CFP01] and [ZC04], and exploited further in [CCSGGO09] and [CSGO16]. In the example above, the persistence diagram $\mathsf{PD}_0(f)$ represents the persistent vector space $\mathbf{R}^{\mathrm{op}} \to \mathbf{Vec}_k$ given by composing $T : \mathbf{R}^{\mathrm{op}} \to \mathbf{Set}$ with the free vector space functor $\mathbf{Set} \to \mathbf{Vec}_k$.

Example: homotopical invariants of discrete metric spaces. Suppose that we want to estimate the homology of a compact submanifold $M \subseteq \mathbb{R}^d$ from a finite sample $X \subseteq M$. As a topological space, X is discrete. One way around this problem is to choose a threshold $r \in \mathbf{R}$ and to construct a simplicial complex $\mathsf{VR}_\mathsf{c}(X)(r)$, the *Vietoris–Rips complex* of X at distance scale r, as follows. We let the vertex set of $\mathsf{VR}_\mathsf{c}(X)(r)$ be X, and we add an n-simplex $\{x_0,\ldots,x_n\}$ if and only if $d(x_i,x_j) \leq r$ for all $0 \leq i,j \leq n$. We can then geometrically realize this simplicial complex as a topological space which we denote by $\mathsf{VR}(X)(r) \in \mathbf{Top}$. In Fig. 1.7, we give an example of this construction for a set X that is a sample from a circular shape, for four different values of the threshold r. Note that there is a natural inclusion $\mathsf{VR}(X)(r) \to \mathsf{VR}(X)(r')$ whenever $r \leq r'$. This observation turns the Vietoris–Rips complex of X into a persistent space $\mathsf{VR}(X) : \mathbf{R} \to \mathbf{Top}$.

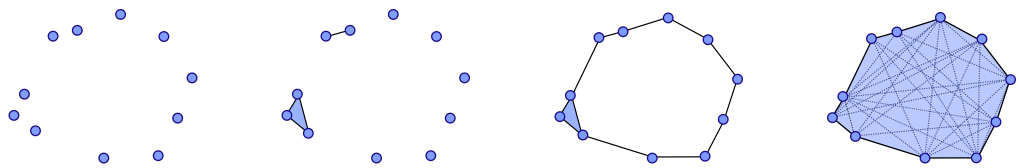

Figure 1.7: The Vietoris–Rips complex of a point cloud at four different stages, starting with $r = 0$ and ending with $r \gg 0$.

If X is a sufficiently good sample of M, one may expect that, for some suitable range of thresholds, the homology of $\mathsf{VR}(X)(r)$ will be a good approximation to the homology of M. Persistence lets us quantify this precisely. The following stability result is a consequence of [CCSGMO09, Theorem 3.1]: if there is $\varepsilon \geq 0$ such that every point of M is at distance at most ε from a point in X, then $d_B\left(\mathsf{PD}_n(X), \mathsf{PD}_n(M)\right) \leq 2\varepsilon$. Here $\mathsf{PD}_n(X)$ denotes the persistence diagram of the persistent vector space $\mathbf{R} \to \mathbf{Vec}_k$, given by composing $\mathsf{VR}(X) : \mathbf{R} \to \mathbf{Top}$ with the n-th homology functor $\mathbf{Top} \to \mathbf{Vec}_k$, and likewise for $\mathsf{PD}_n(M)$. This says that the persistent homology of $\mathsf{VR}(X)$ is a good approximation of the persistent homology of $\mathsf{VR}(M)$.

One then has to relate the persistent homology of $\mathsf{VR}(M)$ to the homology of M. For this, we refer the reader to [Hau95] and [Lat01], and to [KSCRW19] for state of the art results.

Categorification of persistence. The key point in the examples above was the stability of persistence diagrams. We say that a procedure is *stable* if it is uniformly continuous with respect to suitable metrics on its input set and on its output set. Let us now see that, when proving the stability of a procedure in a modular way, category theory comes in handy.

In Fig. 1.8 and Fig. 1.9 we depict the workflow of the first and second example above, respectively.

$$\boxed{\text{density functions}} \to \boxed{\text{persistent topological spaces}} \to \boxed{\text{persistent sets}} \to \boxed{\text{persistent vector spaces}} \to \boxed{\text{persistence diagrams}}$$

Figure 1.8: The workflow of the first example.

$$\boxed{\text{metric spaces}} \to \boxed{\text{persistent topological spaces}} \to \boxed{\text{persistent vector spaces}} \to \boxed{\text{persistence diagrams}}$$

Figure 1.9: The workflow of the second example.

In order to prove that these procedures are stable, one can prove that each of the steps is stable. This is the point of view advocated in [BSS13], [BS14], and [SMS18]. One of the main advantages of this point of view is modularity, that is, the ability to make local changes to the workflows and to combine different workflows without having to reprove the stability of the new workflow from scratch.

In order to prove that each step is stable, one needs a distance for each of the above collections of objects. Category theory gives us distances for all of the collections of objects above, and proofs that the mappings between them are stable.

The main categorical construction is the *interleaving distance*, first introduced for the category of persistent vector spaces in [CCSGGO09]. Note that the three intermediate collections of objects in Fig. 1.8 are the objects of the functor categories $\mathbf{Top}^{\mathbf{R}}$, $\mathbf{Set}^{\mathbf{R}}$, and $\mathbf{Vec}_k^{\mathbf{R}}$ respectively. In its most basic form, the interleaving distance is a metric that lets us compare objects of functor categories of the form $C^{\mathbf{R}}$.

Given a functor $F : \mathbf{R} \to C$ and $\varepsilon \geq 0$, let $F^\varepsilon : \mathbf{R} \to C$ denote the functor F shifted to the left by ε, that is, $F^\varepsilon(r) = F(r+\varepsilon)$ for every $r \in \mathbf{R}$. For $r \leq s \in \mathbf{R}$, let $\varphi^F_{r,s} : F(r) \to F(s)$ denote the structure map of F. We say that that two functors $F, G : \mathbf{R} \to C$ are ε-*interleaved*, if there exist natural transformations $\alpha : F \to G^\varepsilon$ and $\beta : G \to F^\varepsilon$ such that $\beta(r+\varepsilon) \circ \alpha(r) = \varphi^F_{r,r+2\varepsilon}$ and $\alpha(r+\varepsilon) \circ \beta(r) = \varphi^G_{r,r+2\varepsilon}$ for every $r \in \mathbf{R}$. The natural transformations α and β can be thought of as being ε-*approximate isomorphisms*, although the pair (α, β) is often called an ε-*interleaving*. One uses these interleavings to define the interleaving distance, as follows

$$d_I(F, G) = \inf\{\varepsilon \geq 0 : \ F \text{ and } G \text{ are } \varepsilon\text{-interleaved}\}.$$

As observed in [BS14], thanks to functoriality, if $F, G \in C^{\mathbf{R}}$ are ε-interleaved and $H : C \to D$ is any functor, then $H \circ F$ and $H \circ G$ are ε-interleaved as objects of $D^{\mathbf{R}}$. This proves that the intermediate steps in Fig. 1.8 are stable.

Going from persistent vector spaces to persistence diagrams is more subtle. The stability of this construction with respect to the interleaving distance on persistent vector spaces and the bottleneck distance on persistence diagram is known as the *algebraic stability theorem*, and was first introduced in [CCSGGO09]. A categorical proof of the algebraic stability result for pointwise finite-dimensional persistent vector spaces is presented in [BL14], and the fact that barcodes can be seen as the objects of a functor category of the form $C^{\mathbf{R}}$ appears in [EJM15] and [BL20].

To conclude that the workflow presented in Fig. 1.8 is stable, one has to show that the step going from density functions with the ∞-norm to persistent topological spaces with the interleaving distance is stable. This is straightforward, and also has a categorical proof which starts by giving an interpretation of the ∞-norm as an interleaving distance ([SMS18, Section 3.10.1], [Les15, Remark 5.1]).

One may expect the workflow in Fig. 1.9 to work analogously. One first has to choose a distance between metric spaces. The Gromov–Hausdorff distance (Definition 2.2.24) is usually chosen. This distance allows for a very general notion of similarity between metric spaces, and as a consequence it is not necessarily the case that if P and Q are Gromov–Hausdorff close, then $\mathsf{VR}(P)$ and $\mathsf{VR}(Q)$ are ε-interleaved for some small ε.

This issue can be resolved by weakening the notion of interleaving between persistent spaces, as done for filtered simplicial complexes in [Mé17] and [CSO14], and for persistent topological spaces in [BL17]. The solution in [BL17] is particularly natural from a homotopy-theoretic point of view: Blumberg and Lesnick let $X, Y : \mathbf{R} \to \mathbf{Top}$ be ε-*homotopy interleaved* if there exist weakly equivalent persistent topological spaces $X' \simeq X$ and $Y' \simeq Y$ such that X' and Y' are ε-interleaved. Using homotopy interleavings instead of interleavings, they define the *homotopy interleaving distance* d_{HI}, and prove ([BL17, Section 6]):

$$d_{HI}(\mathsf{VR}(P), \mathsf{VR}(Q)) \leq 2d_{GH}(P, Q).$$

The rest of the workflow in Fig. 1.9 is stable, since homology is a homotopy-invariant functor.

Geometry of spaces of persistent objects. When doing statistical inference on spaces of persistent objects, it is necessary to define probability measures on these spaces. It is thus desiderable to know that the space being studied is separable and complete. In practical applications, it is also useful to be able to interpolate between persistent objects, so one is interested in knowing if the space is intrinsic, and in having explicit formulas for constructing paths between points.

This kind of analysis for persistence diagrams has been done in, e.g., [MMH11], [TMMH12], and [FLRWBS14]. Nonetheless, it is often fruitful to study the geometry of spaces of objects other than persistence diagrams. For example, one may need to do statistics directly on spaces of trees, on spaces of multi-dimensional persistent vector spaces, or on spaces of persistent topological spaces, and thus a study of the geometry of these spaces is needed. See, for example, [Les12; BGMP14; KCBRW16; CO17; BSN17; BV18; GMOTWW19; Cru19].

1.2 Contributions

1.2.1 Setup

This thesis proposes an approach for defining distances between objects of a category and provides stability results for these distances, metric results for these distances (such them being complete or geodesic), and ways of combining distances between simple objects to get distances between more structured objects.

The approach is based on two notions of similarity: interleaving and weak equivalence. Interleavings are formalized using enriched category theory, while concepts from categorical homotopy theory are used to handle weak equivalences. One of the main selling points of the approach is that, by framing interleavings using the language of enriched categories, we give ourselves access to a very well developed set of formal tools for working with interleavings. Of particular interest are weighted (co)limits.

We introduce the notion of a *locally persistent category*. A locally persistent category is a category enriched in $\mathbf{Set}^{\mathbf{R}_+}$, where $\mathbf{R}_+ = ([0,\infty),\leq)$ is endowed with the monoidal product given by sum and $\mathbf{Set}^{\mathbf{R}_+}$ is endowed with the monoidal product given by Day convolution. A locally persistent category $\mathscr{C}$ is equivalently a category with extra structure: for each pair of objects $x, y \in \mathscr{C}$, instead of just having a set of morphisms from x to y, we have a persistent set $\mathrm{Hom}_{\mathscr{C}}(x, y) : \mathbf{R}_+ \to \mathbf{Set}$, indexed by the non-negative real numbers. For $x, y \in \mathscr{C}$ and $\varepsilon \in \mathbf{R}_+$, we think of the set $\mathrm{Hom}_{\mathscr{C}}(x, y)_\varepsilon$ as the set of ε-approximate morphisms from x to y. Composition is required to be compatible with this structure, meaning that the composite of an ε-approximate morphism with a δ-approximate morphism is an $(\varepsilon+\delta)$-approximate morphism. By copying the definition of isomorphism, but using approximate morphisms, one gets the notion of interleaving. One can then use interleavings to define an interleaving distance for any locally persistent category. We point out that this distance is really an extended pseudo distance (Definition 2.2.1).

As discussed in Section 1.1, it is sometimes necessary to consider a weaker notion of interleaving, in which we are allowed to replace objects by weakly equivalent ones. Following the methodology of categorical homotopy theory, we do this by considering locally persistent categories together with a class of 0-approximate morphisms that we declare to be *acyclic morphisms*. A *relative locally persistent category* consists of a locally persistent category together with a class of acyclic morphisms. We define a weak version of the interleaving distance, the *quotient interleaving distance*, using a metric quotient: the quotient interleaving distance is defined to be the greatest distance that is bounded above by the (strict) interleaving distance and is invariant under the equivalence relation given by being connected by acyclic morphisms. We point out again that the quotient interleaving distance is an extended pseudo distance.

As one expects, this framework encompasses standard interleaving distances such as the interleaving distance on any functor category of the form $C^{\mathbf{R}}$, and homotopical ones, such as the homotopy interleaving distance on $\mathbf{Top}^{\mathbf{R}}$. More interestingly, distances that at first may not look like interleaving distances arise as quotient interleaving distances. For example, the Gromov–Hausdorff distance and related distances on metric measure spaces are of this form.

1.2.2 Structure of the thesis and main results

In Chapter 2, we give background on metric spaces, enriched categories, and model categories, and study some of the properties of functor categories of the form $C^{\mathbf{R}}$ in detail. The contents of Chapter 2 are either well-known results, or variations of well-known results. The reader may skip Chapter 2 and refer to it when necessary.

In Chapter 3, we develop the category theory of locally persistent categories. We extend diagrammatic notation to this setting and study universal constructions that are particularly relevant when studying metric properties of an interleaving distance, such as weighted pullbacks, weighted sequential limits, and terminal midpoints. In Section 3.3, we give the definition of relative locally persistent category and of its associated quotient interleaving distance.

In Chapter 4, we study the metric properties of quotient interleaving distances. The results in this section are applied in several examples in Chapter 6. We prove the following stability result.

Theorem A (Theorem 4.2.2)**.** *A locally persistent functor between relative locally persistent categories that maps acyclic morphisms to acyclic morphisms is* 1*-Lipschitz with respect to the quotient interleaving distances.*

In order to state some consequences of the main results in Chapter 4 concisely, we introduce the following concepts. We say that an extended pseudo distance is: *complete* (Definition 2.2.9) if every Cauchy sequence has a limit, *intrinsic* (Definition 2.2.15) if the distance between two points at finite distance is the infimum of the lengths of paths between these two points, and *geodesic* (Definition 2.2.14) if the distance between two points at finite distance is equal to the length of some path between the points.

We say that a locally persistent category $\mathscr{C}$ is *powered by representables* (Definition 3.2.6) if, for every $y \in \mathscr{C}$ and $\varepsilon \in \mathbf{R}_+$, there exists $y^\varepsilon \in \mathscr{C}$ such that the persistent set $\mathrm{Hom}_{\mathscr{C}}(x, y)_{\varepsilon+(-)}$ is naturally isomorphic to $\mathrm{Hom}_{\mathscr{C}}(x, y^\varepsilon)$ for every $x \in \mathscr{C}$. For $\mathscr{C}$ a relative locally persistent category, we let d_{QI} denote its associated quotient interleaving distance and we let $\mathscr{C}_0$ denote the underlying category of $\mathscr{C}$, that is, the category whose objects are the objects of $\mathscr{C}$ and such that $\mathrm{Hom}_{\mathscr{C}_0}(x, y) = \mathrm{Hom}_{\mathscr{C}}(x, y)_0$. We fix $\mathscr{C}$ a relative locally persistent category that is powered by representables and such that powers preserve the acyclic morphisms and limits of $\mathscr{C}_0$. For $x, y \in \mathscr{C}$, we denote the fact that x and y are connected by a zig-zag of acyclic morphisms by $x \simeq y$.

We give a characterization of the quotient interleaving distance.

Theorem B (Theorem 4.1.4 using Lemma 4.1.5)**.** *If $\mathscr{C}_0$ admits pullbacks, and acyclic morphisms are stable under pullback, then*

$$d_{QI}(x, y) = \inf\{\varepsilon : \exists x' \simeq x, \exists y' \simeq y \text{ such that } x' \text{ and } y' \text{ are } \varepsilon\text{-interleaved}\}.$$

We give sufficient conditions under which d_{QI} is complete, which can be seen as a generalization of the completeness result in [Cru19] to quotient interleaving distances.

Theorem C (Theorem 4.3.3 using Proposition 3.2.15 and Lemma 4.1.5)**.** *If $\mathscr{C}_0$ admits pullbacks and sequential limits, and acyclic morphisms are stable under pullbacks and closed under sequential limits, then d_{QI} is complete.*

We give sufficient conditions under which d_{QI} is intrinsic, which generalizes the fact, proven in [CSGO16, Section 3.4], that interleaving distance on persistent vector spaces is intrinsic. We also give conditions under which d_{QI} is geodesic, but these are more technical.

Theorem D (Theorem 4.4.2 using Proposition 3.2.19)**.** *If $\mathscr{C}_0$ admits finite limits, then d_{QI} is intrinsic.*

We import some well-known constructions on categories of persistent objects to the theory of locally persistent categories. In Section 4.6 we generalize the interpolation framework of [BSN17] to locally persistent categories and show that the original framework factors through this generalization in a precise sense. We also show that the category of locally persistent categories contains the category of metric spaces as a full subcategory and that, in a rather trivial way, every metric arises as an interleaving distance. The value of locally persistent categories comes from having extra categorical structure that is compatible with the metric structure. In Section 4.7, we generalize the notion of observable category of [CCBS14] to locally persistent categories, and we relate the observable category of a category of persistent objects to a category of persistent objects satisfying a sheaf condition.

In Chapter 5, we give formal ways of constructing locally persistent categories. Of special interest is the construction of a locally persistent category of persistent objects $\mathscr{C}^{\mathbf{R}}$ for $\mathscr{C}$ a locally persistent category, described in Section 5.1.2. The underlying category of $\mathscr{C}^{\mathbf{R}}$ is the category of functors $\mathbf{R} \to \mathscr{C}_0$, but its locally persistent structure takes into account both the shifts of these functors and the pointwise locally persistent structure of $\mathscr{C}$. We later use this construction to define the Gromov–Hausdorff-interleaving distance on persistent metric spaces. We also relate categories with a flow [SMS18] to locally persistent categories. We show that to each category with a flow one can functorially associate a locally persistent category with the same objects and the same interleaving distance, thus showing that locally persistent categories can be seen as a generalization of categories with a flow. We argue that, although in most applications both frameworks apply, the language of locally persistent categories more closely matches the language of category theory, and avoids coherence arguments with 2-cells. We view categories with a flow as a streamlined way of constructing locally persistent categories.

In Chapter 6 we give applications of our main results. In Section 6.1, we extend the homotopy interleaving distance of [BL17] to persistent objects of a model category, and prove some metric properties of this distance.

Theorem E (Theorem 6.1.7)**.** *Let $\mathscr{M}$ be a cofibrantly generated model category. Then the quotient interleaving distance on the locally persistent category $\mathscr{M}^{\mathbf{R}}$ is intrinsic and complete, and satisfies*

$$d_{QI}(x,y) = \inf\{\delta \geq 0 : \exists x' \simeq x, \exists y' \simeq y,\ x' \text{ and } y' \text{ are } \delta\text{-interleaved}\}.$$

In Section 6.2, we show that the category of metric spaces, and more generally, the category of dissimilarity spaces, has the structure of a relative locally persistent category such that the quotient interleaving distance coincides with twice the Gromov–Hausdorff distance, and we recover well-known facts about the Gromov–Hausdorff distance. A dissimilarity space consists of a set X together with a function $X \times X \to [0,\infty]$. We use the characterization of the Gromov–Hausdorff distance as a quotient interleaving distance to prove the following stability result for invariants of metric spaces, which can also be interpreted as a universal property of the Gromov–Hausdorff distance. Let **epMet** denote the collection of extended pseudo metric spaces. Given $X_1 = (X, d_1)$ and $X_2 = (X, d_2)$ extended pseudo metric spaces with the same underlying set, let $d_\infty(X_1, X_2) = ||d_1 - d_2||_\infty$. This metric extends to an extended pseudo metric on **epMet** by declaring the distance between metric spaces with different underlying sets to be infinity.

Theorem F (Proposition 6.2.21)**.** *Let P be an extended pseudo metric space and let $V : \mathbf{epMet} \to P$ be a function. Assume that V is uniformly continuous (resp. 1-Lipschitz) with respect to d_∞ and the metric on P. If for every surjective and distance preserving* *map $X \to Y \in$ **epMet** we have $d_P(V(X), V(Y)) = 0$, then V is uniformly continuous (resp. 2-Lipschitz) with respect to the Gromov–Hausdorff distance and the metric on P.*

As a corollary of Theorem F, we recover, in Section 6.3, the homotopy stability of the Vietoris–Rips filtration, proven in [BL17].

In Section 6.4, we study the quotient interleaving distance on the category of persistent dissimilarity spaces (which in particular contains all persistent metric spaces, and thus, all filtered metric spaces). We refer to this distance as the *Gromov–Hausdorff-interleaving distance*. We explain in what way this distance generalizes previous distances on filtered metric spaces ([CM10c]) and on dynamic metric spaces ([KM20]). This distance is a useful abstraction: for example, we have the following.

Theorem G (Proposition 5.1.11)**.** *Let V : **epMet** $\to C^{\mathbf{R}}$ be a locally persistent functor that maps surjective and distance preserving maps to isomorphisms. Then, V is Lipzchitz with respect to the Gromov–Hausdorff distance and the interleaving distance, and the functor V_* : $\mathbf{epMet}^{\mathbf{R}} \to C^{\mathbf{R}\times\mathbf{R}}$, obtained by applying V pointwise, is Lipschitz with respect to the Gromov–Hausdorff-interleaving distance and the interleaving distance.*

The material in Section 6.5 is joint work with Alex Rolle. We define a bi-filtration of metric measure spaces that generalizes the degree-Rips bi-filtration ([LW15]): for any suitable kernel K (Definition 6.5.5), metric measure space (X, d_X, μ_X), and $s, k > 0$, we let the *kernel density filtration* of X at s, k be:

$$\mathsf{KDF}(X)(s,k) = \left\{ x \in X \, : \, \int_{x' \in X} K\left(\frac{d_X(x,x')}{s}\right) \mathrm{d}\mu_X \geq k \right\} \subseteq X.$$

We show that this filtration extends to a functor from compact metric probability spaces to bi-persistent metric spaces and prove the following stability result.

Theorem H (Theorem 6.5.1)**.** *The kernel density filtration is uniformly continuous with respect to the Gromov–Hausdorff–Prokhorov distance on compact metric probability spaces and the Gromov–Hausdorff-interleaving distance on bi-persistent metric spaces.*

Theorem G and Theorem H imply that the persistent homology of the kernel density filtration is a stable invariant of compact metric probability spaces.

In Section 6.6, we review some distances on the collection of hierarchical clusterings given in the literature. We show that the category of multi-dimensional hierarchical clusterings has the structure of a relative locally persistent category such that the quotient interleaving distance recovers known distances on hierarchical clusterings.

In Section 6.7, we study the locally persistent category $\mathbf{Set}^{\mathbf{R}}$ of persistent sets. This category contains many useful subcategories, such as the category of dendrograms and the category of merge trees. Adapting the definition of q-tame persistent vector space of [CSGO16], we say that a persistent set is *q-tame* if the image of every non-identity structure map is a finite set. We prove the following.

Theorem I (Theorem 6.7.2)**.** *The interleaving distance on q-tame persistent sets is geodesic and complete.*

In Section 6.8, we show that the distance on finite filtered simplicial complexes defined by Mémoli in [Mé17] is the quotient interleaving distance of a relative locally persistent category structure on the category of finite filtered simplicial complexes. We use the tools developed in this thesis to recover the fact that Mémoli's distance is geodesic. We also show that, after applying geometric realization, Mémoli's distance in general does not coincide with the homotopy interleaving distance of Blumberg and Lesnick, and that, in particular, it is not homotopy invariant.

In Section 6.9, we show that the Wasserstein distances between persistence diagrams can be recovered as the interleaving distance of suitable locally persistent categories.

1.3 Related work

As mentioned in Section 1.1, using category theory to frame and work with interleavings has been the subject of much recent work. Bubenik, de Silva, and Scott study interleaving distances in the context of categories of generalized persistent modules ([BSS13]), while de Silva, Munch, and Stefanou define an interleaving distance in any category with a flow ([SMS18]). In the context of categories with a flow, Cruz studies metric properties of the interleaving distance ([Cru19]).

In parallel, there have been many efforts in defining homotopically meaningful versions of interleaving distances, meaning interleaving distances that are homotopy invariant for some notion of weak equivalence. Of relevance to this thesis are Blumberg and Lesnick's homotopy interleaving distance ([BL17]), and Mémoli's distance on finite filtered simplicial complexes ([Mé17]). Thus far, there has been no study of the interplay between interleavings and weak equivalences in a general categorical framework such as generalized persistent modules or categories with a flow.

A main contribution of this thesis is offering yet another categorical interpretation of the theory of interleavings, namely locally persistent categories. We hope to demonstrate that this language is simpler and closer to usual category theory than previous approaches, and that it admits a clean and useful homotopical enhancement, namely the theory relative locally persistent categories, which permits a formal study of the interplay between weak equivalences and interleavings. Moreover, in Chapter 5, we show that previous categorical approaches to the theory of interleavings, such as generalized persistent modules and categories with a flow, can be seen as convenient ways of *constructing* locally persistent categories. This makes the theory of locally persistent categories automatically applicable to many important examples considered in the literature.

This thesis is also influenced by the categorical interpretation of metric spaces of Lawvere ([Law73]). Locally persistent categories are a categorification of Lawvere metric spaces in the same way that categories are a categorification of partially ordered sets.

Chapter 2

Background

In this chapter, we introduce the necessary background needed to state and prove the results in this thesis. The material in Sections 2.2 to 2.5 is mostly standard, but we recall it here for convenience and to establish notation. The material in Section 2.6 is an application of standard results in category theory to the theory of persistence. It is most likely known to experts but, to best of the author's knowledge, there is no reference explaining it.

2.1 Basic notation, categories, and size issues

We will assume familiarity with the language of category theory, and, in particular, with the notions of *category*, *functor*, *natural transformation*, *(co)limit*, *(co)end*, *adjunction*, *Kan extension*, and *monad*. We recommend the references [Lan98] and [Rie17]. We will generally denote categories by C, D, etc. , and use $\mathscr{C}$, $\mathscr{D}$, etc. for categories enriched in a monoidal category different from **Set** (Section 2.4). If x and y are isomorphic objects of a category C, we write $x \cong y$, and reserve the notation $\simeq$ for weaker notions, such as weak equivalence (Section 3.3). If η is a natural transformation between functors F and G, we write $\eta : F \Rightarrow G$ if we are regarding F and G as diagrams, and $\eta : F \to G$ if we are regarding F and G as objects of a functor category.

Particularly relevant kinds of limits and colimits are *pullbacks* and *pushouts*, and *sequential limits* and *sequential colimits*, that is, limits indexed by the category

$$\cdots \to \bullet \to \bullet \to \bullet,$$

and colimits indexed by the opposite of the above category. We say that a class

W of morphisms of a category C is **closed under sequential limits** if the induced morphisms from the sequential limit of a sequential diagram of morphisms in W to each of the objects in the diagram are in W. Dually, W is closed under sequential colimits if, seen as a class of morphisms of C^{op}, it is closed under sequential limits.

In this thesis, we will consider small and large sets. One can make this notion precise by working with a Grothendieck universe U and letting the small sets be the U-small sets and the large sets be the sets that are not necessarily U-small. Since in the arguments and results of this thesis there are no hidden size issues or subtleties, we will not be more precise than this, and we will point out that a certain set is large or small only when it matters. Every notion that requires an underlying set (such as the notions of set, metric space, and topological space) gives rise to two possible collections of instances. For example, there is a category of small sets, which we denote by **Set**, and a category of large sets, which we denote by **SET**. We will use this notational convention throughout the thesis. For example, we will talk about the categories of small and large extended pseudo metric spaces, denoted by **epMet** and **epMET** respectively, and about the categories of small and large topological spaces, denoted by **Top** and **TOP** respectively.

2.2 Extended pseudo metric spaces

The contents of this section are standard concepts in metric geometry; a good reference is [BBI01]. The only difference between [BBI01] and the exposition here is that we work with extended pseudo metric spaces, a simple generalization of metric spaces.

2.2.1 Elementary notions

We start with the definitions of extended pseudo metric space and of metric space.

Definition 2.2.1. *An **extended pseudo metric space (ep metric space)** (P, d_P) consists of a set P and a function $d_P : P \times P \to [0,\infty]$ such that*

- $d_P(p,p) = 0$ *for all* $p \in P$ *(**reflexivity**);*
- $d_P(p,p') = d_P(p',p)$ *for all* $p,p' \in P$ *(**symmetry**);*
- $d_P(p,p'') \leq d_P(p,p') + d_P(p',p'')$ *(**triangle inequality**).*

A function d_P satisfying the properties above is called an ***ep metric****.*

Remark 2.2.2. Although ep metrics are not as standard as metrics, they induce a topology in exactly the same way. Namely, given an ep metric space (P, d_P), one gets a topology by considering the topology generated by the family of open balls $\{B(p,\varepsilon)\}_{p\in P, \varepsilon\in(0,\infty)}$.

Definition 2.2.3. *A* ***metric space*** *consists of an ep metric space (P, d_P) such that for every $p, p' \in P$, $d_P(p, p') = 0$ implies $p = p'$, and such that d_P doesn't take the value ∞.*

Next we consider morphisms between ep metric spaces.

Definition 2.2.4. *A* ***distance non-increasing map*** *(or* ***1-Lipschitz map****) between ep metric spaces (P, d_P) and (Q, d_Q) consists of a function of sets $f : P \to Q$ such that $d_P(p, p') \geq d_Q(f(p), f(p'))$.*

We can then form a category.

Definition 2.2.5. *The* ***category of ep metric spaces****, denoted by* **epMet***, is the category whose objects are ep metric spaces and whose morphisms are distance non-increasing maps.*

Morphisms that don't increase or decrease the metric will play an important role when studying the category **epMet**.

Definition 2.2.6. *A* ***distance preserving map*** *between ep metric spaces is a morphism of ep metric spaces $f : P \to Q$ that satisfies $d_P(p, p') = d_Q(f(p), f(p'))$ for all $p, p' \in P$.*

2.2.2 Properties of a metric

In this section, we define the notions of an ep metric space being complete, compact, totally bounded, geodesic, and intrinsic. These are natural extensions of the corresponding notions for metric spaces. By an abuse of language, we will sometimes say that a metric is complete, meaning that the underlying ep metric space is complete. Similarly, we may say that a metric is compact, totally bounded, geodesic, or intrinsic.

Definition 2.2.7. *A sequence $\{x_i\}_{i\in\mathbb{N}}$ of elements of an ep metric space (P, d_P) is* ***Cauchy*** *if for every $\varepsilon > 0$ there exists an $n \in \mathbb{N}$ such that, if $i, j \geq n$, then $d_P(x_i, x_j) < \varepsilon$.*

Definition 2.2.8. *A sequence $\{x_i\}_{i\in\mathbb{N}}$ of elements of an ep metric space (P, d_P) is **convergent** if there exists an element $x \in P$ such that $d_P(x_i, x) \to 0$ as $i \to \infty$. When this is the case, we say that x is a **limit** of the sequence $\{x_i\}$.*

Note that in an ep metric space, a convergent sequence may have multiple distinct limits. But, of course, all of these have to be at distance zero from each other.

Definition 2.2.9. *An ep metric space (P, d_P) is **complete** if every Cauchy sequence is convergent.*

Definition 2.2.10. *An ep metric space (P, d_P) is **totally bounded** if, for every $\varepsilon > 0$, there exist finitely many points $\{p_i\} \subseteq P$ such that $P \subseteq \bigcup_i B(p_i, \varepsilon)$.*

Definition 2.2.11. *An ep metric space (P, d_P) is **compact** if every sequence has a convergent subsequence.*

The following lemma has exactly the same proof as its analogue for metric spaces.

Lemma 2.2.12. *An ep metric space is compact if and only if it is complete and totally bounded.* □

We now turn our attention to geodesic and intrinsic distances. In order to define these concepts, we need the notion of length of a continuous path in an ep metric space.

Definition 2.2.13. *Let (P, d_P) be an ep metric space, and let $f : [a, b] \to P$ be a continuous map. The **length** of the path f is defined to be the supremum of*

$$\sum_{i=1}^{N} d_P(f(y_{i-1}), f(y_i))$$

over all finite collections of points $a = y_0 \leq y_1 \leq \cdots \leq y_N = b$.

Note that the length of a curve can be infinite.

Definition 2.2.14. *An ep metric space (P, d_P) is **geodesic** if for every $p, p' \in P$ with $d_P(p, p') < \infty$, there is a continuous path $f : [a, b] \to P$ such that $f(a) = p$, $f(b) = p'$ and the length of f is equal to $d_P(p, p')$.*

Definition 2.2.15. *An ep metric space (P, d_P) is **intrinsic** if for every $p, p' \in P$, the distance $d_P(p, p')$ is equal to the infimum over all paths between p and p' of the length of the path.*

Of course, every geodesic metric is intrinsic. When proving that a metric has a certain property, it is often convenient to prove that a larger space has this property, and to show that the original space is a retract of this larger space. We conclude this section by formalizing this situation in the case of ep metrics.

Definition 2.2.16. *An ep metric space P is a* ***pseudo retract*** *of an ep metric space Q if there exist distance non-increasing maps $s : P \to Q$ and $r : Q \to P$ such that, for all $p \in P$, we have $d_P(p, r(s(p))) = 0$.*

Lemma 2.2.17. *If P is a pseudo retract of an ep metric space Q, and Q is complete (resp. intrinsic, geodesic), then P is complete (resp. intrinsic, geodesic).*

Proof. Let $\{x_n\}$ be a Cauchy sequence in P. Then, $\{s(x_n)\}$ is Cauchy in Q. As such, it has a limit y. This implies that $r(y)$ is a limit for the sequence $\{r(s(x_n))\}$ in P. And, since $d_P(r(s(x_n)), x_n) = 0$ for all n, the point $r(y)$ must also be a limit for the original sequence $\{x_n\}$.

The same line of reasoning proves the claim about the distance being intrinsic or geodesic. □

2.2.3 Quotients of metrics

The reason why we consider ep metrics and not just metrics will become clear when we define interleaving distances. But even without this motivation, we can already show one of its advantages, namely, that we can take quotients of metrics by equivalence relations without having to take a quotient of the underlying set of the metric space.

Definition 2.2.18. *Assume given an ep metric space (X, d) and an equivalence relation $R \subseteq X \times X$. We say that the metric d is* ***R-invariant*** *if $d(x, y) = d(x', y')$ whenever $(x, x'), (y, y') \in R$.*

There is a universal way of turning an ep metric into an R-invariant one, as the next proposition shows.

Proposition 2.2.19. *Given an ep metric space (X, d) and an equivalence relation $R \subseteq X \times X$, there is a unique ep metric $d_{/R} : X \times X \to [0, \infty]$ satisfying the following.*

1. *$d_{/R}(x, y) \leq d(x, y)$ for all $x, y \in X$;*

2. *$d_{/R}$ is R-invariant;*

3. *for any other metric d' satisfying* (1) *and* (2)*, we have $d' \leq d_{/R}$.*

Proof. Condition (3) guarantees that there is at most one ep metric satisfying all the conditions. Let D be the set of all metrics satisfying (1) and (2). Note that this set is non-empty, since the metric that is constantly 0 belongs to it. Let $d_{/R}(x, y) = \sup_{d' \in D} d'(x, y)$. It is straightforward to check that this metric satisfies all the requirements. □

We give a name to this universal construction.

Definition 2.2.20. *The metric determined by conditions* (1)*,* (2)*, and* (3) *above is called the **quotient ep metric** of d by R.*

Arbitrary quotients of a metric are generally not very well behaved. An exception to this is the fact that any quotient of an intrinsic metric is intrinsic. To prove this, we need the following characterization of the quotient metric.

Lemma 2.2.21. *Let (P, d_P) be an intrinsic ep metric space and let R be an equivalence relation on P. For any $p, p' \in P$, let*

$$d(p, p') = \inf\left\{ \sum_{i=1}^{N} d_P(y_i, y_i') : y_0, \ldots, y_N, y_0', \ldots, y_N' \in P, \right.$$

$$\left. pRy_0, y_i' R y_{i+1}, y_N' R p' \right\}.$$

Then $(d_P)_{/R} = d$.

Proof. It is easy to check that d is an R-invariant ep metric that is bounded above by d_P.

Now, let d' be an R-invariant ep metric bounded by d_P. Let $y_0, \ldots, y_N, y_0', \ldots, y_N' \in P$ such that pRy_0 ,$p'Ry_N'$, and $y_i'Ry_{i+1}$. By the triangle inequality, and the fact that d' is R-invariant, it follows that $d'(p, p') \leq \sum_{i=1}^{N} d_P(y_i, y_i')$, so $d' \leq d$.

This means that d satisfies the universal property of the quotient distance, and thus $(d_P)_{/R} = d$. □

We can now prove that any quotient of an intrinsic metric is intrinsic.

Proposition 2.2.22. *Let (P, d_P) be an intrinsic ep metric space and let R be an equivalence relation on P. Then $(d_P)_{/R}$ is intrinsic.*

Proof. By Lemma 2.2.21, we have

$$(d_P)_{/R}(p,p') = \inf\left\{ \sum_{i=1}^{N} d_P(y_i, y_i') : y_0,\dots,y_N, y_0',\dots,y_N' \in P, \right.$$
$$\left. pRy_0, y_i'Ry_{i+1}, y_N'Rp' \right\}.$$

Consider, for each $n \in \mathbb{N}$, sequences $\{y_i^n\}$ and $\{y_i'^n\}$ with $0 \le i \le N_n$ such that $\sum_{i=1}^{N_n} d_P(y_i^n, y_i'^n)$ converges to $(d_P)_{/R}(p,p')$ as $n \to \infty$. Since d_P is intrinsic, there is a path $f_i^n : [a_n, b_n] \to P$ between y_i^n and $y_i'^n$ such that its length is less than $d_P(y_i^n, y_i'^n) + 1/(nN_n)$. Since $(d_P)_{/R}(y_i', y_{i+1}) = 0$, these paths can be glued to a continuous path in $(P, (d_P)_{/R})$ whose length is at most

$$\sum_{i=1}^{N_n} \left(d_P(y_i^n, y_i'^n) + \frac{1}{nN_n} \right) = \sum_{i=1}^{N_n} \left(d_P(y_i^n, y_i'^n) \right) + \frac{1}{n}$$

The proposition follows by taking the limit $n \to \infty$. □

2.2.4 The Gromov–Hausdorff distance

Definition 2.2.23. *Let P be a metric space and let $A, B \subseteq P$ be subsets. The* ***Hausdorff distance*** *between A and B is defined by*

$$d_H^P(A,B) = \inf\left\{\varepsilon \ge 0 : B \subseteq A^\varepsilon, A \subseteq B^\varepsilon\right\},$$

where, for a subset $A \subseteq P$ and $\varepsilon \ge 0$, we let $A^\varepsilon = \{p \in P : \exists a \in A, d_P(a,p) < \varepsilon\}$.

Definition 2.2.24. *The* ***Gromov–Hausdorff distance*** *between metric spaces P and Q is defined by*

$$d_{GH}(P,Q) = \inf_{\substack{i:P\to Z \\ j:Q\to Z}} d_H^Z(i(P), j(Q)),$$

where the infimum is taken over all distance preserving inclusions i and j into a common metric space Z.

Note that, although one usually restricts P and Q to be compact, the definition makes sense for general P and Q ([BBI01, Definition 7.3.10]), even if they are ep metric spaces. At this level of generality, the Gromov–Hausdorff distance is actually just an ep metric. The triangle inequality is easy to prove by glueing metrics [BBI01, Proposition 7.3.16].

In this short section, we give an equivalent characterization of the Gromov–Hausdorff distance in terms of correspondences. This characterization is well-known, see for example [BBI01, Section 7.3.3].

Definition 2.2.25. *A **correspondence** R between sets X and Y is a subset $R \subseteq X \times Y$ such that the induced projections $R \to X$ and $R \to Y$ are surjective.*

Definition 2.2.26. *Let R be a correspondence between two metric spaces P and Q. The **distortion** of R is defined as*

$$\operatorname{dist}(R) = \sup\left\{|d_P(p,p') - d_Q(q,q')| : (p,q),(p',q') \in R\right\}.$$

Theorem 2.2.27. *For any metric spaces P and Q, we have*

$$2d_{GH}(P,Q) = \inf\left\{\operatorname{dist}(R) : R \subseteq P \times Q \text{ a correspondence}\right\}.$$

Proof. If $d_{GH}(X,Y) < r$, then there is a metric space Z and distance preserving inclusions $i : P \to Z$ and $j : Q \to Z$ such that $d_H^Z(i(P), j(Q)) < r$. Let $R \subseteq P \times Q$ be given by the pairs (p,q) such that $d_Z(i(p), j(q)) < r$. This is a correspondence, since, by hypothesis, $d_H^Z(i(P), j(Q)) < r$. Now, if $(p,q),(p',q') \in R$, then

$$|d_P(p,p') - d_Q(q,q')| \le d_Z(i(p), j(q)) + d_Z(i(p'), j(q')) < 2r,$$

by the triangle inequality of Z and the fact that i and j are embeddings. So $\operatorname{dist}(R) < 2r$.

Going the other way, assume given a correspondence $R \subseteq P \times Q$ such that $\operatorname{dist}(R) = 2r$. Consider the metric space Z with underlying set $P \coprod Q$ and metric given by

$$\begin{aligned} d_Z(p,p') &= d_P(p,p') \\ d_Z(q,q') &= d_Q(q,q') \\ d_Z(p,q) &= \inf\left\{d_P(p,p') + d_Q(q',q) + r : (p',q') \in R\right\}. \end{aligned}$$

To see that this satisfies the triangle inequality, it suffices to check that given $p_1, p_2 \in P$ and $q \in Q$, we have $d_Z(p_1,p_2) \le d_Z(p_1,q) + d_Z(q,p_2)$. We do this by applying the triangle inequality a few times and the fact that $\mathrm{dist}(R) = 2r$:

$$\begin{aligned}
& d_Z(p_1, q) + d_Z(q, p_2) \\
= & \inf_{(p_1', q') \in R} (d_P(p_1, p_1') + d_Q(q, q')) + \inf_{(p_2', q'') \in R} (d_P(p_2, p_2') + d_Q(q, q'')) + 2r \\
\geq & \inf_{(p_1', q') \in R} \inf_{(p_2', q'') \in R} d_P(p_1, p_1') + d_P(p_2, p_2') + d_Q(q', q'') + 2r \\
\geq & \inf_{(p_1', q') \in R} \inf_{(p_2', q'') \in R} d_P(p_1, p_1') + d_P(p_2, p_2') + d_P(p_1', p_2') \\
\geq & d_P(p_1, p_2) = d_Z(p_1, p_2).
\end{aligned}$$

Finally, we must see that the images of P and Q in Z are at Hausdorff distance less than or equal to r. This follows from the definition of d_Z and the fact that the projections $R \to P$ and $R \to Q$ are surjective. □

2.3 Monoidal categories

In this section we define closed symmetric monoidal categories. For more details, we refer the reader to [Kel82]. We remark that an understanding of the contents of this section is not strictly necessary to understand the main concepts in this thesis: although these concepts are inspired by enriched category theory, we unfold the main definitions to avoid heavy categorical language whenever possible. It is nonetheless very helpful to interpret the topics in this thesis using enriched category theory.

In Section 2.6.1, we describe our main example of monoidal category, the category of persistent sets.

Definition 2.3.1. *A* ***monoidal category*** *is a category $\mathcal{V}$ together with*

1. *a functor $\otimes : \mathcal{V} \times \mathcal{V} \to \mathcal{V}$ called the* ***tensor product****;*

2. *an object $1 \in \mathcal{V}$ called the* ***tensor unit****;*

3. *for each $x, y, z \in \mathcal{V}$, a natural isomorphism*

$$\alpha_{x,y,z} : (x \otimes y) \otimes x \xrightarrow{\simeq} x \otimes (y \otimes z),$$

called the ***associator****;*

4. *for every $x \in \mathcal{V}$, a natural isomorphism*

$$\lambda_x : 1 \otimes x \xrightarrow{\cong} x,$$

*called the **left unitor**;*

5. *for every $x \in \mathcal{V}$, a natural isomorphism*

$$\rho_x : x \otimes 1 \xrightarrow{\cong} x,$$

*called the **right unitor**;*

such that all of the diagrams of the following two forms commute in $\mathcal{V}$:

▷ *the **triangle identity**:*

$$\begin{tikzcd} (x \otimes 1) \otimes y \arrow[rr, "\alpha_{x,1,y}"] \arrow[dr, "\rho_x \otimes \mathrm{id}_y"'] & & x \otimes (1 \otimes y) \arrow[dl, "\mathrm{id}_x \otimes \lambda_y"] \\ & x \otimes y & \end{tikzcd}$$

▷ *the **pentagon identity**:*

$$\begin{tikzcd} & (w \otimes x) \otimes (y \otimes z) \arrow[dr, "\alpha_{w,x,y \otimes z}"] & \\ ((w \otimes x) \otimes y) \otimes z \arrow[ur, "\alpha_{w \otimes x,y,z}"] \arrow[d, "\alpha_{w,x,y} \otimes \mathrm{id}_z"'] & & (w \otimes (x \otimes (y \otimes z))) \\ (w \otimes (x \otimes y)) \otimes z \arrow[rr, "\alpha_{w,x \otimes y,z}"'] & & w \otimes ((x \otimes y) \otimes z). \arrow[u, "\mathrm{id}_w \otimes \alpha_{x,y,z}"'] \end{tikzcd}$$

There are a few different natural notions of functor between monoidal categories. We will use lax monoidal functors and strong monoidal functors.

Definition 2.3.2. *A **lax monoidal functor** between two monoidal categories $(\mathcal{V}, \otimes_{\mathcal{V}}, 1_{\mathcal{V}})$ and $(\mathcal{W}, \otimes_{\mathcal{W}}, 1_{\mathcal{W}})$ is given by*

1. *a functor $F : \mathcal{V} \to \mathcal{W}$;*

2. *a morphism* $\varepsilon : 1_{\mathscr{W}} \to F(1_{\mathcal{V}})$*;*

3. *for all* $x, y \in \mathcal{V}$*, a morphism*

$$\mu_{x,y} : F(x) \otimes_{\mathscr{W}} F(y) \to F(x \otimes_{\mathcal{V}} y),$$

natural in x *and* y*;*

such that all the following two kinds of diagrams commute in $\mathscr{W}$*:*

1. *associativity diagram:*

$$\begin{tikzcd}
(Fx \otimes_{\mathscr{W}} Fy) \otimes_{\mathscr{W}} Fz \arrow[r, "\alpha_{Fx,Fy,Fz}"] \arrow[d, "\mu_{x,y} \otimes_{\mathscr{D}} \mathsf{id}"'] & Fx \otimes_W (Fy \otimes_{\mathscr{W}} Fz) \arrow[d, "\mathsf{id} \otimes \mu_{y,z}"] \\
F(x \otimes_{\mathcal{V}} y) \otimes_{\mathscr{W}} Fz \arrow[d, "\mu_{x \otimes_{\mathcal{V}} y, z}"'] & Fx \otimes_W F(y \otimes_{\mathcal{V}} z) \arrow[d, "\mu_{x, y \otimes_{\mathcal{V}} z}"] \\
F((x \otimes_{\mathcal{V}} y) \otimes_{\mathcal{V}} z) \arrow[r, "F(\alpha_{x,y,z})"] & F(x \otimes_{\mathcal{V}} (y \otimes_{\mathcal{V}} z))
\end{tikzcd}$$

2. *unitality diagrams:*

$$\begin{tikzcd}
1_{\mathscr{W}} \otimes_{\mathscr{W}} Fx \arrow[r, "\varepsilon \otimes \mathsf{id}"] \arrow[d, "\lambda_{Fx}"'] & F(1_{\mathcal{V}}) \otimes_{\mathscr{W}} Fx \arrow[d, "\mu_{1_{\mathcal{V}},x}"] \\
Fx & F(1_{\mathcal{V}} \otimes x), \arrow[l, "F(\lambda_x)"]
\end{tikzcd}$$

$$\begin{tikzcd}
Fx \otimes_{\mathscr{W}} 1_{\mathscr{W}} \arrow[r, "\mathsf{id} \otimes \varepsilon"] \arrow[d, "\rho_{Fx}"'] & Fx \otimes_{\mathscr{W}} F(1_{\mathcal{V}}) \arrow[d, "\mu_{x,1_{\mathcal{V}}}"] \\
Fx & F(x \otimes 1_{\mathcal{V}}). \arrow[l, "F(\rho)"]
\end{tikzcd}$$

If ε *and* $\mu_{x,y}$ *are isomorphisms, then* F *is called a* ***strong monoidal functor****.*

The corresponding notion of natural transformation is the following.

Definition 2.3.3. *Let* $(\mathcal{V}, \otimes_{\mathcal{V}}, 1_{\mathcal{V}})$ *and* $(\mathcal{W}, \otimes_{\mathcal{W}}, 1_{\mathcal{W}})$ *be monoidal categories, and let* $(F, \varepsilon^F, \mu^F)$ *and* $(G, \varepsilon^G, \mu^G)$ *be lax monoidal functors from* $\mathcal{V}$ *to* $\mathcal{W}$*. A* ***monoidal natural transformation*** η *from* F *to* G *consists of a natural transformation* $\eta : F \Rightarrow G$ *between the underlying functors, such that all of the following two kinds of diagrams commute in* $\mathcal{W}$*:*

1. *respect for monoidal product:*

$$\begin{tikzcd}
F(x) \otimes_{\mathcal{W}} F(y) \arrow[r, "\eta_x \otimes_{\mathcal{W}} \eta_y"] \arrow[d, "\mu^F_{x,y}"'] & G(x) \otimes_{\mathcal{W}} G(y) \arrow[d, "\mu^G_{x,y}"] \\
F(x \otimes_{\mathcal{V}} y) \arrow[r, "\eta_{x \otimes_{\mathcal{V}} y}"'] & G(x \otimes_{\mathcal{V}} y),
\end{tikzcd}$$

2. *respect for units:*

$$\begin{tikzcd}
& 1_{\mathcal{W}} \arrow[dl, "\varepsilon^F"'] \arrow[dr, "\varepsilon^G"] & \\
F(1_{\mathcal{V}}) \arrow[rr, "\eta_{1_{\mathcal{V}}}"] & & G(1_{\mathcal{V}}).
\end{tikzcd}$$

The following definition formalizes the notion of a monoidal product being commutative.

Definition 2.3.4. *A* ***symmetric monoidal category*** *consists of a monoidal category* $\mathcal{V}$ *together with, for every* $x, y \in \mathcal{V}$*, a natural isomorphism* $B_{x,y} : x \otimes y \to y \otimes x$ *such that* $B_{y,x} \circ B_{x,y} = \mathrm{id}_{x \otimes y}$ *and such that all of the diagrams of the following form commute in* $\mathcal{V}$*: the* ***hexagon identity****:*

$$\begin{tikzcd}
(x \otimes y) \otimes x \arrow[r, "\alpha_{x,y,z}"] \arrow[d, "B_{x,y} \otimes \mathrm{id}_z"'] & x \otimes (y \otimes z) \arrow[r, "B_{x,y \otimes z}"] & (y \otimes z) \otimes x \arrow[d, "\alpha_{y,z,x}"] \\
(y \otimes x) \otimes z \arrow[r, "\alpha_{y,x,z}"] & y \otimes (x \otimes z) \arrow[r, "\mathrm{id}_y \otimes B_{x,z}"] & y \otimes (z \otimes x).
\end{tikzcd}$$

Finally, the notion of closedness formalizes the idea of having internal homs, that is, of having objects that represent the collection of morphisms between two objects of the category.

Definition 2.3.5. *A **closed symmetric monoidal category** is a symmetric monoidal category $\mathcal{V}$ such that, for all objects $x \in \mathcal{V}$, the functor $- \otimes x : \mathcal{V} \to \mathcal{V}$ has a right adjoint functor $[x, -] : \mathcal{V} \to \mathcal{V}$. For $x, y \in \mathcal{V}$, the object $[x, y] \in \mathcal{V}$ is called the **internal hom** of x and y.*

2.4 Enriched categories

Monoidal categories serve as a basis for enriching categories. Again, for more details about enriched category theory, we refer the reader to [Kel82].

2.4.1 Elementary notions

A category enriched in a monoidal category $\mathcal{V}$ is, informally, a category where the hom objects are not sets, but objects of $\mathcal{V}$.

Definition 2.4.1. *Let $\mathcal{V}$ be a monoidal category. A $\mathcal{V}$**-category** $\mathcal{C}$ (or $\mathcal{V}$-enriched category) consists of*

- *a collection of **objects**, denoted by $\mathsf{obj}(\mathcal{C})$;*
- *for $x, y \in \mathsf{obj}(\mathcal{C})$, an object $\mathrm{Hom}(x, y) \in \mathsf{obj}(\mathcal{V})$, called the **hom-object** from x to y;*
- *for each $x, y, z \in \mathsf{obj}(\mathcal{C})$, a morphism*

$$\circ_{x,y,z} : \mathrm{Hom}(y, z) \otimes \mathrm{Hom}(x, y) \to \mathrm{Hom}(x, z),$$

  *called the **composition morphism**;*
- *for each object $x \in \mathsf{obj}(\mathcal{C})$, a morphism $\iota_x : 1 \to \mathrm{Hom}(x, x)$, called the **identity morphism** of x;*

such that all of the diagrams of the following two forms commute in $\mathcal{V}$:

1. *the associativity diagram:*

$$\begin{tikzcd}
(\mathrm{Hom}(y,z) \otimes \mathrm{Hom}(x,y)) \otimes \mathrm{Hom}(w,x) \arrow[rr, "\alpha"] \arrow[d, "\circ_{x,y,z} \otimes \mathsf{id}"'] & & \mathrm{Hom}(y,z) \otimes (\mathrm{Hom}(x,y) \otimes \mathrm{Hom}(w,x)) \arrow[d, "\mathsf{id} \otimes \circ_{w,x,y}"] \\
\mathrm{Hom}(x,z) \otimes \mathrm{Hom}(w,x) \arrow[r, "\circ_{w,x,z}"] & \mathrm{Hom}(w,z) & \mathrm{Hom}(y,z) \otimes \mathrm{Hom}(w,y), \arrow[l, "\circ_{w,y,z}"']
\end{tikzcd}$$

2. *the unitality diagram:*

$$\begin{tikzcd}
\mathrm{Hom}(x,x)\otimes\mathrm{Hom}(w,x) \arrow[rr, "\circ_{w,x,x}"] & & \mathrm{Hom}(w,x) & & \mathrm{Hom}(w,x)\otimes\mathrm{Hom}(w,w) \arrow[ll, "\circ_{w,w,x}"'] \\
1\otimes\mathrm{Hom}(w,x) \arrow[u, "\iota_x\otimes\mathsf{id}"] \arrow[urr, "\lambda"'] & & & & \mathrm{Hom}(w,x)\otimes 1. \arrow[u, "\mathsf{id}\otimes\iota_w"'] \arrow[ull, "\rho"]
\end{tikzcd}$$

The following construction lets us compare categories that are enriched over different monoidal categories and will be important we constructing examples of our main object of study: locally persistent categories.

Definition 2.4.2. *Let $\mathscr{V}$ and $\mathscr{W}$ be monoidal categories, let $F:\mathscr{V}\to\mathscr{W}$ be a lax monoidal functor, and let $\mathscr{C}$ be a $\mathscr{V}$-enriched category. The* ***change of enrichment*** *of $\mathscr{C}$ along F is the $\mathscr{W}$-enriched category whose objects are the same as the objects of $\mathscr{C}$, and whose hom-object $\mathrm{Hom}(x,y)$ is given by $F(\mathrm{Hom}_{\mathscr{C}}(x,y))$. Identities and composition are defined using the (lax) monoidal structure of F.*

Example 2.4.3. Fix $\mathscr{V}$ a monoidal category. There is a lax monoidal functor $\mathscr{V}\to$ **Set** given by mapping v to $\mathrm{Hom}_{\mathscr{V}}(1,v)$. The change of enrichment gives us a **Set**-enrichment for $\mathscr{C}$. The (ordinary) category thus obtained is called the **underlying category** of $\mathscr{C}$, and is denoted by $\mathscr{C}_0$.

Remark 2.4.4. The change of enrichment construction for a lax monoidal functor $F:\mathscr{V}\to\mathscr{W}$ between monoidal categories $\mathscr{V}$ and $\mathscr{W}$ provides us with a change of enrichment functor

$$F:\mathscr{V}\text{-}\mathbf{Cat}\to\mathscr{W}\text{-}\mathbf{Cat}.$$

This construction respects natural transformations, that is, given $F,G:\mathscr{V}\to\mathscr{W}$ a lax monoidal functors and $\eta:F\Rightarrow G$ a monoidal natural transformation, we get a natural transformation $\eta:F\Rightarrow G$ as functors $\mathscr{V}\text{-}\mathbf{Cat}\to\mathscr{W}\text{-}\mathbf{Cat}$.

Definition 2.4.5. *Let $\mathscr{V}$ be a monoidal category. Given $\mathscr{V}$-enriched categories $\mathscr{C}$ and $\mathscr{D}$, a* ***$\mathscr{V}$-enriched functor*** *$F:\mathscr{C}\to\mathscr{D}$ consists of*

- *an mapping $F:\mathsf{obj}(\mathscr{C})\to\mathsf{obj}(\mathscr{D})$;*
- *for every $x,y\in\mathscr{C}$, a morphism*

$$F_{x,y}:\mathrm{Hom}_{\mathscr{C}}(x,y)\to\mathrm{Hom}_{\mathscr{D}}(Fx,Fy);$$

such that all of the following two kinds of diagrams commute in $\mathcal{V}$:

1. *respect for composition:*

$$\begin{tikzcd}
\mathrm{Hom}_{\mathscr{C}}(y,z)\otimes\mathrm{Hom}_{\mathscr{C}}(x,y) \arrow[r, "\circ_{x,y,z}"] \arrow[d, "F_{y,z}\otimes F_{x,y}"'] & \mathrm{Hom}_{\mathscr{C}}(x,z) \arrow[d, "F_{x,z}"] \\
\mathrm{Hom}_{\mathscr{D}}(Fy,Fz)\otimes\mathrm{Hom}_{\mathscr{C}}(Fx,Fy) \arrow[r, "\circ_{Fx,Fy,Fz}"] & \mathrm{Hom}_{\mathscr{D}}(Fx,Fz)
\end{tikzcd}$$

2. *respect for units:*

$$\begin{tikzcd}
& 1 \arrow[dl, "\iota_x"'] \arrow[dr, "\iota_{Fx}"] & \\
\mathrm{Hom}_{\mathscr{C}}(x,x) \arrow[rr, "F_{x,x}"] & & \mathrm{Hom}_{\mathscr{D}}(Fx,Fx).
\end{tikzcd}$$

The collection of all $\mathcal{V}$-enriched categories forms a category in its own right.

Definition 2.4.6. *Let $\mathcal{V}$ be a monoidal category. The* ***category of $\mathcal{V}$-enriched categories*** *is the category whose objects are $\mathcal{V}$-enriched categories and whose morphisms are $\mathcal{V}$-enriched functors. We denote this category by $\mathcal{V}$-**Cat** or $_{\mathcal{V}}$**Cat**.*

2.4.2 Constructions with enriched categories

In this section we assume that $\mathcal{V}$ is locally small, complete and cocomplete. As in the non-enriched case, we can take the opposite of an enriched category.

Definition 2.4.7. *Let $\mathcal{V}$ be a monoidal category. For any $\mathcal{V}$-enriched category $\mathscr{C}$, define the* ***opposite category*** *$\mathscr{C}^{\mathrm{op}}$ to be the $\mathcal{V}$-enriched category with the same objects as $\mathscr{C}$, hom-object $\mathrm{Hom}_{\mathscr{C}^{\mathrm{op}}}(x,y)$ given by $\mathrm{Hom}_{\mathscr{C}}(y,x)$ for every $x,y\in\mathscr{C}$, identities given by the identities of $\mathscr{C}$, and composition given by swapping the arguments of the composition of $\mathscr{C}$.*

We can also consider the (tensor) product of two enriched categories.

Definition 2.4.8. *Let $\mathcal{V}$ be a monoidal category and let $\mathscr{C}$ and $\mathscr{D}$ be $\mathcal{V}$-enriched categories. The* ***tensor product*** *of $\mathscr{C}$ and $\mathscr{D}$, denoted by $\mathscr{C}\otimes\mathscr{D}$, is the $\mathcal{V}$-enriched category* *with objects* $\mathsf{obj}(\mathscr{C}\otimes\mathscr{D})=\mathsf{obj}(\mathscr{C})\times\mathsf{obj}(\mathscr{D})$, *hom-object* $\mathrm{Hom}_{\mathscr{C}\otimes\mathscr{D}}((x,y),(x',y'))$ *given by* $\mathrm{Hom}_{\mathscr{C}}(x,x')\otimes\mathrm{Hom}_{\mathscr{D}}(y,y')$, *identities given by* $\iota_x\otimes_{\mathscr{V}}\iota_y:1\to\mathrm{Hom}((x,y),(x,y))$, *and composition given by tensoring the composition of* $\mathscr{C}$ *and* $\mathscr{D}$ *as morphisms in* $\mathscr{V}$.

In order to define enriched natural transformations between enriched functors, we must define (universal) extranatural transformations and ends.

Definition 2.4.9. *Let* $\mathscr{V}$ *be a monoidal category, let* $\mathscr{C}$ *and* $\mathscr{D}$ *be* $\mathscr{V}$*-enriched categories, and let* $F:\mathscr{C}^{\mathsf{op}}\otimes\mathscr{C}\to\mathscr{D}$ *be a* $\mathscr{V}$*-enriched functor. A* $\mathscr{V}$***-extranatural transformation*** $\theta:d\xrightarrow{\bullet}F$ *from* $d\in\mathscr{D}$ *to* F *consists of, for every* $c\in\mathscr{C}$, *a morphism* $\theta_c:d\to F(c,c)$ *in* $\mathscr{D}_0$, *such all of the diagrams of the following form commute:*

$$\begin{tikzcd}
\mathrm{Hom}_{\mathscr{C}}(x,y) \arrow[r, "F(x,-)"] \arrow[d, "F(-,y)"'] & \mathrm{Hom}_{\mathscr{D}}(F(x,x),F(x,y)) \arrow[d, "\mathrm{Hom}_{\mathscr{D}}(\theta_x,\mathsf{id})"] \\
\mathrm{Hom}_{\mathscr{D}}(F(y,y),F(x,y)) \arrow[r, "\mathrm{Hom}_{\mathscr{D}}(\theta_y,\mathsf{id})"'] & \mathrm{Hom}_{\mathscr{D}}(d,F(x,y)).
\end{tikzcd}$$

Definition 2.4.10. *Let* $\mathscr{V}$ *be a monoidal category, let* $\mathscr{C}$ *be a* $\mathscr{V}$*-enriched category, let* $F:\mathscr{C}^{\mathsf{op}}\otimes\mathscr{C}\to\mathscr{V}$ *be a* $\mathscr{V}$*-enriched functor, and let* $v\in\mathscr{V}$. *A* $\mathscr{V}$*-extranatural transformation* $\theta:v\xrightarrow{\bullet}F$ *is **universal** if every* $\mathscr{V}$*-extranatural transformation* $\alpha:v'\xrightarrow{\bullet}F$ *is given by* $\alpha_x=\theta_x\circ f$ *for a unique morphism* $f:v'\to v$ *in* $\mathscr{V}$.

*When such a universal transformation exists, it is called an **end** of* F, *and it is denoted by* $\int_{c\in\mathscr{C}}F(c,c)$.

It is easy to check that any two ends of the same functor are isomorphic.

We are ready to define enriched natural transformations between enriched functors.

Definition 2.4.11. *Let* $\mathscr{V}$ *be a monoidal category, let* $\mathscr{C}$ *and* $\mathscr{D}$ *be* $\mathscr{V}$*-enriched categories, and let* $F,G:\mathscr{C}\to\mathscr{D}$ *be* $\mathscr{V}$*-enriched functors. The* $\mathscr{V}$***-object of natural transformations*** *between* F *and* G *is defined to be* $\int_{c\in\mathscr{C}}\mathrm{Hom}_{\mathscr{D}}(F(c),G(c))$, *whenever it exists. It is denoted by* $\mathrm{Hom}_{[\mathscr{C},\mathscr{D}]}(F,G)$.

If the enriched category $\mathscr{C}$ is small (i.e. , the collection $\mathsf{obj}(\mathscr{C})$ is a set), then the end exists.

We now wish to extend the above definition to the definition of the enriched functor category of enriched functors between two fixed enriched categories. In order to do this, consider, for each $c \in \mathscr{C}$

$$E_c : \mathrm{Hom}_{[\mathscr{C},\mathscr{D}]}(F,G) = \int_{c' \in \mathscr{C}} \mathrm{Hom}_{\mathscr{D}}(F(c'),G(c')) \to \mathrm{Hom}_{\mathscr{D}}(F(c),G(c)),$$

the canonical morphism out of the end (often referred to as the counit). Observe that the composite

$$\begin{aligned}\mathrm{Hom}_{[\mathscr{C},\mathscr{D}]}(F,G) \otimes \mathrm{Hom}_{[\mathscr{C},\mathscr{D}]}(H,F) &\xrightarrow{E_c \otimes E_d} \mathrm{Hom}_{\mathscr{D}}(F(c),G(c)) \otimes \mathrm{Hom}_{\mathscr{D}}(H(c),F(c)) \\ &\xrightarrow{\circ} \mathrm{Hom}_{\mathscr{D}}(H(c),G(c))\end{aligned}$$

forms an extranatural transformation. This provides us with a morphism

$$\circ_{H,F,G} : \mathrm{Hom}_{[\mathscr{C},\mathscr{D}]}(F,G) \otimes \mathrm{Hom}_{[\mathscr{C},\mathscr{D}]}(H,F) \to \mathrm{Hom}_{[\mathscr{C},\mathscr{D}]}(H,G).$$

Finally, the identity morphisms of $\mathscr{D}$ give us extranatural transformations

$$\iota_F : 1 \to \mathrm{Hom}_{\mathscr{D}}(F(c),F(c)).$$

Definition 2.4.12. *Let $\mathscr{V}$ be a monoidal category, let $\mathscr{C}$ and $\mathscr{D}$ be $\mathscr{V}$-enriched categories, and let $F,G : \mathscr{C} \to \mathscr{D}$ be $\mathscr{V}$-enriched functors. The **functor $\mathscr{V}$-category** $[\mathscr{C},\mathscr{D}]$ has as objects $\mathscr{V}$-enriched functors between $\mathscr{C}$ and $\mathscr{D}$, as hom-objects the ones defined in Definition 2.4.11, and identity morphisms and composition as defined above.*

2.4.3 (Co)powers and weighted (co)limits

Since the hom-objects of a $\mathscr{V}$-enriched category $\mathscr{C}$ are objects of $\mathscr{V}$, given $x, y \in \mathscr{C}$ and $k \in \mathscr{V}$ it makes sense to talk about k-morphisms between x and y. These are just morphisms $k \to \mathrm{Hom}_{\mathscr{C}}(x,y)$ in $\mathscr{V}$. A 1-morphism (or just morphism) between x and y is then a morphism $1 \to \mathrm{Hom}_{\mathscr{C}}(x,y)$, or equivalently, a morphism form x to y in the underlying category $\mathscr{C}_0$.

It is often useful to represent k-morphisms between x and y as 1-morphisms between related objects. This is what the next definition accomplishes.

Definition 2.4.13. *Let $\mathscr{V}$ be a closed symmetric monoidal category and let $\mathscr{C}$ be a $\mathscr{V}$-enriched category. The **copower** of an object $x \in \mathscr{C}$ by an object $k \in \mathscr{V}$ consists of an* *object $k \cdot x \in \mathscr{C}$ together with, for every $y \in \mathscr{C}$, a natural isomorphism*

$$\mathrm{Hom}_{\mathscr{C}}(k \cdot x, y) \cong [k, \mathrm{Hom}_{\mathscr{C}}(x, y)],$$

where $[-,-]$ denotes the internal hom of $\mathscr{V}$.

*Dually, the **power** of an object $y \in \mathscr{C}$ by k consists of an object $y^k \in \mathscr{C}$ together with, for every $x \in \mathscr{C}$, a natural isomorphism*

$$\mathrm{Hom}_{\mathscr{C}}\left(x, y^k\right) \cong [k, \mathrm{Hom}_{\mathscr{C}}(x, y)].$$

We conclude this section with an extension of the notion of (co)limit to the enriched case. The main idea is that, since we are working with an enriched category and thus have a notion of k-morphism for every $k \in \mathscr{V}$, the indexing diagram of a (co)limit should come with *weights* that specify what kind of morphism should be used when constructing (co)cones for the diagram.

For intuition about the notion of weighted (co)limit we recommend [Shu06] for weighted (co)limits in the context of homotopy theory, [Rut98] for weighted (co)limits in the context of (Lawvere) metric spaces, and Section 3.2.7, where we interpret some universal constructions that are relevant to the theory of interleavings as weighted limits.

Definition 2.4.14. *Let $\mathscr{V}$ be a closed symmetric monoidal category and let K and $\mathscr{C}$ be $\mathscr{V}$-enriched categories. A **weighted limit** over an enriched functor $F : K \to \mathscr{C}$ with respect to a weight $W : K \to \mathscr{V}$ consists of an object $\lim^W F \in \mathscr{C}$ together with, for every $c \in \mathscr{C}$, a natural isomorphism*

$$\mathrm{Hom}_{\mathscr{C}}\left(c, \lim^W F\right) \cong \mathrm{Hom}_{[K,V]}\left(W, \mathrm{Hom}_{\mathscr{C}}(c, F(-))\right).$$

*Dually, a **weighted colimit** over F with respect to W consists of an object $\mathrm{colim}^W F \in \mathscr{C}$ together with, for every $c \in \mathscr{C}$, a natural isomorphism*

$$\mathrm{Hom}_{\mathscr{C}}\left(\mathrm{colim}^W F, c\right) \cong \mathrm{Hom}_{[K,V]}\left(W, \mathrm{Hom}_{\mathscr{C}}(F(-), c)\right).$$

*An enriched category is **(co)complete** if it admits all small (co)limits with arbitrary weights.*

2.4.4 Enriched Kan extensions

We now define Kan extensions in the context of enriched category theory. Although this subject has some subtleties ([Dub70], [Kel82, Section 4]), this won't be a problem for us, as we will only use very basic facts and definitions.

Definition 2.4.15. *Let $\mathscr{V}$ be a monoidal category, let $\mathscr{C}$, $\mathscr{C}'$, and $\mathscr{D}$ be $\mathscr{V}$-enriched categories with $\mathscr{C}$ and $\mathscr{C}'$ small and $\mathscr{D}$ complete, and let $G : \mathscr{C} \to \mathscr{C}'$ and $F : \mathscr{C} \to \mathscr{D}$ be $\mathscr{V}$-enriched functors. The* ***pointwise right Kan extension*** *of F along G, denoted by $\mathrm{Ran}_G F : \mathscr{C}' \to \mathscr{D}$, is defined by the weighted limit:*

$$(\mathrm{Ran}_G F)(c') = \lim^{\mathrm{Hom}_{\mathscr{C}'}(c', G(-))} F.$$

Dually, if $\mathscr{D}$ is cocomplete, the ***pointwise left Kan extension*** *of F along G is defined by the weighted colimit:*

$$(\mathrm{Lan}_G F)(c') = \mathrm{colim}^{\mathrm{Hom}_{\mathscr{C}'}(G(-), c')} F.$$

From the definition, we get canonical natural transformations $\mathrm{Ran}_G F \circ G \Rightarrow F$ and $F \Rightarrow \mathrm{Lan}_G F \circ G$.

A $\mathscr{V}$-enriched functor $G : \mathscr{C} \to \mathscr{C}'$ between $\mathscr{V}$-enriched categories is **fully faithful** if, for every $x, y \in \mathscr{C}$, it induces an isomorphism $\mathrm{Hom}_{\mathscr{C}}(x, y) \to \mathrm{Hom}_{\mathscr{C}'}(Fx, Fx')$ in $\mathscr{V}$.

The following result is standard.

Proposition 2.4.16. *Let $\mathscr{V}$ be a monoidal category, let $\mathscr{C}$, $\mathscr{C}'$, and $\mathscr{D}$ be $\mathscr{V}$-enriched categories with $\mathscr{C}$ and $\mathscr{C}'$ small and $\mathscr{D}$ complete (resp. cocomplete), and let $G : \mathscr{C} \to \mathscr{C}'$ and $F : \mathscr{C} \to \mathscr{D}$ be $\mathscr{V}$-enriched functors. If G is fully faithful, then the natural transformation $\mathrm{Ran}_G F \circ G \Rightarrow F$ (resp. $F \Rightarrow \mathrm{Lan}_G F \circ G$) is a natural isomorphism.* □

2.5 Model categories

In this section, we describe the very basics of the theory of model categories, and give the examples we are interested in. For details, we refer the reader to [Qui67] for the original description of the theory, and to [Hov07] and [Hir09] for more modern accounts of it.

We start with the notion of lifting property.

Definition 2.5.1. *Let $i : A \to B$ and $p : X \to Y$ be morphisms in a category C. We say that i has the* ***left lifting property with respect to p*** *and that p has the* ***right lifting property with respect to i*** *if, for every commutative square*

$$\begin{tikzcd} A \arrow[r, "f"] \arrow[d, "i"'] & Y \arrow[d, "p"] \\ B \arrow[r, "g"] & X \end{tikzcd}$$

there is a morphism $h : B \to Y$ such that $h \circ i = f$ and $p \circ h = g$.

The following definition formalizes the notion of functorial factorization. For a category C, we let $C^{\to}$ denote the category of morphisms of C, that is, the category of functors from the category freely generated by two objects and a single morphism between them, to C.

Definition 2.5.2. *A* ***functorial weak factorization system*** *on a category C consists of a pair $(\mathfrak{L}, \mathfrak{R})$ of classes of morphisms of C such that there exist functors $F_L, F_R : C^{\to} \to C^{\to}$ such that*

- *For every $f \in C^{\to}$, we have $F_L(f) \in \mathfrak{L}$ and $F_R(f) \in \mathfrak{R}$ and $F_R(f) \circ F_L(f) = f$.*
- *The class $\mathfrak{L}$ is precisely the class of morphisms having the left lifting property against every morphism in $\mathfrak{R}$, and the class $\mathfrak{R}$ is precisely the class of morphisms having the left lifting property against every morphism in $\mathfrak{L}$.*

We can now give a concise definition of model structure.

Definition 2.5.3. *A* ***model structure*** *on a category C consists of three classes of morphisms of C called* ***weak equivalences*** *(denoted by $\mathscr{W}$),* ***cofibrations*** *(denoted by* **Cof***), and* ***fibrations*** *(denoted by* **Fib***) such that*

1. *(2-out-of-3) If f and g are composable morphisms of C and two of f, g and $g \circ f$ are weak equivalences, then so is the third.*
2. *$(\mathbf{Cof}, \mathbf{Fib} \cap \mathscr{W})$ and $(\mathbf{Cof} \cap \mathscr{W}, \mathbf{Fib})$ form two functorial factorization systems of C.*

The morphisms in $\mathbf{Cof} \cap \mathscr{W}$ are called **trivial cofibrations** and the morphisms in $\mathbf{Fib} \cap \mathscr{W}$ are called **trivial fibrations**.

It follows from the definition that a model structure, if it exists, it is completely determined by the weak equivalences and one of the classes **Cof** or **Fib**.

Definition 2.5.4. *A **model category** consists of a complete and cocomplete category together with a model structure.*

Arbitrary model structures can be badly behaved. Many useful model structures are cofibrantly generated, which intuitively means that one has to check relatively few things when proving that a certain morphism is a fibration or a trivial fibration. We now give the formal definitions.

Definition 2.5.5. *Let C be cocomplete and let I be a class of morphisms of C.*

- ▹ *We write* $\mathbf{cell}(I)$ *for the class of morphisms obtained by transfinite composition of pushouts of coproducts of elements in I.*
- ▹ *We write* $\mathbf{cof}(I)$ *for the class of retracts of elements of* $\mathbf{cell}(I)$.

Definition 2.5.6. *A model category is **cofibrantly generated** if there are small sets of morphisms I and J of C such that*

- ▹ $\mathbf{cof}(I)$ *is precisely the class of cofibrations of C;*
- ▹ $\mathbf{cof}(J)$ *is precisely the class of trivial cofibrations of C;*
- ▹ *I and J admit the small object argument, meaning that the domains of morphisms of I are small relative to* $\mathbf{cell}(I)$ *and the domains of morphisms of J are small relative to* $\mathbf{cell}(J)$.

Of great use are model structures on functor categories. Thanks to compositionality, it is often the case that a model structure on a category C induces a model structure on a functor category C^D. In fact, there are two canonical choices, the projective model structure, and the injective model structure. We will be interested in the projective case.

Definition 2.5.7. *Let C be a model category, let D be a small category, and consider the functor category C^D. A **projective weak equivalence** is a morphism of C^D that is an objectwise weak equivalence. A **projective fibration** is a morphism of C^D that is an objectwise fibration. The **projective model structure** on C^D is the model structure whose weak equivalences are the projective weak equivalences and whose fibrations are the projective fibrations, provided it exists.*

The following is well-known and appears, for example, in [Hir09, Section 11.6].

Theorem 2.5.8. *If C is a cofibrantly generated model category and D is a small category, then the projective model structure on C^D exists.* □

We finish this section with the examples most relevant to us. These are some of the most well-known examples of model categories and go back to Quillen's work ([Qui67]). They appear as Theorems 2.3.11, 2.4.19, and 3.6.5 of [Hov07], respectively.

Example 2.5.9. Let R be a commutative ring. There is a cofibrantly generated model structure on $\mathbf{Ch}(R)$ such that the weak equivalences are the homology isomorphisms, and such that every object is fibrant.

Example 2.5.10. There is a cofibrantly generated model structure on **Top** such that the weak equivalences are the continuous functions inducing isomorphisms in all homotopy groups, and such that every object is fibrant.

Example 2.5.11. There is a cofibrantly generated model structure on **sSet** such that the weak equivalences are the simplicial maps whose geometric realization is a weak equivalence of topological spaces, and such that every object is cofibrant.

2.6 Persistent objects

In this section, we study the notion of *persistent object*, and prove some basic facts about categories of persistent objects. Persistent objects have also been referred to as *generalized persistent modules* ([BSS13]). Although some of the results as stated cannot be found in the literature, they are consequences of well-known facts in category theory, and are sometimes used implicitly in the persistence literature.

We regard posets as categories. A **monoidal poset** $(P, \otimes, 1)$ consists of a poset P together with a binary operation $- \otimes - : P \times P \to P$ and a unit object $1 \in P$ that underly a (necessarily unique) monoidal structure when interpreting the poset as a category.

The poset $([0,\infty), \leq)$ will be denoted by $\mathbf{R}_+$. This is a monoidal poset, with monoidal product given by addition of real numbers. The poset $((-\infty,\infty), \leq)$ will be denoted by $\mathbf{R}$, and we will reserve $\mathbb{R}$ for the metric space given by the real numbers. The poset $\mathbf{R}$ is also monoidal, with monoidal product given by addition of real numbers. If P is a poset and $r, s \in P$, we write $s < r$ whenever $s \leq r$ and $s \neq r$.

The main objects of study of this thesis are persistent objects, (i.e. objects of a functor category C^P) and categories enriched in persistent objects. Persistent objects can be shifted, as follows.

Definition 2.6.1. *Let $(P, \otimes, 1)$ be a monoidal poset and let C be a category. Given $p \in P$ and a functor $F : P \to C$, we define the* ***p-shift to the left*** *of F as the functor $F^p : P \to C$ given by*

$$F^p(q) = F(p \otimes q).$$

In order to enrich a category over a category of persistent objects, we must give a monoidal structure for the category of persistent objects. The monoidal product is given by Day convolution, which we now explain.

2.6.1 Day convolution of persistent objects

Let $(P, \otimes, 1)$ be a small monoidal category. The category of functors indexed by P with values in **Set** inherits a monoidal product called **Day convolution** ([DK69], [Day70]). Given two functors $F, G : P \to$ **Set**, the Day convolution $F \otimes_{\mathsf{Day}} G : P \to$ **Set** is defined by

$$(F \otimes_{\mathsf{Day}} G)(p) = \int^{(p_1,p_2)\in P\times P} \mathrm{Hom}_P(p_1 \otimes p_2, p) \times F(p_1) \times G(p_2).$$

The structure morphisms are defined in a straightforward way, using the universal property of coends. Since we will only use the definition in the case where the indexing monoidal category is a poset, we now specialize the above definition to that case.

Let $(P, \leq, \otimes)$ be a monoidal poset. The Day convolution of two functors $F, G : P \to$ **Set** is given by

$$(F \otimes_{\mathsf{Day}} G)(r) = \int^{s\otimes t \leq r} F(s) \times G(t). \tag{2.6.2}$$

The indexing poset of the coend is the subposet of $P \times P$ spanned by pairs (s, t) such that $s \otimes t \leq r$.

Day convolution automatically gives a closed symmetric monoidal structure. We now give the formula for the internal hom in the case where the indexing category is a poset. The internal hom $[F, G]_{\mathsf{Day}} : P \to$ **Set** between two functors $F, G : P \to$ **Set** is given by

$$[F, G]_{\mathsf{Day}}(r) = \mathrm{Nat}(F, G^r), \tag{2.6.3}$$

where $G^r : P \to$ **Set** is the r-shift to the left of G.

Our motivating example is the following.

Example 2.6.4. The main case of interest to us is when P is $\mathbf{R}_+$. In that case, the category $\mathbf{Set}^{\mathbf{R}_+}$ generalizes both dendrograms and ultra metric spaces (Section 6.7.2).

Categories enriched in $\mathbf{Set}^{\mathbf{R}_+}$ are the main object of study of this thesis, and are called **locally persistent categories**.

When working in a functor category $\mathbf{Set}^P$, the representable functors are of special interest. These are the functors in the (essential) image of the Yoneda embedding, which we now specialize to our case.

Definition 2.6.5. *Let P be a poset and let $r \in P$. Let $\mathscr{Y}(r) : P \to \mathbf{Set}$ be such that $\mathscr{Y}(r)(s)$ is a singleton set if $r \leq s$ and the empty set if $r > s$. The mapping $r \mapsto \mathscr{Y}(r)$ provides us with a functor $\mathscr{Y} : P^{\mathrm{op}} \to \mathbf{Set}^P$, called the **coYoneda embedding**, or the **Yoneda embedding** for simplicity.*

We now state a few consequences of the Yoneda lemma. By Eq. (2.6.3) and the Yoneda lemma, we have

$$G^r \cong [\mathscr{Y}(r), G]_{\mathsf{Day}} \tag{2.6.6}$$

and by adjunction and Eq. (2.6.3), we have

$$[F, G]_{\mathsf{Day}}(r) = \mathrm{Nat}(F, G^r) \cong \mathrm{Nat}(F, [\mathscr{Y}(r), G]_{\mathsf{Day}}) \cong \mathrm{Nat}(F \otimes_{\mathsf{Day}} \mathscr{Y}(r), G). \tag{2.6.7}$$

We now give an important result about Day convolution that says that, in a sense, Day convolution is the most natural monoidal structure on a functor category of the form $\mathbf{Set}^P$. We specialize it to the case of poset-indexed functors, but the results holds for general indexing monoidal categories.

Lemma 2.6.8. *Let P be a monoidal poset. Then, the (co)Yoneda embedding $\mathscr{Y} : P^{\mathrm{op}} \to \mathbf{Set}^P$ is strong monoidal.*

Proof. Let $r, u, v \in P$. Using Eq. (2.6.2), we see that

$$(\mathscr{Y}(u) \otimes_{\mathsf{Day}} \mathscr{Y}(v))(r) = \int^{s \otimes t \leq r} \mathscr{Y}(u)(s) \times \mathscr{Y}(v)(t).$$

By definition of $\mathscr{Y}(u)$, we have that $\mathscr{Y}(u)(s)$ is a singleton if $u \leq s$ and the empty set otherwise. Similarly, $\mathscr{Y}(v)(t)$ is a singleton if $v \leq t$ and the empty set otherwise. This implies that

$$\int^{s \otimes t \leq r} \mathscr{Y}(u)(s) \times \mathscr{Y}(v)(t) \cong \mathscr{Y}(u + v),$$

as required. □

2.6.2 Continuity of persistent objects

This section presents two classes of well-behaved persistent objects, and interprets this well-behaved condition as a (co)sheaf condition. This condition is a completeness condition, and will become important when proving that an interleaving distance is geodesic (Section 4.5).

Definition 2.6.9. *Let C be a cocomplete category and let P be a poset. A functor $F : P \to C$ is* ***left continuous*** *if for every $r \in P$ the canonical morphism*

$$\operatorname{colim} F_{<r} \to F(r)$$

is an isomorphism. Here $F_{<r} : \{r' \in P : r' < r\} \to C$ denotes the restriction of F to the subposet of P given by all elements strictly smaller than r.

Dually, let C be a complete category. A functor $F : P \to C$ is ***right continuous*** *if for every $r \in P$ the canonical morphism*

$$F(r) \to \lim F_{>r}$$

is an isomorphism.

Right and left continuous functors enjoy some useful closure properties.

Proposition 2.6.10. *Let P be a monoidal poset, and let $F, G : P \to$ **Set**. If G is right continuous, then the internal hom $[F, G]_{\mathsf{Day}}$ is right continuous.*

Proof. We simply compute

$$\begin{aligned}
\lim_{r'>r} [F, G]_{\mathsf{Day}}(r') &\cong \lim_{r'>r} \operatorname{Nat}(F, G^{r'}) \\
&\cong \operatorname{Nat}(F, \lim_{r'>r} G^{r'}) \\
&\cong \operatorname{Nat}(F, G^{r}) \\
&\cong [F, G]_{\mathsf{Day}}(r).
\end{aligned}$$

□

Universal property of $(-)^\#$. When P is the poset $\mathbf{R}_+$ or the poset $\mathbf{R}$, there is a universal way of turning a functor $F : \mathbf{R}_+ \to C$ into a left or right continuous functor. We explain the case of right continuity, the case of left continuity being dual. We only consider the poset $\mathbf{R}_+$ for simplicity; the following discussion generalizes to products of $\mathbf{R}_+$ and $\mathbf{R}$.

Let C be a complete category. Let $C^{\mathbf{R}_+}_{\mathsf{right}}$ be the subcategory of $C^{\mathbf{R}_+}$ consisting of right continuous functors. Consider the functor $C^{\mathbf{R}_+} \to C^{\mathbf{R}_+}_{\mathsf{right}}$ that maps a functor $F : \mathbf{R}_+ \to C$ to the functor $F^{\#} : \mathbf{R}_+ \to C$ defined by

$$F^{\#} = \lim_{r>0} F^r,$$

where F^r denotes the r-shift to the left of F, as in Definition 2.6.1. For every $F : \mathbf{R}_+ \to C$, there is a natural morphism $\eta^{\#}_F : F \to F^{\#}$. Moreover, $F^{\#}$ is right continuous since

$$\lim_{s'>s} F^{\#}(s') = \lim_{s'>s}\lim_{r>0} F^r(s') = \lim_{s'>s}\lim_{r>0} F(s'+r) \cong \lim_{s''>s} F(s'') \cong F^{\#}(s),$$

where in the first isomorphism we used the fact the poset map $(s,\infty) \times (0,\infty) \to (s,\infty)$ given by mapping (t,t') to $t+t'$ is coinitial, i.e. for every $u > s$, there is $(t,t') \in (s,\infty) \times (0,\infty)$ such that $t+t' < u$. It is easy to see that $(-)^{\#} : C^{\mathbf{R}_+} \to C^{\mathbf{R}_+}$ exhibits $C^{\mathbf{R}_+}_{\mathsf{right}}$ as a reflective subcategory of $C^{\mathbf{R}_+}$, that is, that there is a natural bijection $\mathrm{Nat}(F^{\#},G) \cong \mathrm{Nat}(F,G)$ for $F,G : \mathbf{R}_+ \to C$ and G right continuous, given by precomposition with $\eta^{\#}_F$.

Monoidality of $(-)^{\#}$. Consider the isomorphism $\varepsilon : \mathscr{Y}(0) \to \mathscr{Y}(0)^{\#}$ given by the fact that $\mathscr{Y}(0)$ is right continuous, and the natural transformation $\mu_{F,G} : F^{\#} \otimes_{\mathsf{Day}} G^{\#} \to (F \otimes_{\mathsf{Day}} G)^{\#}$ that corresponds to the natural morphism $\eta_{F\otimes_{\mathsf{Day}} G} : F \otimes_{\mathsf{Day}} G \to (F \otimes_{\mathsf{Day}} G)^{\#}$ under the composite isomorphism

$$\begin{aligned}
\mathrm{Nat}(F^{\#} \otimes_{\mathsf{Day}} G^{\#}, (F \otimes_{\mathsf{Day}} G)^{\#}) &\cong \mathrm{Nat}(F^{\#}, [G^{\#}, (F \otimes_{\mathsf{Day}} G)^{\#}]_{\mathsf{Day}}) \\
&\cong \mathrm{Nat}(F, [G^{\#}, (F \otimes_{\mathsf{Day}} G)^{\#}]_{\mathsf{Day}}) \\
&\cong \mathrm{Nat}(F \otimes_{\mathsf{Day}} G^{\#}, (F \otimes_{\mathsf{Day}} G)^{\#}) \\
&\cong \mathrm{Nat}(G^{\#}, [F, (F \otimes_{\mathsf{Day}} G)^{\#}]_{\mathsf{Day}}) \\
&\cong \mathrm{Nat}(G, [F, (F \otimes_{\mathsf{Day}} G)^{\#}]_{\mathsf{Day}}) \\
&\cong \mathrm{Nat}(F \otimes_{\mathsf{Day}} G, (F \otimes_{\mathsf{Day}} G)^{\#}),
\end{aligned}$$

where we used Proposition 2.6.10 and the fact that $(-)^{\#}$ is a reflection into right continuous functors. The following proposition is then a consequence of Day's reflection theorem ([Day72]) and Proposition 2.6.10.

Proposition 2.6.11. *The morphisms ε and μ exhibit the functor $(-)^{\#} : C^{\mathbf{R}_+} \to C^{\mathbf{R}_+}$ as a lax monoidal functor, and the natural transformation $\eta^{\#} : \mathrm{id}_{C^{\mathbf{R}_+}} \Rightarrow (-)^{\#}$ as a monoidal natural transformation.* □

Other properties of $(-)^{\#}$**.** The following result says that, up to arbitrarily small shifts, F and $F^{\#}$ are indistinguishable.

Proposition 2.6.12. *Let C be a complete category. Let $F : \mathbf{R}_+ \to C$ and let $\varepsilon > 0$. There are morphisms $F \to (F^{\#})^{\varepsilon}$ and $F^{\#} \to F^{\varepsilon}$ such that the composites are equal to the natural maps $F \to F^{2\varepsilon}$ and $F^{\#} \to (F^{\#})^{2\varepsilon}$.*

Proof. On the one hand, we have $F \to F^{\#}$ given by η. On the other hand, for any $\varepsilon > 0$, we have $F^{\#} \to F^{\varepsilon}$ by construction. It is enough to show that these maps compose to the natural maps $F \to F^{\varepsilon}$ and $F^{\#} \to (F^{\#})^{\varepsilon}$, and this follows from the universal property of $F^{\#}$. □

Given $F : \mathbf{R}_+ \to \mathbf{Set}$ and $s \leq r$, let $\varphi^F_{s,r} : F(s) \to F(r)$ denote the structure morphism of F. We deduce the following.

Lemma 2.6.13. *Let $F : \mathbf{R}_+ \to \mathbf{Set}$. Given $s < r \in \mathbf{R}_+$ and $a, b \in \mathbf{R}_+(s)$, if $\eta^{\#}_F(a) = \eta^{\#}_F(b) \in F^{\#}(s)$, then $\varphi^F_{s,r}(a) = \varphi^F_{s,r}(b)$.* □

Continuity as sheaf condition. We conclude this section by interpreting right (resp. left) continuity as a sheaf (resp. cosheaf) condition, in the case where the indexing poset is $\mathbf{R}_+$ or $\mathbf{R}$.

Remark 2.6.14. Note that $\mathbf{R}_+$ is a full subcategory of $(\overline{\mathbf{R}}_+, \leq) = ([0,\infty], \leq)$, which is a frame, that is, a poset with all joins and all finite meets, and such that binary meets distribute over arbitrary joins.

Any frame is naturally equipped with the structure of a site, where a family of morphism $\{U_i \to U\}$ is covering precisely if $\bigvee_i U_i = U$.

This means that, if C is complete and cocomplete, there is a well-defined sheaf condition for functors $\overline{\mathbf{R}}_+ \to C$. Note that, given $r \in \overline{\mathbf{R}}_+$, there are exactly two covering sieves for r: $\{r' : r' > r\}$ and $\{r' : r' \geq r\}$. This means that a functor $F : \overline{\mathbf{R}}_+ \to C$ is a sheaf exactly if $F(r) \to \lim F_{>r}$ is an isomorphism, so exactly if it is right continuous.

The category $\mathbf{R}_+$, being a dense subsite of $\overline{\mathbf{R}}_+$, inherits a site structure with the same sheaf condition. With this site structure, being left continuous is the same as being a sheaf. Under this interpretation, the functor $(-)^{\#}$ is just sheafification.

The discussion above dualizes, so that being right continuous is equivalent to being a cosheaf. It also applies to the category $\mathbf{R}$, by seeing it as a dense subsite of $[-\infty, \infty]$.

Chapter 3

Locally persistent categories

To any category C, one can assign the equivalence relation on its collection of objects where two objects are related exactly if they are isomorphic. Any equivalence relation on a set has an associated extended pseudo metric, where the distance between two elements of the set is 0 if they are related and ∞ if they are not. In particular, to every category C, one can assign an extended pseudo metric on its collection of objects. This distance is rather discrete, but it completely characterizes the equivalence relation given by isomorphism.

In this chapter, we study *locally persistent categories*. These are categories with extra structure that allows one to define a notion of approximate isomorphism. In these categories, for each pair of objects and each $\varepsilon \in \mathbf{R}_+$, there is a set of ε-approximate isomorphisms between them, which are usually referred to as *ε-interleavings*. When one composes an ε-approximate isomorphism with a δ-approximate isomorphism, one obtains an $(\varepsilon+\delta)$-isomorphism. What in the case of categories was a discrete distance, becomes, in this case, a more interesting distance. This is the *interleaving distance* associated to a locally persistent category. Concretely, a locally persistent category is a category enriched in the functor category $\mathbf{Set}^{\mathbf{R}_+}$.

Category	**Locally persistent category**
morphism	ε-approximate morphism
isomorphism	ε-interleaving
equivalence relation given by isomorphism	interleaving distance

Table 3.1: Translating a few basic notions between categories and locally persistent categories.

In many cases, the interleaving distance is “too strict”. For example, as observed in [BL17], in the locally persistent category of persistent topological spaces, there exist homotopy equivalent persistent topological spaces whose interleaving distance is infinite. This problem has been approached in more than one way in the literature. In [Les12], the author considers the interleaving distance in a homotopy category of persistent topological spaces, and similarly, in [FLM17], the authors consider the interleaving distance in a homotopy category of $\mathbf{R}$-filtered topological spaces. In [BL17], the authors relax the notion of interleaving to a notion of *homotopy interleaving*, and define the *homotopy interleaving distance* in the category of persistent topological spaces. In [Mé17], the author defines a distance between $\mathbf{R}$-filtered finite simplicial complexes that shares many similarities with the homotopy interleaving distance of [BL17]. In order to incorporate such interleaving distances into our framework, we draw inspiration from the solutions of [Mé17] and [BL17] and consider locally persistent categories with additional homotopical structure that allows one to define a kind of homotopy interleaving distance, which we call the *quotient interleaving distance*. The extra homotopical structure comes in the form of a class of morphisms of our locally persistent category which we regard as weak equivalences, or acyclic morphisms.

A category together with a class of morphisms containing all identities is usually called a relative category. Following this convention, we call locally persistent categories together with the extra homotopical structure *relative locally persistent categories*.

Relative category	**Relative locally persistent category**
homotopy class of morphisms	homotopy class of ε-approximate morphism
weak equivalence	ε-quotient interleaving
equivalence relation given by weak equivalence	quotient interleaving distance

Table 3.2: Translating a few basic notions between relative categories and relative locally persistent categories.

The framework presented in this chapter is not just a way of organizing concepts: we will see in Chapter 4 that, in a locally persistent category, simple categorical structure (such as weighted (co)limits) gives rise to useful metric structure.

The chapter is structured as follows. In Section 3.1, we introduce locally persistent categories and the interleaving distance on their collections of objects. In Section 3.2 we study the category theory of locally persistent categories by describing diagrammatic reasoning and universal constructions relevant to this setting. Finally, in Section 3.3, we introduce quotient interleaving distances.

Note that, although we use the language of enriched category theory to motivate some definitions, the proofs in this chapter don't rely on any results of enriched category theory.

3.1 Main definitions

Recall that $\mathbf{R}_+ = ([0,\infty), \le, +)$ and $\mathbf{R} = ((-\infty,\infty), \le)$. We start by defining our main object of study. In the language of Section 2.4 and Section 2.6.1, we will be studying categories enriched in $\mathbf{Set}^{\mathbf{R}_+}$. Nonetheless, we will unfold definitions as much as possible, and, in this section, we will not rely on the theory of enriched categories for our definitions and proofs.

The following is an unpacking of the definition of category enriched in persistent sets, as defined in Example 2.6.4.

Definition 3.1.1. *A **locally persistent category** $\mathscr{C}$ consists of the following data:*

- ▷ *a collection of **objects**, denoted by* $\mathsf{obj}(\mathscr{C})$*;*
- ▷ *for each $x, y \in \mathsf{obj}(\mathscr{C})$ and each $\varepsilon \in \mathbf{R}_+$, a collection of **$\varepsilon$-approximate morphisms**, denoted by* $\mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon$*;*
- ▷ *for each $x \in \mathsf{obj}(\mathscr{C})$, an **identity morphism**, denoted by* $\mathsf{id}_x \in \mathrm{Hom}_{\mathscr{C}}(x,x)_0$*;*
- ▷ *for each $x, y, z \in \mathsf{obj}(\mathscr{C})$ and each $\varepsilon, \delta \in \mathbf{R}_+$, a **composition** operation*

$$- \circ - : \mathrm{Hom}_{\mathscr{C}}(y,z)_\delta \times \mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon \to \mathrm{Hom}_{\mathscr{C}}(x,z)_{\varepsilon+\delta};$$

- ▷ *for each $x, y \in \mathsf{obj}(\mathscr{C})$ and each $\varepsilon \le \delta \in \mathbf{R}_+$, a **shift** operation*

$$\mathbf{S}_{\varepsilon,\delta} : \mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon \to \mathrm{Hom}_{\mathscr{C}}(x,y)_\delta;$$

such that

- *identity morphisms act as the identity, in the sense that* $f \circ \mathrm{id} = \mathrm{id} \circ f = f$ *for any approximate morphism* f*;*

- *composition is associative, so that* $f \circ (g \circ h) = (f \circ g) \circ h$ *for any approximate morphisms;*

- *the shift* $\mathbf{S}$ *is a functor, in the sense that* $\mathbf{S}_{\varepsilon,\varepsilon} : \mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon \to \mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon$ *is the identity function, and for* $\varepsilon \leq \delta \leq \gamma$*, we have* $\mathbf{S}_{\delta,\gamma} \circ \mathbf{S}_{\varepsilon,\delta} = \mathbf{S}_{\varepsilon,\gamma} : \mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon \to \mathrm{Hom}_{\mathscr{C}}(x,y)_\gamma$*;*

- *the functor* $\mathbf{S}$ *respects composition, in the sense that for* $\varepsilon \leq \varepsilon' \in \mathbf{R}_+$ *and* $\delta \leq \delta' \in \mathbf{R}_+$*, and objects* $x, y, z \in \mathscr{C}$*, the following diagram commutes*

$$\begin{tikzcd}
\mathrm{Hom}_{\mathscr{C}}(y,z)_\delta \times \mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon \arrow[r, "\circ"] \arrow[d, "\mathbf{S}_{\delta,\delta'} \times \mathbf{S}_{\varepsilon,\varepsilon'}"'] & \mathrm{Hom}_{\mathscr{C}}(x,z)_{\varepsilon+\delta} \arrow[d, "\mathbf{S}_{\varepsilon+\delta,\varepsilon'+\delta'}"] \\
\mathrm{Hom}_{\mathscr{C}}(y,z)_{\delta'} \times \mathrm{Hom}_{\mathscr{C}}(x,y)_{\varepsilon'} \arrow[r, "\circ"] & \mathrm{Hom}_{\mathscr{C}}(x,z)_{\varepsilon'+\delta'}.
\end{tikzcd}$$

Notation 3.1.2. For conciseness, we will often refer to ε-approximate morphisms simply as ε-morphisms, and to approximate morphisms simply as morphisms.

The above definition deserves a few remarks. Firstly, note that the objects of a locally persistent category $\mathscr{C}$ together with the 0-morphisms form a usual category. We refer to this category as the **underlying category** of $\mathscr{C}$, and we denote it by $\mathscr{C}_0$. In this sense, a locally persistent category can be interpreted as a category with extra structure.

Secondly, note that being a locally persistent category is a self-dual notion. This means that every locally persistent category $\mathscr{C}$ has an associated opposite locally persistent category $\mathscr{C}^{\mathrm{op}}$ with the same objects, and such that $\mathrm{Hom}_{\mathscr{C}^{\mathrm{op}}}(x,y)_\varepsilon = \mathrm{Hom}_{\mathscr{C}}(y,x)_\varepsilon$ for every $\varepsilon \in \mathbf{R}_+$ and $x, y \in \mathrm{obj}(\mathscr{C})$. This allows one to dualize the universal constructions described in Section 3.2 and the results proven in Chapter 4.

We now give some examples of locally persistent categories. These and other examples are studied in depth in Chapter 6.

One of the most natural examples, and the main motivating example for many of the categorical approaches to interleaving distances, is the case of the category of persistent objects in a category C.

Example 3.1.3. Let C be a category. The **category of persistent objects** of C is the functor category $C^{\mathbf{R}}$. For a persistent object $X \in C^{\mathbf{R}}$ and $\varepsilon \in \mathbf{R}_+$, let $X^\varepsilon \in C^{\mathbf{R}}$ be the shift of X to the left by ε, as in Definition 2.6.1, that is $X^\varepsilon(t) = X(t+\varepsilon)$. We now endow the category $C^{\mathbf{R}}$ with a locally persistent category structure. An ε-morphism from a persistent object X to a persistent object Y consists of a natural transformation $X \to Y^\varepsilon$. The shift operator, $\mathrm{Nat}(X, Y^\varepsilon) \to \mathrm{Nat}(X, Y^{\varepsilon'})$ for $\varepsilon' \geq \varepsilon$ simply postcomposes with the natural transformation $Y^\varepsilon \to Y^{\varepsilon'}$ given by the structure morphisms of Y. Composition and identities work as in the category $C^{\mathbf{R}}$. Note that the underlying category of this locally persistent category is precisely the original functor category.

Example 3.1.4. Metric spaces form a locally persistent category, where the ε-morphisms are the morphisms that don't increase the distance more than ε. Concretely, we endow the category **epMet** with the following locally persistent category structure. An ε-morphism between ep metric spaces $P, Q \in$ **epMet** consists of a function $f : P \to Q$ between the underlying sets such that $d_P(p, p') + \varepsilon \geq d_Q(f(p), f(p'))$ for all $p, p' \in P$. The shift operator, sending ε-morphisms to ε'-morphisms for $\varepsilon \leq \varepsilon'$, is an inclusion in this case. The identity morphisms are given by the identity functions, and composition is just composition of functions. Note that a 0-morphism is precisely a 1-Lipschitz map, and thus the underlying category of this locally persistent category is our original category **epMet**.

The following example is a great source of applications. Any category with a flow, in the sense of [SMS18], gives rise to a locally persistent category. Here we outline the main idea; details about the constructions are given in Section 5.2.

Example 3.1.5. Let $(D, \mathcal{T})$ be a category with a flow, that is, a category D together with a lax monoidal functor $\mathcal{T} : \mathbf{R}_+ \to \mathrm{End}(D)$. The collection of ε-morphisms between objects $x, y \in D$ is given by the set $\mathrm{Hom}_D(x, \mathcal{T}_\varepsilon(y))$. The shift operator is given by postcomposition with the morphism $\mathcal{T}_\varepsilon(y) \to \mathcal{T}_{\varepsilon'}(y)$ for $\varepsilon \leq \varepsilon'$ induced by the flow. Identities are defined similarly. Composition is a bit more subtle and we don't describe it now. One should note that, in this case, the underlying category of the locally persistent category that we get does not coincide with D in general. The underlying category is in fact equivalent to the Kleisli category of the monad given by $\mathcal{T}_0 : D \to D$. This is not really an issue, as, in practice, the monad $\mathcal{T}_0$ is often naturally isomorphic to the identity, and, in that case, its Kleisli category is just D.

We now unfold the definition of enriched functor between categories enriched in $\mathbf{Set}^{\mathbf{R}_+}$. This is the natural notion of morphism between locally persistent categories.

Definition 3.1.6. *A **locally persistent functor** $F : \mathscr{C} \to \mathscr{D}$ between locally persistent categories $\mathscr{C}$ and $\mathscr{D}$ is given by the following data:*

- ▷ *a mapping $F : \mathsf{obj}(\mathscr{C}) \to \mathsf{obj}(\mathscr{D})$;*
- ▷ *for each $x, y \in \mathscr{C}$ and each $\varepsilon \in \mathbf{R}_+$, a mapping $F : \mathrm{Hom}_{\mathscr{C}}(x, y)_\varepsilon \to \mathrm{Hom}_{\mathscr{D}}(F(x), F(y))_\varepsilon$;*

such that

- ▷ *F respects identities, $F(\mathsf{id}_x) = \mathsf{id}_{F(x)}$;*
- ▷ *F respects composition, $F(f \circ g) = F(f) \circ F(g)$;*
- ▷ *F respects shifts, in the sense that for a morphism $f \in \mathrm{Hom}(x, y)_\varepsilon$ and $\varepsilon \le \delta$, we have $\mathbf{S}_{\varepsilon,\delta}(F(f)) = F(\mathbf{S}_{\varepsilon,\delta}(f))$.*

Small locally persistent categories together with locally persistent functors form a category that we denote by **lpCat**. Similarly, large locally persistent categories together with locally persistent functors form a category **lpCAT**.

Many important constructions in the theory of persistence can be interpreted as locally persistent functors. The following two examples are expanded upon in Section 6.3.

Example 3.1.7. The Vietoris–Rips and the Čech filtrations give locally persistent functors $\mathsf{VR}, \check{\mathsf{C}} : \mathbf{epMet} \to \mathbf{Top}^{\mathbf{R}}$.

As explained in the introduction to Chapter 3, the main reason to consider locally persistent categories is that this extra structure allows for the definition of a kind of approximate isomorphism, which in turn gives a notion of distance between the objects of the category. Approximate isomorphisms are called interleavings, and, as we shall see, they share many properties with isomorphisms.

Definition 3.1.8. *Let $\mathscr{C}$ be a locally persistent category and let $\varepsilon, \delta \in \mathbf{R}_+$. An **$(\boldsymbol{\varepsilon},\boldsymbol{\delta})$-interleaving** between objects $x, y \in \mathscr{C}$ is given by morphisms $f \in \mathrm{Hom}_{\mathscr{C}}(x, y)_\varepsilon$ and $g \in \mathrm{Hom}_{\mathscr{C}}(y, x)_\delta$ such that $g \circ f = \mathbf{S}_{0,\varepsilon+\delta}(\mathsf{id}_x)$ and $f \circ g = \mathbf{S}_{0,\varepsilon+\delta}(\mathsf{id}_y)$. A **$\boldsymbol{\delta}$-interleaving** is a (δ, δ)-interleaving.*

Note that a 0-interleaving in a locally persistent category $\mathscr{C}$ is precisely an isomorphism in the underlying category $\mathscr{C}_0$.

Remark 3.1.9. Interleavings in a locally persistent category correspond to what sometimes is referred to as *weak interleavings* ([SMS18]). There is a way to define *strong interleavings* in a locally persistent category with extra structure. This is discussed in Section 5.2.2.

In the rest of this section we define the interleaving distance and we show that locally persistent functors are distance non-increasing.

The following is a simple application of the composition law of a locally persistent category.

Lemma 3.1.10. *If $\mathscr{C}$ is a locally persistent category, $x, y \in \mathscr{C}$ are ε-interleaved, and $y, z \in \mathscr{C}$ are δ-interleaved, then x and z are $(\varepsilon+\delta)$-interleaved.*

Proof. Let $f \in \mathrm{Hom}(x,y)_\varepsilon$ and $g \in \mathrm{Hom}(y,x)_\varepsilon$ witness the fact that x and y are ε-interleaved, and let $h \in \mathrm{Hom}(y,z)_\delta$ and $i \in \mathrm{Hom}(z,y)_\delta$ witness the fact that y and z are δ-interleaved.

Consider the composites $h \circ f \in \mathrm{Hom}(x,z)_{\varepsilon+\delta}$ and $g \circ i \in \mathrm{Hom}(z,x)_{\varepsilon+\delta}$. In order to see that these form an $(\varepsilon+\delta)$-interleaving between x and z, we compute

$$\begin{aligned}(g\circ i)\circ(h\circ f) &= g \circ \mathbf{S}_{0,2\delta}(\mathsf{id}_y)\circ f\\ &= g\circ \mathbf{S}_{\varepsilon,\varepsilon+2\delta}(f)\\ &= \mathbf{S}_{2\varepsilon,2\varepsilon+2\delta}(g\circ f)\\ &= \mathbf{S}_{0,2\varepsilon+2\delta}(\mathsf{id}_x).\end{aligned}$$

An analogous computation shows that $(h\circ f)\circ(g\circ i) = \mathbf{S}_{0,2\varepsilon+2\delta}(\mathsf{id}_y)$. □

As a consequence, the interleaving distance, which we now introduce, satisfies the triangle inequality.

Definition 3.1.11. *Let $\mathscr{C}$ be a locally persistent category. Define the* ***interleaving distance*** $d_I^{\mathscr{C}} : \mathsf{obj}(\mathscr{C}) \times \mathsf{obj}(\mathscr{C}) \to [0,\infty]$ *as*

$$d_I^{\mathscr{C}}(x,y) = \inf\left\{\delta \in \mathbf{R}_+ :\ x \text{ and } y \text{ are } \delta\text{-interleaved}\right\},$$

with the convention that the infimum of the empty subset of $\mathbf{R}_+$ is ∞. This is an ep metric on $\mathsf{obj}(\mathscr{C})$.

The proof of the following result is an immediate application of the definitions, but it is nonetheless one of the most useful results of the theory of interleaving distances.

Theorem 3.1.12. *Let $F : \mathscr{C} \to \mathscr{D}$ be a locally persistent functor. Then F is distance non-increasing, in the sense that for all $x, y \in \mathscr{C}$ we have $d_I^{\mathscr{C}}(x, y) \geq d_I^{\mathscr{D}}(Fx, Fy)$.*

Proof. By functoriality of F, a δ-interleaving between $x, y \in \mathscr{C}$ maps to a δ-interleaving between $F(x), F(y) \in \mathscr{D}$. □

3.2 Category theory of locally persistent categories

In this section, we develop the category theory of locally persistent categories. One of the most useful tools of categorical reasoning are diagrams. For this reason, we start by extending the notion of a diagram of objects and morphisms in a category, to a diagram of objects and approximate morphisms in a locally persistent category.

We then identify three kinds of universal constructions that are particularly relevant when studying distances, and use them to prove properties about the interleaving distance. These are weighted pullbacks, weighted sequential limits, and terminal midpoints. Although we don't base any of our arguments on the theory of enriched categories, we show that the language of enriched category theory is useful in understanding these constructions, as they are all instances of weighted limits (Definition 2.4.14).

3.2.1 Diagrams

In this section, we describe how one can use diagrams to reason about locally persistent categories.

Notation 3.2.1. An ε-morphism in a locally persistent category can be drawn as an arrow with index ε, in the following way. For a locally persistent category $\mathscr{C}$, objects $x, y \in \mathscr{C}$, and $\varepsilon \in \mathbf{R}_+$, the notation $f : x \to_\varepsilon y$ means that f is an element of $\mathrm{Hom}_{\mathscr{C}}(x, y)_\varepsilon$. Note that we may write ε either as a superscript or as a subscript. Furthermore, to keep additional notation to a minimum, we will avoid explicitly writing the index 0 for 0-morphisms.

This notation can be extended to diagrams to get a well-defined notion of commutative diagram, as follows. Given a locally persistent category $\mathscr{C}$, objects $x, y, z \in \mathscr{C}$, $\varepsilon, \delta, \gamma \in \mathbf{R}_+$, and morphisms $f \in \mathrm{Hom}(x, y)_\varepsilon$, $g \in \mathrm{Hom}(y, z)_\delta$, and $h \in \mathrm{Hom}(x, z)_\gamma$, we say that the diagram

$$\begin{tikzcd} x \arrow[r, "h", "\gamma"'] \arrow[d, "f"', "\varepsilon"] & z \\ y \arrow[ur, "g", "\delta"'] & \end{tikzcd}$$

is **commutative** (or **commutes**) if $\mathbf{S}_{\gamma,\max(\gamma,\varepsilon+\delta)}(h) = \mathbf{S}_{\varepsilon+\delta,\max(\gamma,\varepsilon+\delta)}(g \circ f)$. In the case when $\delta + \varepsilon = \gamma$, this just means that $h = g \circ f$.

Note that the above notion of commutative diagram is well-defined even if $\varepsilon + \delta$ is not equal to γ. Next, we give a convention to avoid writing subscripts for the shift operation.

Notation 3.2.2. When it is clear from the context, we may omit the explicit shift of a morphism, so that, if $f : x \to_\varepsilon y$ and $\delta \geq \varepsilon$, the morphism $\mathbf{S}(f) : x \to_\delta y$ denotes $\mathbf{S}_{\varepsilon,\delta}(f)$.

Since we will use interleavings often, we introduce notation for them.

Notation 3.2.3. Given a locally persistent category $\mathscr{C}$, objects $x, y \in \mathscr{C}$, and $\varepsilon, \delta \in \mathbf{R}_+$, the notation

$$f : x\, {}_\delta\!\longleftrightarrow_\varepsilon\, y : g$$

means that there are morphisms $f : x \to_\varepsilon y$ and $g : y \to_\delta x$ forming an (ε,δ)-interleaving.

When labeling these morphisms in a diagram, we will use the convention that the upper label corresponds to the left-to-right morphism, and the lower label to the right-to-left morphism. So, if the interleaving above appears in a diagram, we will label it as follows

$$\begin{tikzcd} x \arrow[r, leftrightarrow, "f", "g"'] & y. \end{tikzcd}$$

3.2.2 Diagrams as functors

As is usual in category theory, a diagram in a locally persistent category $\mathscr{C}$ can be equivalently described by a locally persistent functor $\mathscr{D} \to \mathscr{C}$, for $\mathscr{D}$ an indexing locally persistent category. It is often useful to use diagrams indexed by a locally persistent category that is freely generated by a *locally persistent graph*. Let us make this formal.

A **locally persistent graph** G consists of a set $\mathsf{obj}(G)$ and, for each $x, y \in \mathsf{obj}(G)$, a persistent set $\mathrm{Arr}_G(x, y) \in \mathbf{Set}^{\mathbf{R}_+}$. A morphism $f : G \to H$ between locally persistent graphs consists of a map $f : \mathrm{obj}(G) \to \mathrm{obj}(H)$ together with, for every $x, y \in \mathrm{obj}(G)$, a natural transformation between persistent sets $f : \mathrm{Arr}_G(x, y) \to \mathrm{Arr}_H(f(x), f(y))$. Let **lpGph** be the category of locally persistent graphs.

Every locally persistent category can be seen as a locally persistent graph, by forgetting the composition and the identities. This gives a forgetful functor $U : \mathbf{lpCat} \to \mathbf{lpGph}$. This functor has a left adjoint $F : \mathbf{lpGph} \to \mathbf{lpCat}$, which is defined in [Wol74, Proposition 2.2] for a general enriching category $\mathscr{V}$. We describe F in our case.

Let $G \in \mathbf{lpGph}$. Let $F(G) \in \mathbf{lpCat}$ have the same objects as G and define

$$\mathrm{Hom}_{F(G)}(x, y) = \coprod_{\substack{e_0,\ldots,e_n \in G \\ e_0 = x, e_n = y}} \mathrm{Arr}_G(e_0, e_1) \otimes_{\mathsf{Day}} \mathrm{Arr}_G(e_1, e_2) \otimes_{\mathsf{Day}} \cdots \otimes_{\mathsf{Day}} \mathrm{Arr}_G(e_{n-1}, e_n),$$

when $x \neq y$ and

$$\mathrm{Hom}_{F(G)}(x, y) = \left(\coprod_{\substack{e_0,\ldots,e_n \in G \\ e_0 = x, e_n = y}} \mathrm{Arr}_G(e_0, e_1) \otimes_{\mathsf{Day}} \mathrm{Arr}_G(e_1, e_2) \otimes_{\mathsf{Day}} \cdots \otimes_{\mathsf{Day}} \mathrm{Arr}_G(e_{n-1}, e_n) \right) \coprod \mathscr{Y}(0),$$

when $x = y$. Here $\otimes_{\mathsf{Day}}$ is the Day convolution tensor product between objects of $\mathbf{Set}^{\mathbf{R}_+}$, defined in Section 2.6.1.

Identities are given by the $\mathscr{Y}(0)$ summand in the above equation, and composition is formal, using the fact that $\mathrm{Hom}_{F(G)}(x, z)$ contains a summand $\mathrm{Arr}_G(x, y) \otimes_{\mathsf{Day}} \mathrm{Arr}_G(y, z)$ for every $y \in G$.

Note that there is a canonical morphism $u_G : G \to U(F(G))$ given by mapping each object of G to itself and $\mathrm{Arr}_G(x, y)$ to the corresponding summand in $\mathrm{Arr}_{U(F(G))}(x, y)$.

The most important property of F is the following proposition, which tells us how to map out of a freely generated locally persistent category. The proposition is a particular case of [Wol74, Proposition 2.2].

Proposition 3.2.4. *The functor $F : \mathbf{lpGph} \to \mathbf{lpCat}$ is left adjoint to $U : \mathbf{lpCat} \to \mathbf{lpGph}$, with unit given by u. In particular, for every $G \in \mathbf{lpGph}$ and every $\mathscr{C} \in \mathbf{lpCat}$, there is a natural isomorphism*

$$\mathrm{Hom}_{\mathbf{lpGph}}(G, U(\mathscr{C})) \cong \mathrm{Hom}_{\mathbf{lpCat}}(F(G), \mathscr{C}). \qquad \square$$

3.2.3 (Co)powers

In this section we specialize the notion of copowers and powers (Definition 2.4.13) to locally persistent categories. Although we use enriched category theory to motivate the definition, the definition that we give (Definition 3.2.6) and that we use in the rest of this thesis does not depend on the notion of enrichment.

Recall that a locally persistent category is a category enriched in $\mathbf{Set}^{\mathbf{R}_+}$, where the monoidal structure of $\mathbf{Set}^{\mathbf{R}_+}$ is given by Day convolution. Recall from Definition 2.3.5 that given functors $F, G : \mathbf{R}_+ \to \mathbf{Set}$, the functor $[F, G] : \mathbf{R}_+ \to \mathbf{Set}$ denotes the internal hom from F to G.

According to Definition 2.4.13, a locally persistent category $\mathscr{C}$ is copowered if for every $F \in \mathbf{Set}^{\mathbf{R}_+}$ and every $x, y \in \mathscr{C}$, there is an object $F \cdot x$, and an isomorphism

$$\mathrm{Hom}_{\mathscr{C}}(F \cdot x, y) \cong [F, \mathrm{Hom}_{\mathscr{C}}(x, y)],$$

natural in x, y, and F. Dually, a power of y by F is an object y^F that satisfies

$$\mathrm{Hom}_{\mathscr{C}}(x, y^F) \cong [F, \mathrm{Hom}_{\mathscr{C}}(x, y)].$$

In practice, one may be interested in categories that are (co)powered only by a certain class of functors $\mathbf{R}_+ \to \mathbf{Set}$. Since locally persistent categories are categories enriched in a copresheaf category, we are especially interested in copowering and powering by representables. Recall from Definition 2.6.5 that given $\varepsilon \in \mathbf{R}_+$ we let $\mathscr{Y}(\varepsilon) \in \mathbf{Set}^{\mathbf{R}_+}$ denote its corresponding representable functor. Concretely, this functor behaves as follows: given $r \in \mathbf{R}_+$ we have

$$\mathscr{Y}(\varepsilon)(r) = \begin{cases} \emptyset & \text{if } r < \varepsilon \\ \{*\} & \text{if } r \geq \varepsilon, \end{cases}$$

with the only possible structure morphisms.

Notation 3.2.5. For simplicity, we denote $\mathscr{Y}(\varepsilon) \cdot x$ by $\varepsilon \cdot x$ and $x^{\mathscr{Y}(\varepsilon)}$ by x^{ε}.

By definition of copower, and the formula Eq. (2.6.7), if $\mathscr{C}$ is copowered by a representable $\mathscr{Y}(\varepsilon)$, we have, for every $r \in \mathbf{R}_+$,

$$\mathrm{Hom}_{\mathscr{C}}(\varepsilon \cdot x, y)_r \cong [\mathscr{Y}(\varepsilon), \mathrm{Hom}_{\mathscr{C}}(x, y)]_r \cong \mathrm{Nat}(\mathscr{Y}(\varepsilon) \otimes \mathscr{Y}(r), \mathrm{Hom}_{\mathscr{C}}(x, y)).$$

Since the Yoneda embedding is monoidal (Lemma 2.6.8), we have $\mathscr{Y}(\varepsilon)\otimes\mathscr{Y}(r)\cong\mathscr{Y}(\varepsilon+r)$. Also, by the Yoneda lemma, natural transformations $\mathscr{Y}(\varepsilon)\to F$ correspond to elements in $F(\varepsilon)$, so

$$\mathrm{Hom}_{\mathscr{C}}(\varepsilon\cdot x,y)_r\cong\mathrm{Nat}(\mathscr{Y}(\varepsilon)\otimes\mathscr{Y}(r),\mathrm{Hom}_{\mathscr{C}}(x,y))\cong\mathrm{Hom}_{\mathscr{C}}(x,y)_{\varepsilon+r},$$

and this is natural in ε, r, x, and y. Dually, if $\mathscr{C}$ is powered by representables,

$$\mathrm{Hom}_{\mathscr{C}}(x,y^{\varepsilon})_r\cong\mathrm{Hom}_{\mathscr{C}}(x,y)_{\varepsilon+r}.$$

We use this as our definition.

Definition 3.2.6. *Let $\mathscr{C}$ be a locally persistent category. We say that $\mathscr{C}$ **is copowered by representables** if for every $x,y\in\mathscr{C}$ and $\varepsilon\in\mathbf{R}_+$, there exists $\varepsilon\cdot x\in\mathscr{C}$, and an isomorphism of functors*

$$\mathrm{Hom}_{\mathscr{C}}(\varepsilon\cdot x,y)\cong\mathrm{Hom}_{\mathscr{C}}(x,y)_{\varepsilon+(-)},$$

natural in ε, x, and y.

*Dually, we say that $\mathscr{C}$ **is powered by representables** if for every $x,y\in\mathscr{C}$ and $\varepsilon\in\mathbf{R}_+$, there exists $y^{\varepsilon}\in\mathscr{C}$, and an isomorphism of functors*

$$\mathrm{Hom}_{\mathscr{C}}(x,y^{\varepsilon})\cong\mathrm{Hom}_{\mathscr{C}}(x,y)_{\varepsilon+(-)},$$

natural in ε, x, and y.

Equivalently, a locally persistent category $\mathscr{C}$ is powered by representables exactly if there is a natural isomorphism $\mathrm{Hom}_{\mathscr{C}}(x,y^{\varepsilon})\cong\mathrm{Hom}_{\mathscr{C}}(x,y)^{\varepsilon}$ for all $x,y\in\mathscr{C}$ and $\varepsilon\in\mathbf{R}_+$.

Categories of persistent objects are always copowered and powered by representables, as the following example shows.

Example 3.2.7. Let C be a category, and consider the locally persistent category $C^{\mathbf{R}}$ of persistent objects of C, as in Example 3.1.3.

Given $\varepsilon\in\mathbf{R}_+$, any persistent object $X\in C^{\mathbf{R}}$ can be shifted to the left and to the right by ε, by letting $X^{\varepsilon}(r)=X(r+\varepsilon)$ and $(\varepsilon\cdot X)(r)=X(r-\varepsilon)$ respectively. These shifts give a power and a copower of X respectively. This is because, in the locally persistent category $C^{\mathbf{R}}$, we have

$$\mathrm{Hom}_{C^{\mathbf{R}}}(\varepsilon\cdot X,Y)_{\delta}\cong\mathrm{Nat}(X((-)-\varepsilon),Y((-)+\delta))\cong\mathrm{Hom}_{C^{\mathbf{R}}}(X,Y)_{\varepsilon+\delta}\cong\mathrm{Hom}_{C^{\mathbf{R}}}(X,Y^{\varepsilon})_{\delta},$$

natural in X, Y, δ, and ε, by definition of the locally persistent structure.

Copowers and powers by representables are especially useful when working with ε-approximate morphisms for $\varepsilon > 0$, as we now explain. Fix a locally persistent category $\mathscr{C}$ that is copowered and powered by representables, $x, y \in \mathscr{C}$, and $\varepsilon \in \mathbf{R}_+$. There are natural bijections between $\mathrm{Hom}(x, y)_\varepsilon$, $\mathrm{Hom}(\varepsilon \cdot x, y)_0$, and $\mathrm{Hom}(x, y^\varepsilon)_0$. In this sense, working with ε-approximate morphisms for $\varepsilon > 0$ can be reduced to working with 0-morphisms.

We conclude this section by introducing a handy notation to work with copowers and powers by representables.

Notation 3.2.8. Since the isomorphisms

$$\mathrm{Hom}(x, y)_\varepsilon \cong \mathrm{Hom}(\varepsilon \cdot x, y)_0 \cong \mathrm{Hom}(x, y^\varepsilon)_0$$

are natural in x, y, and ε, we will often use them implicitly, so that, for a morphism $f : x \to_\varepsilon y$, the corresponding morphisms $\varepsilon \cdot x \to y$ and $x \to y^\varepsilon$ will also be denoted by f.

Finally, if no confusion can arise, we may sometimes omit copowers and powers of morphisms, as follows.

Notation 3.2.9. Given $x, y \in \mathscr{C}$, $f : x \to_\varepsilon y$, and $\delta \in \mathbf{R}_+$, if no confusion can arise, we may denote the morphisms $\delta \cdot f : \delta \cdot x \to_\varepsilon \delta \cdot y$ and $f^\delta : x^\delta \to_\varepsilon y^\delta$, given by the functorial action of copowers and powers by representables, by $f : \delta \cdot x \to_\varepsilon \delta \cdot y$ and $f : x^\delta \to_\varepsilon y^\delta$, respectively.

3.2.4 Weighted pullbacks

In this section, we introduce a universal construction that lets us talk about pullbacks of a 0-morphism along an approximate morphism. We start by specializing the notion of commutative diagram in a locally persistent category, Notation 3.2.1, to the case of squares.

Let $\mathscr{C}$ be a locally persistent category, and assume given a diagram in $\mathscr{C}$ of the following form

$$
\begin{tikzcd}
p \arrow[r, "f", "\varepsilon"'] \arrow[d, "j"'] & b \arrow[d, "k"] \\
a \arrow[r, "h", "\varepsilon"'] & c.
\end{tikzcd}
$$

Recall that, according to our notation, j and k are 0-morphisms, and that the above diagram is commutative if $k \circ f = h \circ j$.

We now define pullbacks of 0-morphisms along approximate morphisms, and prove some basic properties of this construction.

It is interesting to note that essentially the same definition would work to define pullbacks of approximate morphisms along approximate morphisms. The reason why we don't state it in this generality is because we won't make use of it.

Definition 3.2.10. *Let $\mathscr{C}$ be a locally persistent category. A diagram of the form*

$$
\begin{tikzcd}
p \arrow[r, "f", "\varepsilon"'] \arrow[d, "j"'] & b \arrow[d, "k"] \\
a \arrow[r, "h", "\varepsilon"'] & c
\end{tikzcd}
$$

*is a **weighted pullback of h and k** if it satisfies the following universal property. For every object $p' \in \mathscr{C}$, $\gamma \in \mathbf{R}_+$, and morphisms $j' : p' \to_\gamma a$ and $f' : p' \to_{\varepsilon+\gamma} b$ making the following diagram commute, there exists a unique morphism $u : p' \to_\gamma p$ completing the diagram*

$$
\begin{tikzcd}
p' \arrow[drr, bend left, "f'", "\varepsilon+\gamma"'] \arrow[ddr, bend right, "j'"', "\gamma"] \arrow[dr, dashed, "u", "\gamma"'] & & \\
& p \arrow[r, "f", "\varepsilon"'] \arrow[d, "j"'] & b \arrow[d, "k"] \\
& a \arrow[r, "h", "\varepsilon"'] & c.
\end{tikzcd}
$$

In the situation of the previous definition, we refer to j as the **weighed pullback of k along h**. Dually, one defines weighted pushouts.

A very basic, yet very useful, result in category theory says that the pullback of an isomorphism along any morphism is also an isomorphism. A similar statement holds for locally persistent categories, if we replace isomorphism by interleaving.

Proposition 3.2.11. *Let $\mathscr{C}$ be a locally persistent category, and assume given a weighted pullback*

$$\begin{tikzcd}
p \arrow[r, "f", "{}_{\varepsilon}"' near end] \arrow[d, "j"'] & b \arrow[d, "k"] \\
a \arrow[r, "h", "{}_{\varepsilon}"' near end] & c.
\end{tikzcd}$$

If there exists $\delta \in \mathbf{R}_+$ and a morphism $i : c \to_\delta a$ forming an (ε,δ)-interleaving $h : a\,{}_{\delta}\!\longleftrightarrow_{\varepsilon} c : i$, then there exists a unique morphism $g : b \to_\delta p$ forming an interleaving $f : p\,{}_{\delta}\!\longleftrightarrow_{\varepsilon} b : g$ and rendering the following square commutative

$$\begin{tikzcd}
p \arrow[r, leftrightarrow, "f", "{}_{\delta}\; g\; {}_{\varepsilon}"'] \arrow[d, "j"'] & b \arrow[d, "k"] \\
a \arrow[r, leftrightarrow, "h", "{}_{\delta}\; i\; {}_{\varepsilon}"'] & c.
\end{tikzcd}$$

Proof. By the universal property, we have a unique morphism $g : b \to_\delta p$ rendering the following diagram commutative

$$\begin{tikzcd}
b \arrow[dr, dashed, "g", "{}_{\delta}"' near end] \arrow[drr, bend left, "\mathbf{S}_{0,\varepsilon+\delta}(\mathrm{id}_b)", "{}^{\varepsilon+\delta}"' near end] \arrow[ddr, bend right, "i \circ k"', "{}_{\delta}" near end] & & \\
& p \arrow[r, "f", "{}_{\varepsilon}"' near end] \arrow[d, "j"'] & b \arrow[d, "k"] \\
& a \arrow[r, "h", "{}_{\varepsilon}"' near end] & c.
\end{tikzcd}$$

This establishes uniqueness. We must show that $f : p\,{}_{\delta}\!\longleftrightarrow_{\varepsilon} b : g$ is in fact an interleaving. By the commutativity of the above diagram, all that remains to be shown is that

$g \circ f = \mathbf{S}_{0,\varepsilon+\delta}(\mathrm{id}_p)$. To see this, consider the following commutative diagram.

$$\begin{tikzcd}
p \arrow[dr, "f", "\varepsilon"' near end] \arrow[ddrrr, bend left, "\mathbf{S}_{\varepsilon,2\varepsilon+\delta}(f)", "2\varepsilon+\delta"' near end] \arrow[dddrr, bend right, "\mathbf{S}_{0,\varepsilon+\delta}(j)"', "\varepsilon+\delta" near end] & & & \\
& b \arrow[dr, "g", "\delta"' near end] \arrow[drr, bend left, "\mathbf{S}(\mathrm{id})", "\varepsilon+\delta"' near end] \arrow[ddr, bend right, "i \circ k"', "\delta" near end] & & \\
& & p \arrow[r, "f", "\varepsilon"' near end] \arrow[d, "j"'] & b \arrow[d, "k"] \\
& & a \arrow[r, "h", "\varepsilon"' near end] & c.
\end{tikzcd}$$

The uniqueness part of the universal property of p implies that $g \circ f = \mathbf{S}_{0,\varepsilon+\delta}(\mathrm{id}_p)$, since $\mathbf{S}_{0,\varepsilon+\delta}(\mathrm{id}_p)$ renders the diagram commutative too. □

We conclude this section by giving an explicit construction of weighted pullbacks.

Proposition 3.2.12. *Let $\mathscr{C}$ be a locally persistent category that is powered by representables and such that pullbacks exist in its underlying category. If powers by representables preserve these pullbacks, then weighted pullbacks along approximate morphisms exist in $\mathscr{C}$.*

In particular, if $\mathscr{C}$ is copowered and powered by representables, and pullbacks exist in $\mathscr{C}_0$, then $\mathscr{C}$ admits weighted pullbacks.

In the two cases above, the pullback of $k : b \to c$ along $h : a \to_\varepsilon c$ is computed as the pullback of $k : b^\varepsilon \to c^\varepsilon$ along $h : a \to c^\varepsilon$.

Proof. The second statement follows from the first one by noting that, if $\mathscr{C}$ is copowered and powered by representables, then powers preserve all limits that exist in the underlying category of $\mathscr{C}$, since, in this case, powers are right adjoints.

We now prove the first statement. The statement about the construction of weighted pullbacks will be clear by construction. Consider the following pullback

$$\begin{tikzcd}
p \arrow[r, "f"] \arrow[d, "j"'] & b^\varepsilon \arrow[d, "k"] \\
a \arrow[r, "h"] & c^\varepsilon.
\end{tikzcd}$$

This gives a commutative square

$$\begin{tikzcd}
p \arrow[r, "f"] \arrow[d, "j"'] & b \arrow[d, "k"] \\
a \arrow[r, "h"] & c.
\end{tikzcd}$$

We now proceed to prove its universal property. Given a commutative square

$$\begin{tikzcd}
p' \arrow[r, "f'"] \arrow[d, "j'"'] & b \arrow[d, "k"] \\
a \arrow[r, "h"] & c,
\end{tikzcd}$$

we use the universal property of p^γ to obtain a unique dashed morphism rendering the following diagram commutative

$$\begin{tikzcd}
p' \arrow[dr, dashed, "u"] \arrow[drr, bend left, "f'"] \arrow[ddr, bend right, "j'"'] & & \\
& p^\gamma \arrow[r, "f^\gamma"] \arrow[d, "j"'] & b^{\varepsilon+\gamma} \arrow[d, "k^{\varepsilon+\gamma}"] \\
& a^\gamma \arrow[r, "h^\gamma"] & c^{\varepsilon+\gamma}.
\end{tikzcd}$$

Note that p^γ is a pullback by the assumption that powers preserve all pullbacks of the underlying category of $\mathscr{C}$. This gives the unique $u : p' \to_\gamma p$ rendering the following diagram commutative, concluding the proof.

$$
\begin{tikzcd}
p' \arrow[drr, bend left, "f'", "\varepsilon+\gamma"'] \arrow[dr, dashed, "u", "\gamma"'] \arrow[ddr, bend right, "j'"', "\gamma"] & & \\
& p \arrow[r, "f", "\varepsilon"'] \arrow[d, "j"'] & b \arrow[d, "k"] \\
& a \arrow[r, "h", "\varepsilon"'] & c.
\end{tikzcd}
$$

□

3.2.5 Weighted sequential limits

We now turn our attention to a universal construction that will let us prove (metric) completeness of certain interleaving distances. This universal construction is a straightforward generalization of the notion of sequential limit in usual (**Set**-enriched) categories.

Definition 3.2.13. *Let $\mathscr{C}$ be a locally persistent category. Assume given $\varepsilon \in \mathbf{R}_+$ and $\varepsilon_i \in \mathbf{R}_+$ for each $i \in \mathbb{N}$ such that $\sum_i \varepsilon_i = \varepsilon$, objects $x_i \in \mathscr{C}$ for each $i \in \mathbb{N}$, and morphisms $f_i : x_{i+1} \to_{\varepsilon_i} x_i$ for each $i \in \mathbb{N}$. We depict this situation as follows:*

$$\cdots \xrightarrow{f_i}_{\varepsilon_i} x_i \xrightarrow{f_{i-1}}_{\varepsilon_{i-1}} \cdots \xrightarrow{f_1}_{\varepsilon_1} x_1 \xrightarrow{f_0}_{\varepsilon_0} x_0.$$

*Let $\overline{\varepsilon}_i = \varepsilon - \sum_{j<i} \varepsilon_j$. A **weighted sequential limit** of the above diagram is given by an object $l \in \mathscr{C}$ and morphisms $g_i : l \to_{\overline{\varepsilon}_i} x_i$, such that $f_i \circ g_{i+1} = g_i$, and satisfying the following universal property.*

For any object $l' \in \mathscr{C}$, $\gamma \in \mathbf{R}_+$, and morphisms $g'_i : l' \to_{\overline{\varepsilon}_i+\gamma} x_i$, such that $f_i \circ g'_{i+1} = g'_i$, there exists a unique morphism $u : l' \to_\gamma l$ rendering the following diagram commutative

$$
\begin{tikzcd}
l' \arrow[d, dashed, "u", "\gamma"'] \arrow[ddr, bend left, "g'_i", "\overline{\varepsilon}_i+\gamma"'] \arrow[ddrrr, bend left, "g'_1", "\overline{\varepsilon}_1+\gamma"'] \arrow[ddrrrr, bend left, "g'_0", "\overline{\varepsilon}_0+\gamma"'] & & & & \\
l \arrow[dr, "g_i"', "\overline{\varepsilon}_i"] \arrow[drrr, "g_1"', "\overline{\varepsilon}_1"] \arrow[drrrr, "g_0", "\overline{\varepsilon}_0"'] & & & & \\
\cdots \arrow[r, "f_i"', "\varepsilon_i"] & x_i \arrow[r, "f_{i-1}"', "\varepsilon_{i-1}"] & \cdots \arrow[r, "f_1"', "\varepsilon_1"] & x_1 \arrow[r, "f_0"', "\varepsilon_0"] & x_0.
\end{tikzcd}
$$

Analogously to the case of weighted pullbacks, we recall from classical category theory that a sequential limit of a diagram where each of the morphisms is an isomorphism induces an isomorphism between the limit of the diagram and each of the objects of the diagram, that is, isomorphisms are closed under sequential limits. Informally, the following result says that a weighted sequential limit of a diagram where each of the morphisms forms part of an interleaving induces an interleaving between the limit an each of the objects of the diagram. Its proof is analogous to the proof of Proposition 3.2.11.

Proposition 3.2.14. *Let $\mathscr{C}$ be a locally persistent category and let*

$$\begin{tikzcd}
l \arrow[drr, bend left, "g_i", "\overline{\varepsilon_i}"' near end] \arrow[drrrr, bend left, "g_1", "\overline{\varepsilon_1}"' near end] \arrow[drrrrr, bend left, "g_0", "\overline{\varepsilon_0}"' near end] & & & & & \\
& \cdots \arrow[r, "f_i"', "\varepsilon_i" near end] & x_i \arrow[r, "f_{i-1}"', "\varepsilon_{i-1}" near end] & \cdots \arrow[r, "f_1"', "\varepsilon_1" near end] & x_1 \arrow[r, "f_0"', "\varepsilon_0" near end] & x_0 .
\end{tikzcd}$$

be a weighted sequential limit.

Consider $\delta \in \mathbf{R}_+$ and $\delta_i \in \mathbf{R}_+$ for each $i \in \mathbb{N}$ such that $\sum_i \delta_i = \delta$, and let $\overline{\delta_i} = \sum_{j<i} \delta_j$. If there are morphisms $h_i : x_i \to_{\delta_i} x_{i+1}$ such that $f_i : x_{i+1}\,{}_{\delta_i}\!\longleftrightarrow_{\varepsilon_i} x_i : h_i$ forms an $(\varepsilon_i, \delta_i)$-interleaving for every $i \in \mathbb{N}$, then there exist unique morphisms $k_i : x_i \to_{\overline{\delta_i}} l$ such that $g_i : l\,{}_{\overline{\delta_i}}\!\longleftrightarrow_{\overline{\varepsilon_i}} x_i : k_i$ forms an interleaving for every $i \in \mathbb{N}$ and such that following diagram is commutative:

$$\begin{tikzcd}
l \arrow[drr, leftrightarrow, bend left, "g_i", "k_i"'] \arrow[drrrr, leftrightarrow, bend left, "g_1", "k_1"'] \arrow[drrrrr, leftrightarrow, bend left, "g_0", "k_0"'] & & & & & \\
& \cdots \arrow[r, leftrightarrow, "h_i", "f_i"'] & x_i \arrow[r, leftrightarrow, "h_{i-1}", "f_{i-1}"'] & \cdots \arrow[r, leftrightarrow, "h_1", "f_1"'] & x_1 \arrow[r, leftrightarrow, "h_0", "f_0"'] & x_0 .
\end{tikzcd}$$

□

We conclude this section by giving sufficient conditions for the existence of weighted sequential limits. The proof of the following result is analogous to the proof of Proposition 3.2.12.

Proposition 3.2.15. *Let $\mathscr{C}$ be a locally persistent category that is powered by representables and such that sequential limits exist in its underlying category. If powers by representables preserve these sequential limits, then weighted sequential limits exist in $\mathscr{C}$.*

In particular, if $\mathscr{C}$ is copowered and powered by representables, and sequential limits exist in $\mathscr{C}_0$, then $\mathscr{C}$ admits weighted sequential limits.

In the two cases above, the weighted sequential limit of a diagram

$$\cdots \xrightarrow{f_i}_{\varepsilon_i} x_i \xrightarrow{f_{i-1}}_{\varepsilon_{i-1}} \cdots \xrightarrow{f_1}_{\varepsilon_1} x_1 \xrightarrow{f_0}_{\varepsilon_0} x_0$$

is computed as the limit of the diagram

$$\cdots \xrightarrow{f_i} x_i^{\bar{\varepsilon}_i} \xrightarrow{f_{i-1}} \cdots \xrightarrow{f_1} x_1^{\bar{\varepsilon}_1} \xrightarrow{f_0} x_0^{\bar{\varepsilon}_0}$$

□

3.2.6 Midpoints

We now present a third universal construction. This one is particularly relevant when proving that an interleaving distance is intrinsic or geodesic. Informally, it defines a notion of "best" or, more precisely, universal midpoint of an interleaving.

Definition 3.2.16. *Let $\varepsilon, \gamma, \delta \in \mathbf{R}_+$ be such that $\gamma + \delta = \varepsilon$. A **terminal (γ, δ)-midpoint** of an interleaving $f : x\ {}_{\varepsilon}\!\longleftrightarrow_{\varepsilon}\ y : g$ consists of an object z and morphisms $a : z \to_{\gamma} x$ and $b : z \to_{\delta} y$ rendering the following diagram commutative*

$$\begin{tikzcd} & z \arrow[dl, "a"', "\gamma"] \arrow[dr, "b", "\delta"'] & \\ x \arrow[rr, "f", "\varepsilon"'] & & y, \arrow[ll, "g", "\varepsilon"'] \end{tikzcd}$$

and satisfying the following universal property. For every object z', $\alpha \in \mathbf{R}_+$, and morphisms $a' : z' \to_{\gamma+\alpha} x$ and $b' : z' \to_{\delta+\alpha} y$ making the following diagram commute, there *exists a unique morphism* $u : z' \to_\alpha z$ *completing the diagram*

$$\begin{tikzcd}
 & z' \arrow[d, dashed, "u", "\alpha"'] \arrow[ddl, bend right, "a'"', "\gamma+\alpha"] \arrow[ddr, bend left, "b'", "\delta+\alpha"'] & \\
 & z \arrow[dl, "a"', "\gamma"] \arrow[dr, "b", "\delta"'] & \\
 x \arrow[rr, shift left, "f", "\varepsilon"'] & & y. \arrow[ll, shift left, "g", "\varepsilon"']
\end{tikzcd}$$

Dually, one has the definition of an initial midpoint.

From the universal property of terminal midpoints, it follows that they are midpoints in the sense that the interleaving factors through them, as we now show.

Proposition 3.2.17. *Let* $\varepsilon, \gamma, \delta \in \mathbf{R}_+$ *be such that* $\gamma + \delta = \varepsilon$. *Given a terminal* (γ, δ)*-midpoint for an interleaving* $f : x \,{}_\varepsilon\!\!\longleftrightarrow_\varepsilon y : g$,

$$\begin{tikzcd}
 & z \arrow[dl, "a"', "\gamma"] \arrow[dr, "b", "\delta"'] & \\
 x \arrow[rr, shift left, "f", "\varepsilon"'] & & y, \arrow[ll, shift left, "g", "\varepsilon"']
\end{tikzcd}$$

there exist unique morphisms $c : x \to_\gamma z$ *and* $d : y \to_\delta z$ *forming interleavings with* a *and* b *respectively, and rendering the following diagram commutative*

$$\begin{tikzcd}
 & z \arrow[dl, shift right, "a"', "\gamma"] \arrow[dr, shift left, "b", "\delta"'] & \\
 x \arrow[ur, shift right, "c"', "\gamma"] \arrow[rr, shift left, "f", "\varepsilon"'] & & y. \arrow[ul, shift left, "d", "\delta"'] \arrow[ll, shift left, "g", "\varepsilon"']
\end{tikzcd}$$

Proof. Let us start by constructing $c : x \to_\gamma z$. We do this using the universal property of z, as follows.

$$\begin{tikzcd}
& x \arrow[d, dashed, "c", "\gamma"'] \arrow[ddl, bend right, "\mathbf{S}(\mathrm{id}_x)"', "2\gamma"] \arrow[ddr, bend left, "f", "\gamma+\delta"'] & \\
& z \arrow[dl, "a"', "\gamma"] \arrow[dr, "b", "\delta"'] & \\
x \arrow[rr, leftrightarrow, "f", "g"', "\varepsilon"{near start, below}, "\varepsilon"{near end, below}] & & y.
\end{tikzcd}$$

This also shows that if a morphism c as in the statement exists, it must be unique. To prove that c and a form an interleaving, it remains to be shown that $\mathbf{S}(\mathrm{id}_z) = c \circ a$. We do this in the usual way, using the uniqueness part of the universal property of z. Concretely, consider the diagram

$$\begin{tikzcd}
& z \arrow[d, "2\gamma"'] \arrow[ddl, bend right, "\mathbf{S}(c)"', "3\gamma"] \arrow[ddr, bend left, "\mathbf{S}(f)", "2\gamma+\delta"'] & \\
& z \arrow[dl, "a"', "\gamma"] \arrow[dr, "b", "\delta"'] & \\
x \arrow[rr, leftrightarrow, "f", "g"', "\varepsilon"{near start, below}, "\varepsilon"{near end, below}] & & y.
\end{tikzcd}$$

Choosing the middle vertical map to be $\mathbf{S}_{0,2\gamma}(\mathrm{id}_z)$ or $c \circ a$ renders the diagram commutative, so, by uniqueness, we must have $\mathbf{S}_{0,2\gamma}(\mathrm{id}_z) = c \circ a$.

The construction of d is symmetrical, and the fact that the triangle in the statement commutes follows by definition of d and c. □

The universal property of terminal midpoints also implies that the above factorization is universal, in the sense that we can compose the factorizations through different midpoints. As we shall see, this implies that, in a locally persistent category that admits terminal midpoints, the interleaving distance is intrinsic.

Proposition 3.2.18. *Let $\varepsilon, \gamma, \delta, \gamma', \delta' \in \mathbf{R}_+$ be such that $\gamma + \delta = \varepsilon$, $\gamma' + \delta' = \varepsilon$, $\gamma \leq \gamma'$, and $\delta \geq \delta'$, and let $f : x \;{}_{\varepsilon}\!\longleftrightarrow_{\varepsilon}\; y : g$ be an interleaving.*

Given a terminal (γ, δ)-midpoint z and a terminal (γ', δ')-midpoint z', there are *unique interleavings such that the following diagram commutes*

$$\begin{tikzcd}
x \arrow[r, leftrightarrow, "c", "a"'] & z \arrow[r, leftrightarrow, "j", "k"'] \arrow[rr, leftrightarrow, bend right, "b", "d"'] & z' \arrow[ll, leftrightarrow, bend right, "c'"', "a'"] \arrow[r, leftrightarrow, "b'", "d'"'] & y,
\end{tikzcd}$$

and the composite of the horizontal interleavings is $f : x \,{}_{\varepsilon}\!\longleftrightarrow_{\varepsilon}\, y : g$. *Here the interleavings* $c : x \,{}_{\gamma}\!\longleftrightarrow_{\gamma}\, z : a$, $b : z \,{}_{\delta}\!\longleftrightarrow_{\delta}\, y : d$, $c' : x \,{}_{\gamma'}\!\longleftrightarrow_{\gamma'}\, z' : a'$, *and* $b' : z' \,{}_{\delta'}\!\longleftrightarrow_{\delta'}\, y : d'$ *are the interleavings of Proposition 3.2.17.*

Proof. Since $\delta + \gamma' - \gamma \geq \delta'$, we can form a diagram

$$\begin{tikzcd}
& z' \arrow[d, dashed, "k", "\gamma'-\gamma"'] \arrow[ddl, bend right, "a'"'] \arrow[ddr, bend left, "\mathbf{S}(b')"] & \\
& z \arrow[dl, "a"'] \arrow[dr, "b"] & \\
x \arrow[rr, leftrightarrow, "f", "g"'] & & y,
\end{tikzcd}$$

which defines the morphism k. The morphism $j : z \to_{\gamma'-\gamma} z'$ is defined symmetrically.

All the required commutativities follow from the uniqueness part of the universal properties of z and z'. $\square$

We conclude this section by providing sufficient conditions for the existence of terminal midpoints. This construction is essentially the construction of one parameter families associated to an interleaving given in [CSGO16, Section 3.4].

Proposition 3.2.19. *Let* $\mathscr{C}$ *be a locally persistent category that is powered by representables, and such that binary products and pullbacks exist in its underlying category. If powers by representables preserve these limits, then* $\mathscr{C}$ *admits terminal midpoints.*

In particular, if $\mathscr{C}$ *is copowered and powered by representables, and binary products and pullbacks exist in* $\mathscr{C}_0$, *then* $\mathscr{C}$ *admits terminal midpoints.*

In the two cases above, the terminal (γ, δ)*-midpoint of an interleaving* $f : x \,{}_{\varepsilon}\!\longleftrightarrow_{\varepsilon}\, y : g$ *is computed by taking the pullback of the following diagram in the underlying* *category of $\mathscr{C}$:*

$$\begin{tikzcd}
 & y^{\delta} \arrow[d, "{(g, \mathbf{S}(\mathrm{id}))}"] \\
x^{\gamma} \arrow[r, "{(\mathbf{S}(\mathrm{id}), f)}"'] & x^{\delta+\varepsilon} \times y^{\gamma+\varepsilon}.
\end{tikzcd}$$

Proof. The second statement follows, as usual, from the first one, by noting that the existence of copowers by representables implies that powers by representables preserve all limits that exist in $\mathscr{C}_0$.

Let $\varepsilon, \gamma, \delta \in \mathbf{R}_+$ be such that $\gamma + \delta = \varepsilon$, and let $f : x \,{}_{\varepsilon}\!\!\longleftrightarrow_{\varepsilon}\, y : g$ be an interleaving. Let $z \in \mathscr{C}$ and consider a diagram

$$\begin{tikzcd}
z \arrow[r, "b"] \arrow[d, "a"'] & y^{\delta} \arrow[d, "{(g, \mathbf{S}(\mathrm{id}))}"] \\
x^{\gamma} \arrow[r, "{(\mathbf{S}(\mathrm{id}), f)}"'] & x^{\delta+\varepsilon} \times y^{\gamma+\varepsilon}.
\end{tikzcd}$$

Clearly, the above diagram commutes if and only if the following diagram does:

$$\begin{tikzcd}
 & z \arrow[dl, "a"', "\gamma" near end] \arrow[dr, "b", "\delta"' near end] & \\
x \arrow[rr, leftrightarrow, "f", "g"', "\varepsilon" near start, "\varepsilon"' near end] & & y.
\end{tikzcd}$$

This means that, if the first diagram is a pullback, then the second one must be a terminal (γ, δ)-midpoint diagram. This is because, given $z' \in \mathscr{C}$, a morphism $z' \to_{\alpha} z$ is equivalently given by a cone from z' to the first diagram powered by α, by the assumption that powers preserve binary products and pullbacks. □

3.2.7 Interpretation as weighted (co)limits

In this short, optional section we interpret the three universal constructions introduced in previous sections, namely weighted pullbacks, weighted sequential limits, and terminal midpoints, as weighted limits (Definition 2.4.14). We start with weighted pullbacks.

Consider the locally persistent category K freely generated (Section 3.2.2) by the diagram

$$\begin{tikzcd} & y \arrow[d] \\ x \arrow[r, "\varepsilon"'] & z. \end{tikzcd}$$

Consider further the weight $W : K \to \mathbf{Set}^{\mathbf{R}_+}$ given by

$$\begin{tikzcd} & \mathscr{Y}(\varepsilon) \arrow[d] \\ \mathscr{Y}(0) \arrow[r, "\varepsilon"'] & \mathscr{Y}(\varepsilon). \end{tikzcd}$$

A diagram in a locally persistent category $\mathscr{C}$ of the form

$$\begin{tikzcd} & b \arrow[d, "k"] \\ a \arrow[r, "h", "\varepsilon"'] & c \end{tikzcd}$$

is given by an enriched functor $F : K \to \mathscr{C}$, and its weighted pullback, in the sense of Definition 3.2.10, is precisely its weighted limit $\lim^W F$, in the sense of Definition 2.4.14.

As noted earlier, the definition of weighted pullback of a 0-morphism along an approximate morphism can be generalized further to allow pullbacks of approximate morphisms along approximate morphisms. The above description gives the recipe for the general case.

Weighted sequential limits are also special cases of weighted limits. We give the relevant constructions here. Consider the locally persistent category K freely generated by the diagram

$$\cdots \xrightarrow[\varepsilon_i]{f_i} x_i \xrightarrow[\varepsilon_{i-1}]{f_{i-1}} \cdots \xrightarrow[\varepsilon_1]{f_1} x_1 \xrightarrow[\varepsilon_0]{f_0} x_0.$$

Consider also the weight $W : K \to \mathbf{Set}^{\mathbf{R}_+}$ given by

$$\cdots \to_{\varepsilon_i} \mathscr{Y}(\overline{\varepsilon}_i) \to_{\varepsilon_{i-1}} \cdots \to_{\varepsilon_1} \mathscr{Y}(\overline{\varepsilon}_1) \to_{\varepsilon_0} \mathscr{Y}(\overline{\varepsilon}_0).$$

Then, the weighted sequential limit of a sequential diagram $F : K \to \mathscr{C}$ is precisely the weighted limit $\lim^W F$.

Finally, terminal midpoints are special cases of weighted limits too. Let K be the locally persistent category given by the diagram

$$f : x \,{}_{\varepsilon}\!\longleftrightarrow_{\varepsilon}\, y : g,$$

that is, the unique locally persistent category with set of objects $\{x, y\}$ and such that $\mathrm{Hom}(x, x) = \mathrm{Hom}(y, y) = \mathscr{Y}(0)$ and $\mathrm{Hom}(x, y) = \mathrm{Hom}(y, x) = \mathscr{Y}(\varepsilon)$. Consider further the weight $W : K \to \mathbf{Set}^{\mathbf{R}_+}$ given by

$$\mathscr{Y}(\gamma) \,{}_{\varepsilon}\!\longleftrightarrow_{\varepsilon}\, \mathscr{Y}(\delta).$$

Then, the terminal (γ, δ)-midpoint of an interleaving $F : K \to \mathscr{C}$ is precisely the weighted limit $\lim^W F$.

After these interpretations, a reader with some familiarity with enriched category theory may wonder why we need a special kind of limit and powers that respect that kind of limit to deduce that we have the weighted version of that kind of limit (for example, in Proposition 3.2.12, Proposition 3.2.15, and Proposition 3.2.19). The reason is that, in general, in order to have weighted limits, one needs powers and conical limits. But, if an enriching category $\mathscr{V}$ is not conservative (i.e. the functor $\mathrm{Hom}(1, -) : \mathscr{V} \to \mathbf{Set}$ does not reflect isomorphisms), the existence of conical limits in a $\mathscr{V}$-enriched category does not follow from the existence of limits in the underlying category ([Kel82, Section 3.1]). As a matter of fact, in order to get conical limits from limits in the underlying category, it is enough for the powers to respect the conical limits ([Kel82, Section 3]).

3.3 Relative locally persistent categories

In this section, we consider quotients of interleaving distances. The notion of quotient of an ep metric that we use is the one from Section 2.2.3.

Definition 3.3.1. *Let $\mathscr{C}$ be a locally persistent category and let R be an equivalence relation on the objects of $\mathscr{C}$. Define the* ***quotient interleaving distance*** $(d_I^{\mathscr{C}})_{/R}$*, an ep metric on* $\mathsf{obj}(\mathscr{C})$*, as the quotient of the interleaving distance $d_I^{\mathscr{C}}$ by the equivalence relation R. When no confusion can arise, we sometimes denote it by $d_{QI}^{\mathscr{C}}$.*

Quotient interleaving distances are better behaved when the equivalence relation comes from a notion of acyclic morphism, as this lets us use categorical arguments to prove metric properties.

Definition 3.3.2. *A **relative category** $(C, \mathscr{W})$ is given by a category C together with a class of morphisms $\mathscr{W}$ of C that is closed under composition and contains all identities. The morphisms in the class $\mathscr{W}$ are called **acyclic morphisms**.*

We choose "acyclic morphism" over "weak equivalence" to make it clear that these morphisms need not be the weak equivalences of a model structure.

Definition 3.3.3. *Let $(C, \mathscr{W})$ be a relative category and $x, y \in C$. A **zig-zag** between x and y is given by a finite sequence of morphisms in $\mathscr{W}$ of any of the following forms:*

$$x \to z_1 \leftarrow z_2 \to \cdots \leftarrow z_n \to y,$$
$$x \to z_1 \leftarrow z_2 \to \cdots \to z_n \leftarrow y,$$
$$x \leftarrow z_1 \to z_2 \leftarrow \cdots \leftarrow z_n \to y,$$
$$x \leftarrow z_1 \to z_2 \leftarrow \cdots \to z_n \leftarrow y.$$

Definition 3.3.4. *Given a relative category $(C, \mathscr{W})$ and objects $x, y \in C$, we say that x and y are **weakly equivalent** if they are connected by a zig-zag of acyclic morphisms. In that case, we write $x \simeq_{\mathscr{W}} y$, or $x \simeq y$ if there is no risk of confusion.*

Note that being weakly equivalent is an equivalence relation. We can then use any class of acyclic morphisms in a locally persistent category to define a quotient interleaving distance.

Definition 3.3.5. *A **relative locally persistent category** is given by a locally persistent category $\mathscr{C}$ and a class of 0-morphisms $\mathscr{W}$ such that $(\mathscr{C}_0, \mathscr{W})$ is a relative category. The morphisms in $\mathscr{W}$ are called **acyclic morphisms**.*

Definition 3.3.6. *The **quotient interleaving distance** of a relative locally persistent category $(\mathscr{C}, \mathscr{W})$ is the quotient ep metric obtained by taking the quotient of the interleaving distance of $\mathscr{C}$ by the equivalence relation given by being weakly equivalent. We denote it by $(d_I^{\mathscr{C}})_{/\mathscr{W}}$, or by $(d_I^{\mathscr{C}})_{/\simeq}$, when no confusion can arise.*

Chapter 4

Metric properties of interleaving distances

In this chapter, we show that categorical structure in a locally persistent category can give rise to useful metric structure of its interleaving distance. We give stability results for functors between locally persistent categories (Section 4.2) and conditions under which a quotient interleaving distance is complete (Theorem 4.3.3), intrinsic (Corollary 4.4.5), or geodesic (Section 4.5).

We also give conditions under which a quotient interleaving distance can be computed as an infimum over interleavings. The characterization of an interleaving distance as an infimum over interleavings (as in Definition 3.1.11) is lost for general quotients, but a similar characterization can be recovered under mild hypothesis on the class of acyclic morphisms (Theorem 4.1.4).

In Section 4.6, we study distance non-increasing maps from a metric space into a locally persistent category endowed with its interleaving distance, following the methodology of [BSN17] and generalizing their techniques to the context of locally persistent categories. We also show that ep metric spaces form a full subcategory of the category of locally persistent categories.

In Section 4.7, we define the observable locally persistent category of a locally persistent category, generalizing the methodology of [CCBS14] to locally persistent categories. This construction defines a metrically equivalent locally persistent category such that all of its hom-persistent sets are right continuous (Section 2.6.2).

Completeness of the interleaving distance in the context of categories with a flow ([SMS18]) was studied in [Cru19]; we relate our completeness to their result, and apply some of our results to categories with a flow in Section 5.2.4. Proofs that a certain interleaving distance is complete, intrinsic, or geodesic have been given in many particular examples ([CSEH05], [BSN17], [BV18], [CSGO16]). The techniques that we use to prove the results in this section share many similarities with the techniques used in those references. The main difference is that our results are general results about locally persistent categories that apply in many examples.

The theorems proven in this section are applied in Chapter 6.

4.1 Characterization of the quotient interleaving distance

In this section, we show that, under mild conditions, a quotient interleaving distance admits a description as an infimum over interleavings, similar to the one given for the interleaving distance in Definition 3.1.11. The characterization is given in terms of a weaker notion of interleaving.

Definition 4.1.1. *Let $\mathscr{C}$ be a relative locally persistent category. For $x, y \in \mathscr{C}$ and $\delta \in \mathbf{R}_+$, we say that x and y are **$\boldsymbol{\delta}$-quotient interleaved** if there exist $x' \simeq x$ and $y' \simeq y$ such that x' and y' are δ-interleaved.*

Definition 4.1.2. *In a locally persistent category $\mathscr{C}$, we say that a class E of 0-morphisms is **stable under weighted pullbacks** if the pullback of a morphism in E along any approximate morphism exists and is again a morphism in E.*

Before giving the characterization, we prove a useful lemma about the equivalence relation generated by a class of acyclic morphisms when these are stable under pullbacks.

Lemma 4.1.3. *Let $(C, \mathscr{W})$ be a relative category. If $\mathscr{W}$ is stable under pullbacks, then for any pair of objects $x, y \in C$ we have $x \simeq y$ if and only if there exists $c \in C$ and acyclic morphisms $c \to x$ and $c \to y$.*

Proof. It is clear that if there are acyclic morphisms $c \to x$ and $c \to y$, then $x \simeq y$, by definition of the equivalence relation $\simeq$.

For the converse, assume that x and y are connected by a zig-zag of length greater than two. This implies that the zig-zag starts as $x \to k \leftarrow k' \to \cdots$ or as $x \leftarrow k \to k' \leftarrow$

$k'' \to \cdots$. In the first case, we take the pullback of the cospan $x \to k \leftarrow k'$ and compose the composable morphisms we get, reducing the problem to the second case.

For the second case, we take the pullback of the cospan $k \to k' \leftarrow k''$ and compose the composable morphisms we get. Since $\mathscr{W}$ is closed under composition, we are left with a zig-zag between x and y of shorter length, so the proof follows by induction. □

Theorem 4.1.4. *Let $(\mathscr{C}, \mathscr{W})$ be a relative locally persistent category such that $\mathscr{W}$ is stable under weighted pullbacks. Then*

$$\begin{aligned}(d_I^{\mathscr{C}})_{/\mathscr{W}}(x,y) &= \inf\{\delta \in \mathbf{R}_+ : \ x \textit{ and } y \textit{ are } \delta\textit{-quotient interleaved}\} \\ &= \inf\{\delta \in \mathbf{R}_+ : \exists \textit{ morphisms } x' \to x \textit{ and } y' \to y \textit{ in } \mathscr{W} \\ &\qquad \textit{such that } x' \textit{ and } y' \textit{ are } \delta\textit{-interleaved}\}.\end{aligned}$$

The proof is inspired by [Mé17, Proposition 4.1] and [BL17, Section 4].

Proof. Given $x, y \in \mathscr{C}$, let

$$\begin{aligned}d_1(x,y) &= \inf\{\delta \in \mathbf{R}_+ : \ x \text{ and } y \text{ are } \delta\text{-quotient interleaved}\} \\ d_2(x,y) &= \inf\{\delta \in \mathbf{R}_+ : \exists \text{ morphisms } x' \to x \text{ and } y' \to y \text{ in } \mathscr{W} \\ &\qquad \text{such that } x' \text{ and } y' \text{ are } \delta\text{-interleaved}\}.\end{aligned}$$

We first show that $d_1 = (d_I^{\mathscr{C}})_{/\mathscr{W}}$. Let us start by showing that d_1 is an extended pseudo metric. Reflexivity and symmetry are immediate, so we show the triangle inequality. To prove this, it is enough to show that given $w, x, y, z \in \mathscr{C}$ with $x \simeq y$, an ε-interleaving $w \leftrightarrow x$, and a δ-interleaving $y \leftrightarrow z$, there exist $w' \simeq w$, $z' \simeq z$, and an $(\varepsilon+\delta)$-interleaving $w' \leftrightarrow z'$.

Since the class of acyclic morphisms is stable under pullbacks, there exist $c \in \mathscr{C}$ and acyclic morphisms $e : c \to x$ and $f : c \to y$, by Lemma 4.1.3. To conclude this part of the proof, we use the fact that acyclic morphisms are stable under weighted pullbacks and Proposition 3.2.11 to pull back the interleavings we were given along these maps, as follows

$$\begin{tikzcd}
 & w' \arrow[dl, dashed, "e'"'] \arrow[rr, dashed, leftrightarrow, "\varepsilon\quad\varepsilon"] & & c \arrow[dl, "e"] \arrow[dr, "f"'] \arrow[rr, dashed, leftrightarrow, "\delta\quad\delta"] & & z' \arrow[dr, dashed, "f'"] & \\
w \arrow[rr, leftrightarrow, "\varepsilon\quad\varepsilon"'] & & x & & y \arrow[rr, leftrightarrow, "\delta\quad\delta"'] & & z.
\end{tikzcd}$$

By hypothesis, e' and f' are acyclic morphisms. Composing the top-most interleavings, we get an $(\varepsilon+\delta)$-interleaving between w' and z'. Now, since d_1 is a $\simeq$-invariant metric that is bounded above by d_I, we have $d_1 \leq (d_I^{\mathscr{C}})_{/\mathscr{W}}$ by definition of $(d_I^{\mathscr{C}})_{/\mathscr{W}}$.

On the other hand, suppose that $d_1(x,y) < \delta < \infty$. It follows that there exist $x', y' \in \mathscr{C}$ such that $x' \simeq x$ and $y' \simeq y$ and such that x' and y' are δ-interleaved. By definition of $(d_I^{\mathscr{C}})_{/\mathscr{W}}$, we must have $(d_I^{\mathscr{C}})_{/\mathscr{W}}(x,y) < \delta$, so $(d_I^{\mathscr{C}})_{/\mathscr{W}} \leq d_1$, so $(d_I^{\mathscr{C}})_{/\mathscr{W}} = d_1$.

We now prove that $d_1 = d_2$. Clearly, we have $d_1 \leq d_2$. So it is enough to show that $d_2 \leq d_1$. Suppose that $d_1(x,y) < \delta < \infty$. As before, there are $x_1, y_1 \in \mathscr{C}$ such that $x_1 \simeq x$ and $y_1 \simeq y$ and such that x_1 and y_1 are δ-interleaved. By Lemma 4.1.3, there exist $x_2, y_2 \in \mathscr{C}$ and a diagram

$$\begin{tikzcd}
 & x_2 \arrow[dl] \arrow[dr] & & & y_2 \arrow[dl] \arrow[dr] & \\
x & & x_1 \arrow[r, leftrightarrow, "\delta"' near start, "\delta"' near end] & y_1 & & y,
\end{tikzcd}$$

in which the diagonal morphisms are in $\mathscr{W}$. By taking the following series of (weighted) pullbacks

$$\begin{tikzcd}
 & & x_4 \arrow[dl, dashed] \arrow[dr, dashed] \arrow[d, phantom, "\vee" very near start] \arrow[r, dashed, leftrightarrow, "\delta"' near start, "\delta"' near end] & y_3 \arrow[dr, dashed] \arrow[d, phantom, "\vee" very near start] & & \\
 & x_2 \arrow[dl] \arrow[dr] & & x_3 \arrow[dl, dashed] \arrow[d, phantom, "\vee" very near start] \arrow[r, dashed, leftrightarrow, "\delta"' near start, "\delta"' near end] & y_2 \arrow[dl] \arrow[dr] & \\
x & & x_1 \arrow[r, leftrightarrow, "\delta"' near start, "\delta"' near end] & y_1 & & y,
\end{tikzcd}$$

we obtain acyclic morphisms $\alpha : x_4 \to x$ and $\beta : y_3 \to y$, and a δ-interleaving between x_4 and y_3. It follows that $d_2(x,y) \leq \delta$, so $d_1 = d_2$, concluding the proof. $\square$

We finish this section by giving useful sufficient conditions for acyclic morphisms to be stable under weighted pullback.

Lemma 4.1.5. *Let $(\mathscr{C}, \mathscr{W})$ be a relative locally persistent category. Assume that $\mathscr{C}$ is powered by representables, that the underlying category of $\mathscr{C}$ admits pullbacks, and that powers by representables respect pullbacks in the underlying category of $\mathscr{C}$. If $\mathscr{W}$ is stable under power by representables and pullbacks of the underlying category of $\mathscr{C}$, then $\mathscr{W}$ is stable under weighted pullbacks.*

Proof. By Proposition 3.2.12, the locally persistent category $\mathscr{C}$ admits weighted pullbacks. Moreover, by the same result, the weighted pullback of a morphism is computed as a pullback in the underlying category of $\mathscr{C}$ of a power of this morphism by a representable. The result then follows from the fact that $\mathscr{W}$ is stable under power by representables and under pullbacks in the underlying category of $\mathscr{C}$. □

4.2 Stability

In this short section, we give simple conditions under which a locally persistent functor between relative locally persistent categories is distance non-increasing.

In Theorem 3.1.12 we showed that a locally persistent functor between locally persistent categories induces a distance non-increasing function with respect to the interleaving distances. To generalize this to quotient interleaving distances, we need the following general result about quotient metrics.

Lemma 4.2.1. *Let $f : (X,d) \to (X',d')$ be a distance non-increasing map between ep metric spaces, and let $R \subseteq X \times X$ and $R' \subseteq X' \times X'$ be equivalence relations. If f maps R-related elements to R'-related elements, then $f : (X, d_{/R}) \to (X', d'_{/R'})$ is distance non-increasing.*

Proof. Consider the ep metric d_f on X given by $d_f(x,y) = d'_{/R'}(f(x), f(y))$. It is enough to show that d_f is bounded above by $d_{/R}$.

Since f maps R-related elements to R'-related elements, and $d'_{/R'}$ is R'-invariant, d_f must be R-invariant. Since f is distance non-increasing with respect to d and d', we have that d_f is bounded above by d. So, by the universal property of $d_{/R}$, we have that d_f is bounded above by $d_{/R}$, as required. □

This directly implies the following stability result.

Theorem 4.2.2. *Let $F : \mathscr{C} \to \mathscr{C}'$ be a locally persistent functor between relative locally persistent categories $(\mathscr{C}, \mathscr{W})$ and $(\mathscr{C}', \mathscr{W}')$. If F maps $\mathscr{W}$-related objects to $\mathscr{W}'$-related objects, then F is distance non-increasing with respect to the quotient interleaving distances $(d_I^{\mathscr{C}})_{/\mathscr{W}}$ and $(d_I^{\mathscr{C}'})_{/\mathscr{W}'}$.* □

In particular, if $F : \mathscr{C} \to \mathscr{C}'$ is a locally persistent functor between relative locally persistent categories that maps acyclic morphisms to acyclic morphisms, then F is distance non-increasing with respect to the quotient interleaving distances.

Another useful stability result is the following.

Theorem 4.2.3. *Let $(\mathscr{C}, \mathscr{W})$ be a relative locally persistent category such that $\mathscr{W}$ is stable under weighted pullbacks and let P be an ep metric space. Let $f : \mathsf{obj}(\mathscr{C}) \to P$ be a function. If f maps $\mathscr{W}$-related objects to points at distance 0 and f is distance non-increasing (resp. uniformly continuous) with respect to the interleaving distance on $\mathscr{C}$ and the distance on P, then f is distance non-increasing (resp. uniformly continuous) with respect to the quotient interleaving distance and the distance on P.*

Proof. We start by proving the case of uniform continuity. Given $\varepsilon > 0$, let $\delta > 0$ be such that, if $x, y \in \mathscr{C}$ are δ-interleaved, then $d_P(f(x), f(y)) \leq \varepsilon$.

Now, assume that x and y are such that $\left(d_I^{\mathscr{C}}\right)_{/\sim}(x, y) < \delta$. By Theorem 4.1.4, there exist $x' \simeq x$ and $y' \simeq y$ such that x' and y' are δ-interleaved. It follows that $d_P(f(x), f(y)) = d_P(f(x'), f(y')) \leq \varepsilon$, using the fact that f maps $\mathscr{W}$-related objects to points at distance 0.

For the case of 1-Lipschitz maps, note that, in that case, we can take $\delta = \varepsilon$. □

4.3 Complete interleaving distances

In this section, we give sufficient conditions for a (quotient) interleaving distance to be complete. We use the notion of weighted sequential limit of Section 3.2.5.

Theorem 4.3.1. *Let $\mathscr{C}$ be a locally persistent category. If $\mathscr{C}$ admits weighted sequential limits, then $d_I^{\mathscr{C}}$ is complete.*

As is evident from the proof, is it actually enough for $\mathscr{C}$ to admit weighted sequential limits of morphisms that are part of an interleaving.

Proof. Let $\{x_i\}_{i \in \mathbb{N}} \subseteq \mathscr{C}$ be a Cauchy sequence with respect to the interleaving distance on $\mathscr{C}$. After taking a subsequence, we can assume that there exist $\varepsilon \in \mathbf{R}_+$ and $\varepsilon_i \in \mathbf{R}_+$ for each $i \in \mathbb{N}$ such that $\sum_{i \in \mathbb{N}} \varepsilon_i = \varepsilon$ and $d_I(x_i, x_{i+1}) < \varepsilon_i$.

By definition, we know that there are interleavings $f_i : x_{i+1} {}_{\varepsilon_i}\!\!\longleftrightarrow_{\varepsilon_i} x_i : h_i$, for each $i \in \mathbb{N}$. Let l be the weighted sequential limit of the morphisms f_i. By Proposition 3.2.14, there are interleavings $g_i : l\, {}_{\overline{\varepsilon_i}}\!\!\longleftrightarrow_{\overline{\varepsilon_i}} x_i : k_i$, where $\overline{\varepsilon}_i = \varepsilon - \sum_{j<i} \varepsilon_j$. Since $\overline{\varepsilon}_i \to 0$ as $i \to \infty$, it follows that l is the limit of the sequence $\{x_i\}$ according to the metric $d_I^{\mathscr{C}}$, as needed. □

This is enough to show that the interleaving distance on the category of locally persistent objects of a category that admits sequential limits is complete.

Corollary 4.3.2. *Let C be a category that admits sequential limits. Then the locally persistent category $C^{\mathbf{R}}$ admits weighted sequential limits and the interleaving distance on $C^{\mathbf{R}}$ is complete.*

Proof. By Example 3.2.7, the locally persistent category $C^{\mathbf{R}}$ is copowered and powered by representables. Since C admits sequential limits, so does $C^{\mathbf{R}}$, so $C^{\mathbf{R}}$ admits weighted sequential limits, by Proposition 3.2.15. Theorem 4.3.1 then implies that the interleaving distance is complete. □

To prove that a quotient interleaving distance is complete we need some assumptions about the interaction between approximate morphisms and acyclic morphisms.

Theorem 4.3.3. *Let $(\mathscr{C}, \mathscr{W})$ be a relative locally persistent category. If $\mathscr{C}$ admits weighted sequential limits, and $\mathscr{W}$ is closed under sequential limits in $\mathscr{C}_0$ and weighted pullbacks, then $(d_I^{\mathscr{C}})_{/\simeq}$ is complete.*

Proof. Analogously to the proof of Theorem 4.3.1, let $\{x_i\}_{i\in\mathbb{N}} \subseteq \mathscr{C}$ be a Cauchy sequence with respect to the quotient interleaving distance on $\mathscr{C}$. After taking a subsequence, we can assume that there exist $\varepsilon \in \mathbf{R}_+$ and $\varepsilon_i \in \mathbf{R}_+$ for each $i \in \mathbb{N}$ such that $\sum_{i\in\mathbb{N}} \varepsilon_i = \varepsilon$ and $(d_I^{\mathscr{C}})_{/\simeq}(x_i, x_{i+1}) < \varepsilon_i$.

By Theorem 4.1.4 and Lemma 4.1.3, we may assume that we have objects $\{c_i\}_{i\in\mathbb{N}}$, $\{y_i\}_{i\in\mathbb{N}}$, and $\{z_i\}_{i\in\mathbb{N}}$, interleavings $f_i : z_{i+1} \,{}_{\varepsilon_i}\!\!\longleftrightarrow_{\varepsilon_i} y_i : h_i$ and acyclic morphisms $w_i : c_i \to y_i$, $v_i : c_i \to z_i$, for each $i \in \mathbb{N}$, such that $c_i \simeq x_i$ for every $i \in \mathbb{N}$. Diagrammatically, we have the following

$$\begin{tikzcd}
 & & x_2 \arrow[d, dash, "\wr"] & & & x_1 \arrow[d, dash, "\wr"] & & & x_0 \arrow[d, dash, "\wr"] & \\
 & & c_2 \arrow[dl, "\sim"'] \arrow[dr, "\sim"] & & & c_1 \arrow[dl, "\sim"'] \arrow[dr, "\sim"] & & & c_0 \arrow[dl, "\sim"'] \arrow[dr, "\sim"] & \\
\cdots \arrow[r, leftrightarrow, "f_2", "h_2"'] & y_2 & & z_2 \arrow[r, leftrightarrow, "f_1", "h_1"'] & y_1 & & z_1 \arrow[r, leftrightarrow, "f_0", "h_0"'] & y_0 & & z_0 .
\end{tikzcd} \tag{4.3.4}$$

Let us introduce some notation to explain the rest of the proof. Let D denote the following indexing diagram

$$\begin{tikzcd}
 & & \bullet \arrow[dl] \arrow[dr] & & & \bullet \arrow[dl] \arrow[dr] & & \\
\cdots \arrow[r, leftrightarrow, "\varepsilon_2"' near start, "\varepsilon_2"' near end] & \bullet & & \bullet \arrow[r, leftrightarrow, "\varepsilon_1"' near start, "\varepsilon_1"' near end] & \bullet & & \bullet \arrow[r, leftrightarrow, "\varepsilon_0"' near start, "\varepsilon_0"' near end] & \bullet
\end{tikzcd}$$

Formally, this is a freely generated locally persistent category (Section 3.2.2). Then, Diagram 4.3.4 gives us a locally persistent functor $F_1 : D \to \mathscr{C}$, where we are forgetting about the objects x_i, and the objects c_0 and z_0, since they are weakly equivalent to y_0.

Intuitively, the rest of the proof works as follows. We would like to perform the transfinite composition of the interleavings we were given. We cannot quite do this, since the interleavings are separated by spans of acyclic morphisms. The idea is to pull back all the interleavings along the acyclic morphisms, to obtain a sequential diagram of interleavings, and then to take a limit, as we did in the proof of Theorem 4.3.1.

More precisely, we will proceed inductively, and construct diagrams $F_n : D \to \mathscr{C}$ for $n \in \mathbb{N}$, such that the first n spans of the n-th diagram are spans of identity morphisms. We will moreover construct natural transformations $\theta_n : F_{n+1} \Rightarrow F_n$ that are componentwise acyclic morphisms. Having done that, we will take the limit of the sequential diagram given by the natural transformations θ_n. This can be done, since (co)limits in a functor category are computed pointwise, and $\mathscr{C}$ is assumed to have sequential limits. Moreover, since the natural transformations have acyclic morphisms as components, and these are assumed to be closed under sequential limits, the limit diagram $\lim_n F_n$ comes with a natural transformation $\lim_n F_n \to F_1$ that is an acyclic morphism in each component. Finally, since the first n spans of F_n consist of identity morphisms, it follows that all of the spans of $\lim_n F_n$ consist of identity morphisms. Omitting the identities, we get a sequential diagram of interleavings, with the i-th object weakly equivalent to x_i. Taking the limit of this final diagram produces a limit for the initial sequence, concluding the proof.

The rest of the proof consists of constructing diagrams $F_n : D \to \mathscr{C}$ for $n \geq 1$ with the first $n-1$ spans consisting of identity morphisms, and natural transformations $\theta_n : F_{n+1} \Rightarrow F_n$, that are componentwise acyclic morphisms. Of course the first diagram has to be F_1. In order to do this, it is convenient to depict F_1 as follows:

$$\cdots \xleftarrow[\sim]{} \bullet \xrightarrow[\sim]{} \bullet \longleftrightarrow \bullet \xleftarrow[\sim]{} \bullet \xrightarrow[\sim]{} \bullet \longleftrightarrow \bullet \xleftarrow[\sim]{} \bullet \xrightarrow[\sim]{} \bullet \longleftrightarrow \bullet \xleftarrow[\sim]{} \bullet \xrightarrow[\sim]{} \bullet \longleftrightarrow \bullet$$

Here we are emphasizing the fact that the spans consist of acyclic morphisms, and deemphasizing the specific objects and the "lengths" of the interleavings. We are also giving the whole diagram in one line, so that the next step is clearer. Consider the following diagram, which we obtain from the one above by repeating some of the acyclic morphisms and taking the pullbacks depicted in the diagram:

$$\begin{tikzcd}
\cdots & \bullet \arrow[l, "\sim"'] \arrow[r, "\sim"] \arrow[d, equal] & \bullet \arrow[r, leftrightarrow] \arrow[d, equal] & \bullet \arrow[d, equal] & \bullet \arrow[l, "\sim"'] \arrow[r, "\sim"] \arrow[d, equal] & \bullet \arrow[r, leftrightarrow] \arrow[d, equal] & \bullet \arrow[d, equal] & \bullet \arrow[l, "\sim"'] \arrow[r, dashed, "\sim"] \arrow[d, dashed, "\wr"', "\beta"] & \bullet \arrow[r, dashed, leftrightarrow] \arrow[d, dashed, "\wr"', "\alpha"] & \bullet \arrow[r, equal] \arrow[d, "v_1"', "\wr"] & \bullet \arrow[r, equal] \arrow[d, equal] & \bullet \arrow[r, dashed, leftrightarrow] \arrow[d, "\wr"', "w_1"] & \bullet \arrow[d, dashed, "\wr"] \\
\cdots & \bullet \arrow[l, "\sim"] \arrow[r, "\sim"'] & \bullet \arrow[r, leftrightarrow] & \bullet & \bullet \arrow[l, "\sim"] \arrow[r, "\sim"'] & \bullet \arrow[r, leftrightarrow] & \bullet & \bullet \arrow[l, "\sim", "v_2"'] \arrow[r, "\sim", "w_2"'] & \bullet \arrow[r, leftrightarrow] & \bullet & \bullet \arrow[l, "\sim", "v_1"'] \arrow[r, "\sim", "w_1"'] & \bullet \arrow[r, leftrightarrow] & \bullet
\end{tikzcd}$$

More precisely, the above diagram is constructed by first taking the pullback of the right-most interleaving along w_1 (which can be done by Proposition 3.2.11), then taking the pullback of v_1 along the second to last interleaving, to obtain the acyclic morphism α, and finally taking the pullback of α along w_2 to obtain β. The morphism above v_2 is given by the composite $v_2 \circ \beta$.

We let $F_2 : D \to \mathscr{C}$ be the top-most row in the above diagram, and the natural transformation $\theta_1 : F_2 \Rightarrow F_1$ have as components the vertical morphisms in the above diagram. Note that all of the components are acyclic morphisms.

This process can be repeated now on the first two spans, to obtain a diagram $F_3 : D \to \mathscr{C}$ and a natural transformation with acyclic morphisms as components $\theta_2 : F_3 \Rightarrow F_2$, as follows:

$$\begin{tikzcd}
\cdots & \bullet \arrow[l, "\sim"'] \arrow[r, "\sim"] \arrow[d, equal] & \bullet \arrow[r, leftrightarrow] \arrow[d, equal] & \bullet \arrow[d, equal] & \bullet \arrow[l, "\sim"'] \arrow[r, dashed, "\sim"] \arrow[d, dashed, "\wr"] & \bullet \arrow[r, dashed, leftrightarrow] \arrow[d, dashed, "\wr"] & \bullet \arrow[r, equal] \arrow[d, "\wr"] & \bullet \arrow[r, equal] \arrow[d, equal] & \bullet \arrow[r, dashed, leftrightarrow] \arrow[d, "\wr"] & \bullet \arrow[r, equal] \arrow[d, dashed, "\wr"] & \bullet \arrow[r, equal] \arrow[d, dashed, "\wr"] & \bullet \arrow[r, dashed, leftrightarrow] \arrow[d, dashed, "\wr"] & \bullet \arrow[d, dashed, "\wr"] \\
\cdots & \bullet \arrow[l, "\sim"] \arrow[r, "\sim"'] & \bullet \arrow[r, leftrightarrow] & \bullet & \bullet \arrow[l, "\sim"] \arrow[r, "\sim"'] & \bullet \arrow[r, leftrightarrow] & \bullet & \bullet \arrow[l, "\sim"] \arrow[r, "\sim"'] & \bullet \arrow[r, leftrightarrow] & \bullet \arrow[r, equal] & \bullet \arrow[r, equal] & \bullet \arrow[r, leftrightarrow] & \bullet
\end{tikzcd}$$

In the above diagram, the bottom row represents F_2 and the top row represents F_3. The natural transformation $\theta_2 : F_3 \Rightarrow F_2$ has as components the vertical maps in the diagram. The proof follows by induction, taking $3n$ pullbacks at the n-th stage, as we've been doing. $\square$

4.4 Intrinsic interleaving distances

In this short section, we give sufficient conditions for a (quotient) interleaving distance to be intrinsic. We use the notion of terminal midpoint of Section 3.2.6.

Definition 4.4.1. *Let $\mathscr{C}$ be a locally persistent category and let $\varepsilon, \gamma, \delta \in \mathbf{R}_+$ be such that $\gamma + \delta = \varepsilon$. A **coherent factorization** of an ε-interleaving $f : x \;{}_{\varepsilon}\!\!\longleftrightarrow_{\varepsilon}\; y : g$ consists of the* *following data: for each* $r \in [0,\varepsilon]$*, an object* $z_r \in \mathscr{C}$*, and, for every* $r \le r' \in [0,\varepsilon]$*, an interleaving* $f_{r,r'} : z_r \,{}_{r'-r}\!\longleftrightarrow_{r'-r} z_{r'} : g_{r,r'}$ *such that, for every* $r \le r' \le r'' \in [0,\varepsilon]$*, we have* $f_{r',r''} \circ f_{r,r'} = f_{r,r''}$ *and* $g_{r,r'} \circ g_{r',r''} = g_{r,r''}$*, and such that when* $r = 0$ *and* $r' = \varepsilon$ *we get back the interleaving* $f : x \,{}_{\varepsilon}\!\longleftrightarrow_{\varepsilon} y : g$*.*

The following is then straightforward.

Theorem 4.4.2. *Let* $\mathscr{C}$ *be a locally persistent category that admits a coherent factorizations of every interleaving. Then, for every* $\delta \in \mathbf{R}_+$*, an interleaving* $x \,{}_{\delta}\!\longleftrightarrow_{\delta} y$ *in* $\mathscr{C}$ *induces a distance non-increasing map* $[0,\delta] \to (\mathscr{C}, d_I^{\mathscr{C}})$ *that sends* 0 *to* x *and* δ *to* y*. In particular,* $d_I^{\mathscr{C}}$ *is intrinsic.* □

Here the metric on $[0,\delta]$ is the metric inherited from $\mathbb{R}$. Note that Proposition 3.2.18 tells us that a locally persistent category that admits terminal midpoints necessarily admits coherent factorizations of interleavings. As an immediate corollary of this fact and Theorem 4.4.2, we have the following.

Corollary 4.4.3. *Let* $\mathscr{C}$ *be a locally persistent category that admits terminal midpoints. Then, for every* $\delta \in \mathbf{R}_+$*, an interleaving* $x \,{}_{\delta}\!\longleftrightarrow_{\delta} y$ *in* $\mathscr{C}$ *induces a distance non-increasing map* $[0,\delta] \to (\mathscr{C}, d_I^{\mathscr{C}})$ *that sends* 0 *to* x *and* δ *to* y*. In particular,* $d_I^{\mathscr{C}}$ *is intrinsic.* □

This is enough to show that the interleaving distance on the category of locally persistent objects of a category that admits pullbacks and binary products is intrinsic.

Corollary 4.4.4. *Let* C *be a category that admits pullbacks and binary products. Then the interleaving distance on the locally persistent category* $C^{\mathbf{R}}$ *is intrinsic.*

Proof. By Example 3.2.7, the locally persistent category $C^{\mathbf{R}}$ is copowered and powered by representables. Since C admits pullbacks and binary products, so does $C^{\mathbf{R}}$, so $C^{\mathbf{R}}$ admits terminal midpoints, by Proposition 3.2.19. Corollary 4.4.3 then implies that the interleaving distance is intrinsic. □

The analogous result for quotient interleaving distances is easy to prove in this case. Since a quotient of any intrinsic ep metric is intrinsic (Proposition 2.2.22), the following is immediate.

Corollary 4.4.5. *Let* $\mathscr{C}$ *be a locally persistent category with an equivalence relation on its class of objects. If* $\mathscr{C}$ *admits coherent factorizations of interleavings, then* $d_{QI}^{\mathscr{C}}$ *is intrinsic. In particular, if* $\mathscr{C}$ *admits terminal midpoints, then* $d_{QI}^{\mathscr{C}}$ *is intrinsic.* □

4.5 Geodesic interleaving distances

In this section we give sufficient conditions under which a (quotient) interleaving distance is geodesic. We start by addressing the case of non-quotient interleaving distances, and generalize the result to quotient interleaving distances afterwards.

4.5.1 Geodesic non-quotient interleaving distances

The main question that arises when trying to prove that an interleaving distance is geodesic is the following. Assume that $\mathscr{C}$ is a locally persistent category, and let $x, y \in \mathscr{C}$ satisfy $d_I^{\mathscr{C}}(x, y) = \delta$; under what conditions is it true that x and y are δ-interleaved?

This motivates the following definition.

Definition 4.5.1. *Let $\mathscr{C}$ be a locally persistent category. The distance $d_I^{\mathscr{C}}$ **reflects interleavings** if the following holds for all $x, y \in \mathscr{C}$ and $\delta \in \mathbf{R}_+$: if $d_I^{\mathscr{C}}(x, y) = \delta$, then x and y are δ-interleaved.*

Note that, in the hypothesis of the above definition, if $d_I(x, y) = 0$ and the interleaving distance reflects interleavings, then $x \cong y$. An interleaving distance satisfying the property above is sometimes referred to as a *closed* interleaving distance ([Les15],[BG18]). The connection with being geodesic is established by the following result.

Theorem 4.5.2. *Let $\mathscr{C}$ be a locally persistent category that admits coherent factorizations of interleavings. If $d_I^{\mathscr{C}}$ reflects interleavings, then $d_I^{\mathscr{C}}$ is geodesic.*

Proof. Let $x, y \in \mathscr{C}$ and $\delta \in \mathbf{R}_+$ such that $d_I^{\mathscr{C}}(x, y) = \delta$. Since $d_I^{\mathscr{C}}$ reflects interleavings, there is a δ-interleaving between x and y. By Theorem 4.4.2, we can use a coherent factorizations of this interleavings to construct a path of length δ between x and y, concluding the proof. □

Remark 4.5.3. In [Les15, Theorem 6.1], it is shown that the interleaving distance between finitely presented multi-persistent modules reflects interleavings, and the proof strategy in fact generalizes to the interleaving distance between objects of a functor category $C^{\mathbf{R}^n}$ that are left Kan extensions of finite posets $Q \hookrightarrow \mathbf{R}^n$. In Proposition 4.5.12, we show that the main result of this section, Proposition 4.5.11, generalizes Lesnick's result.

The rest of this subsection is devoted to giving sufficient conditions to deduce that an interleaving distance reflects interleavings. Given a poset P, a category C, a functor $X \in C^P$, and $r \leq s \in P$, recall that we let $\varphi^X_{r,s} : X(r) \to X(s)$ denote the structure morphism of X. The following definition generalizes the notion of q-tame persistent module ([CSGO16]) to other functor categories $C^{\mathbf{R}}$ where the category C has a notion of "small" or "compact" object.

Definition 4.5.4. *Let P be a poset. Let X be an object of $\mathbf{Set}^P$ (resp. $\mathbf{Top}^P$). We say that X is **q-tame** if for every $r, s \in P$ such that $r < s$, the image of $\varphi^X_{r,s} : X(r) \to X(s)$ is finite (resp. compact).*

Example 4.5.5. Let $X \in \mathbf{Set}^P$. Endowing X with the discrete topology objectwise, we get $X' \in \mathbf{Top}^P$ with X as underlying set-valued functor. Then X is q-tame if and only if X' is q-tame.

In order to prove that an interleaving distance reflects interleavings, we need to be able to construct a δ-interleaving out of a sequence of δ_n-interleavings for $\delta_n \to \delta$ from above. Intuitively, we do this in two steps. The first step is to make the sequence of δ_n-interleavings *coherent*, so that in the second step we can take a categorical limit of this coherent sequence and get a δ-interleaving. The notion of coherence is the following.

Definition 4.5.6. *Let P be a poset and let X be an object of $\mathbf{Set}^P$ (resp. $\mathbf{Top}^P$). A **compatible family** for X consists of an element $x_r \in X(r)$ for every $r \in P$, such that $\varphi^X_{r,s}(x_r) = x_s$ whenever $s \leq r$.*

Note that the set of compatible families of a functor X is canonically isomorphic to the (underlying set of the) limit of X. The following result by Stone gives conditions under which a functor $(\mathbb{N}, \geq) \to \mathbf{Top}$ admits a compatible family.

Proposition 4.5.7 ([Sto79, Theorem 2])**.** *Let $X \in \mathbf{Top}^{(\mathbb{N},\geq)}$ be objectwise compact with closed structure morphisms. If X is objectwise non-empty, then there exists a compatible family for X.* □

We interpret the above theorem as constructing a compatible family, the one in the conclusion, out of a non-compatible one, the one that makes the functor objectwise non-empty. The notion of q-tameness allows us to relax the hypothesis of the above theorem. Before proceeding, we introduce the notion of subfunctor.

Definition 4.5.8. *Let P be a poset and let $X : P \to \mathbf{Set}$ (resp. $X : P \to \mathbf{Top}$) be a functor. A* ***subfunctor*** *of X is a given by a functor $Y : P \to \mathbf{Set}$ (resp. $Y : P \to \mathbf{Top}$) such that $Y(r) \subseteq X(r)$ for every $r \in P$ and such that the structure morphisms of Y are the restrictions of the structure morphisms of X. In the case that $X : P \to \mathbf{Top}$, we require Y to have the subspace topology, and say that Y is a* ***closed subfunctor*** *if $Y(r) \subseteq X(r)$ is a closed subspace for every $r \in P$.*

Proposition 4.5.9. *Let $X \in \mathbf{Top}^{(\mathbb{N},\geq)}$ be q-tame with closed structure morphisms. If X is objectwise non-empty, then there exists a compatible family for X.*

Proof. Consider the subfunctor $Y \subseteq X$ given by the image of the structure morphisms, as follows $Y(n) = \varphi^X_{n+1,n}(X(n+1))$. It is enough to construct a compatible family for Y.

Since X is q-tame, Y is objectwise compact. Moreover, by construction, $Y(n)$ is closed in $X(n)$. This implies that the structure morphisms of Y are closed, since a closed set in $Y(n+1)$ is closed in $X(n+1)$, so its image in $X(n)$ is closed, and thus also closed in $Y(n)$.

Finally, Y is objectwise non-empty, since X is. We can then apply Proposition 4.5.7 to obtain a compatible family for Y, concluding the proof. □

The next definition gives a persistent object structure to the collection of all interleavings between a pair of objects.

Definition 4.5.10. *Let $\mathscr{C}$ be a locally persistent category and let $x, y \in \mathscr{C}$. The* ***persistent set of interleavings*** *between x and y is the functor $\mathbb{I}(x, y) : \mathbf{R}_+ \to \mathbf{Set}$ given by*

$$\mathbb{I}(x,y)_\delta = \left\{(f,g) : \ f \text{ and } g \text{ form a } \delta\text{-interleaving between } x \text{ and } y\right\},$$

for every $\delta \in \mathbf{R}_+$, with the structure maps given by the shift $\mathbf{S}$.

The following is a key result when establishing that an interleaving distance reflects interleavings.

Proposition 4.5.11. *Let $\mathscr{C}$ be a locally persistent category and let $x, y \in \mathscr{C}$. Assume that the persistent set of interleavings $\mathbb{I}(x, y)$ is right continuous, and that for each $\delta \in \mathbf{R}_+$ the set $\mathbb{I}(x, y)_\delta$ can be given a topology such that the structure morphisms of $\mathbb{I}(x, y)$ are continuous and closed and such that $\mathbb{I}(x, y) : \mathbf{R}_+ \to \mathbf{Top}$ is q-tame. Let $\delta \in \mathbf{R}_+$. If $d_I^{\mathscr{C}}(x, y) = \delta$, then x and y are δ-interleaved.*

Proof. If $d_I^{\mathscr{C}}(x,y) = \delta$, then there exists a strictly decreasing sequence $S = \{\delta_n\}_{n\in\mathbb{N}} \subseteq \mathbf{R}_+$ converging to δ such that there is a δ_n-interleaving $(f_n, g_n) \in \mathbb{I}(\delta_n)$ for each $n \in \mathbb{N}$. Restrict $\mathbb{I}$ to S, as a subposet of $\mathbf{R}_+$, and denote it by $\mathbb{I}|_S : S \to \mathbf{Top}$. Note that, as a poset, S is isomorphic to $(\mathbb{N}, \geq)$.

We claim that it is enough to construct a compatible family for $\mathbb{I}|_S$. This is because this compatible family gives a compatible family for $\mathbb{I}|_{(\delta,\infty)}$, since S is cofinal in (δ,∞) because $\delta_n \to \delta$. So such a compatible family gives us $(f,g) \in \mathbb{I}(\delta)$, by the right continuity condition of $\mathbb{I}$. By definition of $\mathbb{I}$, the morphisms f and g form a δ-interleaving between x and y, as required.

We finish the proof by constructing a compatible family for $\mathbb{I}|_S$. In order to do this, we use Proposition 4.5.9. To satisfy the hypothesis, notice that $\mathbb{I}|_S$ is q-tame with closed structure morphisms. The elements $(f_n, g_n) \in \mathbb{I}(r_n)$ witness the fact that $\mathbb{I}|_S$ is objectwise non-empty, so Proposition 4.5.9 gives a compatible family for $\mathbb{I}|_S$, as required. □

A simple application is the following.

Proposition 4.5.12. *Let $n \in \mathbb{N}$, $C \in \mathbf{Cat}$, and let $\mathscr{C} \subseteq C^{\mathbf{R}^n}$ be the locally persistent subcategory spanned by objects $X : \mathbf{R}^n \to C$ that are isomorphic to a left Kan extension of a functor $P_X \to C$ for $P_X \subseteq \mathbf{R}^n$ a finite subposet. Then the interleaving distance on $\mathscr{C}$ reflects interleavings.*

Proof. Let $X, Y \in \mathscr{C}$. The functors X and Y can be written as left Kan extensions of functors $P \to C$ for a common finite poset $P \subseteq \mathbf{R}^n$, for example, by letting $P = P_X \cup P_Y$. Since X and Y are left Kan extensions of functors $P \to C$, it follows that $\mathbb{I}(X,Y) : \mathbf{R}_+ \to \mathbf{Set}$ is right continuous, as the value of this functor changes finitely many times and is continuous from the right at these values. By Proposition 4.5.11, it is then enough to show that $\mathbb{I}(X,Y)$ can be lifted to a q-tame persistent topological space.

Let $\{v_1, \ldots, v_k\} \subseteq \mathbf{R}^n$ be the set of values at which $\mathbb{I}(X,Y)$ changes, $\mathbb{I}(X,Y)(r)$ being empty for $r < v_1$. It is then enough to give a topology on $\mathbb{I}(X,Y)(v_i)$ such that the map $\mathbb{I}(X,Y)(v_i) \to \mathbb{I}(X,Y)(v_{i+1})$ is closed and has compact image for each $i < k$. We define the topology inductively. The topology on $\mathbb{I}(X,Y)(v_1)$ is the codiscrete topology, and the topology on $\mathbb{I}(X,Y)(v_{i+1}) = \mathbf{Im}(\varphi_{v_i,v_{i+1}}) \coprod \mathbf{Im}(\varphi_{v_i,v_{i+1}})^c$ is taken to be the coproduct topology of the quotient topology on $\mathbf{Im}(\varphi_{v_i,v_{i+1}})$, induced by $\mathbb{I}(X,Y)(v_{i+1}) \to \mathbf{Im}(\varphi_{v_i,v_{i+1}})$, and the codiscrete topology on $\mathbf{Im}(\varphi_{v_i,v_{i+1}})^c$.

It is clear that the structure maps are closed, and that these images are also compact, since, in fact, all the spaces are compact, as they are binary coproducts of compact spaces. □

We now give a result that allows us to check the hypothesis of Proposition 4.5.11 more easily. In order to prove this result, we need the following technical lemma that identifies a well-behaved class of subfunctors of q-tame functors with closed structure morphisms.

Lemma 4.5.13. *Let $X \in \mathbf{Top}^{(\mathbb{N},\geq)}$ be q-tame with closed structure morphisms. Let $Y \subseteq X$ be a closed subfunctor. Then Y is q-tame and its structure morphisms are closed.*

Proof. Let us start by showing that Y has closed structure morphisms. Since $Y(n+1)$ is closed in $X(n+1)$, its image $\varphi_{n+1,n}(Y(n+1))$ in $X(n)$ must also be closed. And since $Y(n)$ is closed in $X(n)$, the image $\varphi_{n+1,n}(Y(n+1))$ must be closed in $Y(n)$.

To see that Y is q-tame, note that $\varphi_{n+1,n}(Y(n+1))$ is closed in $\varphi_{n+1,n}(X(n+1))$, since it is closed in $X(n)$. Since X is q-tame, $\varphi_{n+1,n}(X(n+1))$ is compact, and thus $\varphi_{n+1,n}(Y(n+1))$ is compact too. □

Theorem 4.5.14. *Let $\mathscr{C}$ be a locally persistent category. Suppose that for every $x, y \in \mathscr{C}$ the functor $\mathrm{Hom}_{\mathscr{C}}(x,y) : \mathbf{R}_+ \to \mathbf{Set}$ is right continuous, and that for every $\delta \in \mathbf{R}_+$ the set $\mathrm{Hom}_{\mathscr{C}}(x,y)_\delta$ admits a T_1 topology such that:*

1. *the structure maps of $\mathrm{Hom}_{\mathscr{C}}(x,y)$ are continuous and closed and the functor $\mathrm{Hom}(x,y) : \mathbf{R}_+ \to \mathbf{Top}$ is q-tame;*

2. *for each $x, y, z \in \mathscr{C}$ and each $\varepsilon, \delta \in \mathbf{R}_+$, the composition operation $\mathrm{Hom}_{\mathscr{C}}(y,z)_\delta \times \mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon \to \mathrm{Hom}_{\mathscr{C}}(x,z)_{\varepsilon+\delta}$ is continuous.*

Then, $d_I^{\mathscr{C}}$ reflects interleavings.

Proof. Consider $H \in \mathbf{Top}^{\mathbf{R}_+}$ given by $H(\delta) = \mathrm{Hom}(x,y)_\delta \times \mathrm{Hom}(y,x)_\delta$, and $K \in \mathbf{Top}^{\mathbf{R}_+}$ given by $K(\delta) = \mathrm{Hom}(x,x)_{2\delta} \times \mathrm{Hom}(y,y)_{2\delta}$. Note that composition gives us a continuous natural transformation $H \to K$. Then, the persistent set of interleavings can be seen as the subfunctor $\mathbb{I}(x,y) \subseteq H$ that is the preimage under the composition map $H \to K$ of the elements $(\mathbf{S}_{0,2\delta}(\mathsf{id}_x), \mathbf{S}_{0,2\delta}(\mathsf{id}_y)) \in K(\delta)$. Since the spaces considered are all T_1, the subfunctor $\mathbb{I}(x,y) \subseteq H$ is closed in H. This implies that $\mathbb{I}(x,y)$ is q-tame with closed structure morphisms, by Lemma 4.5.13. Note also that $\mathrm{Hom}(x,y)$ is right continuous for all $x, y \in \mathscr{C}$, so the functors H and K are right continuous. Since categorical limits commute with categorical limits, the functor $\mathbb{I}(x,y)$ is right continuous too. We can then finish the proof by applying Proposition 4.5.11. □

4.5.2 Geodesic quotient interleaving distances

The case of quotient interleaving distances is only slightly more complicated.

Definition 4.5.15. *Let $\mathscr{C}$ be a relative locally persistent category. The distance $d_{QI}^{\mathscr{C}}$* ***reflects quotient interleavings*** *if the following holds for all $x, y \in \mathscr{C}$ and $\delta \in \mathbf{R}_+$: if $d_{QI}^{\mathscr{C}}(x, y) = \delta$, then x and y are δ-quotient interleaved.*

Note that, in the hypothesis of the above definition, if $(d_I)_{/\simeq}(x, y) = 0$ and the quotient interleaving distance reflects interleavings, then $x \simeq y$. The following is proven in the same way as Theorem 4.5.2, but using the second characterization of $d_{QI}^{\mathscr{C}}$ in Theorem 4.1.4.

Theorem 4.5.16. *Let $(\mathscr{C}, \mathscr{W})$ be a relative locally persistent category that admits coherent factorizations of interleavings and such that $\mathscr{W}$ is stable under weighted pullbacks. If $d_{QI}^{\mathscr{C}}$ reflects quotient interleavings, then $d_{QI}^{\mathscr{C}}$ is geodesic.* □

We conclude this subsection by giving conditions under which a quotient interleaving distance reflects quotient interleavings.

Definition 4.5.17. *Let $(\mathscr{C}, \mathscr{W})$ be a relative locally persistent category such that $\mathscr{W}$ is stable under weighted pullbacks. Let $x, y \in \mathscr{C}$. The* ***persistent set of quotient interleavings*** *between x and y is the functor $\mathbb{QI}(x, y) : \mathbf{R}_+ \to \mathbf{SET}$ given by*

$$\mathbb{QI}(x, y)_\delta = \Big\{(x', y', u, v, f, g) \,:\, x', y' \in \mathscr{C},\, u : x' \to x \text{ and } v : y' \to y \text{ belong to } \mathscr{W}, \\ \text{and } f \text{ and } g \text{ form a } \delta\text{-interleaving between } x \text{ and } y\Big\},$$

for every $\delta \in \mathbf{R}_+$, with the structure morphisms given by shifting the morphisms f and g.

The following theorem is proven in exactly the same way as Proposition 4.5.11.

Theorem 4.5.18. *Let $\mathscr{C}$ be a relative locally persistent category and let $x, y \in \mathscr{C}$. Assume that the persistent set of quotient interleavings $\mathbb{QI}(x, y)$ is right continuous and that for each $\delta \in \mathbf{R}_+$ the set $\mathbb{QI}(x, y)_\delta$ can be given a topology such that the structure morphisms of $\mathbb{QI}(x, y)$ are continuous and closed and such that $\mathbb{QI}(x, y) : \mathbf{R}_+ \to \mathbf{TOP}$ is q-tame. Let $\delta \in \mathbf{R}_+$. If $d_{QI}^{\mathscr{C}}(x, y) = \delta$, then there exist acyclic morphisms $u : x' \to x$ and $v : y' \to y$ such that x' and y' are δ-interleaved.* □

4.6 Higher interpolation in locally persistent categories

In this section, we extend the interpolation framework developed in [BSN17] from categories of persistent objects to general locally persistent categories. Our main motivation is to show that their framework is most natural when seen from the perspective of locally persistent categories, as it concerns a fundamental relationship between ep metric spaces and locally persistent categories. One important consequence of this relationship is Proposition 4.6.4, namely, that ep metric spaces form a full subcategory of the category of locally persistent categories, and thus that every distance can be realized as an interleaving distance, albeit in a rather trivial way.

The key question studied in [BSN17] can be phrased as follows. Let $\mathscr{C}$ be a locally persistent category and let $P \subseteq Q$ be an inclusion of metric spaces. Given a distance non-increasing function $P \to \left(\mathsf{obj}(\mathscr{C}), d_I^{\mathscr{C}}\right)$, under what conditions can this function be extended to a distance non-increasing function $Q \to \left(\mathsf{obj}(\mathscr{C}), d_I^{\mathscr{C}}\right)$? In [BSN17], sufficient conditions for the existence of this extension are given in the case when $\mathscr{C}$ is a category of persistent objects of the form $C^{\mathbf{R}}$, for C a category. The sufficient conditions require the distance non-increasing map $P \to \left(\mathsf{obj}(\mathscr{C}), d_I^{\mathscr{C}}\right)$ to be *coherent* in a certain sense, and C to be complete or cocomplete. In Section 4.6.1, we extend this coherence condition and their main result to locally persistent categories. In Section 4.6.2, we explain in what way our interpolation framework is a generalization of the one presented in [BSN17].

4.6.1 Extensions of maps from metric spaces

Given a locally persistent category $\mathscr{C}$, we denote the (possibly large) ep metric space given by the objects of $\mathscr{C}$ together with the interleaving distance as $\mathsf{met}(C) = \left(\mathsf{obj}(\mathscr{C}), d_I^{\mathscr{C}}\right)$. Theorem 3.1.12 tells us that we have a functor

$$\begin{aligned} \mathsf{met} : \mathbf{lpCAT} &\to \mathbf{epMET} \\ \mathscr{C} &\mapsto \left(\mathsf{obj}(\mathscr{C}), d_I^{\mathscr{C}}\right). \end{aligned}$$

There is a natural functor going the other way, which we now describe. Given an ep metric space P, we can construct a locally persistent category $\mathsf{cat}(P)$ by letting

$\mathsf{obj}(\mathsf{cat}(P))$ be the underlying set of P, and

$$\mathrm{Hom}_{\mathsf{cat}(P)}(p,q)_r = \begin{cases} \{*\}, & \text{if } d_P(p,q) \leq r \\ \emptyset, & \text{if } d_P(p,q) > r, \end{cases}$$

for every $p, q \in P$. Composition is defined in the only possible way, using the triangle inequality of P. This construction gives a functor $\mathsf{cat} : \mathbf{epMET} \to \mathbf{lpCAT}$.

Some interesting locally persistent categories can be constructed in this way; we give an example.

Example 4.6.1. Given $\delta \in \mathbf{R}_+$, let $2_\delta \in \mathbf{epMet}$ be the metric space with underlying set $\{a, b\}$ and such that $d(a,b) = \delta$. Observe that δ-interleavings in a locally persistent category $\mathscr{C}$ are represented by $\mathsf{cat}(2_\delta)$ in the sense that there is a bijection between δ-interleavings in $\mathscr{C}$ and locally persistent functors $\mathsf{cat}(2_\delta) \to \mathscr{C}$.

The functors $\mathsf{met} : \mathbf{lpCAT} \rightleftarrows \mathbf{epMET} : \mathsf{cat}$ do not form an adjunction. In order to get an adjunction, one can consider the full subcategory $\mathbf{lpCAT}_{\mathsf{interl}} \subseteq \mathbf{lpCAT}$ spanned by locally persistent categories where, for every $\varepsilon \in \mathbf{R}_+$, every ε-morphism is part of an ε-interleaving. The inclusion $\iota : \mathbf{lpCAT}_{\mathsf{interl}} \to \mathbf{lpCAT}$ has a right adjoint $\mathsf{core} : \mathbf{lpCAT} \to \mathbf{lpCAT}_{\mathsf{interl}}$ that maps a locally persistent category $\mathscr{C}$ to the locally persistent category $\mathsf{core}(\mathscr{C})$ with the same collection of objects and such that

$$\mathrm{Hom}_{\mathsf{core}(\mathscr{C})}(x,y)_\varepsilon = \left\{ f \in \mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon \;:\; f \text{ is part of an } \varepsilon\text{-interleaving} \right\}.$$

It is clear that met factors through core, and that cat factors through ι, that is, that the following diagram commutes (strictly):

$$\begin{tikzcd}
\mathbf{lpCAT} \arrow[r, bend left, "\mathsf{core}"] \arrow[rr, bend left=60, "\mathsf{met}"] & \mathbf{lpCAT}_{\mathsf{interl}} \arrow[l, bend left, "\iota"] \arrow[r, bend left, "\mathsf{met}_{\mathsf{interl}}"] & \mathbf{epMET}. \arrow[l, bend left, "\mathsf{cat}_{\mathsf{interl}}"] \arrow[ll, bend left=60, "\mathsf{cat}"]
\end{tikzcd}$$

We have the following adjunction.

Proposition 4.6.2. *The functors* $\mathsf{met}_{\mathsf{interl}} : \mathbf{lpCAT}_{\mathsf{interl}} \rightleftarrows \mathbf{epMET} : \mathsf{cat}_{\mathsf{interl}}$ *form an adjunction, with* $\mathsf{met}_{\mathsf{interl}} \dashv \mathsf{cat}_{\mathsf{interl}}$.

Proof. Note, on the one hand, that a distance non-increasing function between metric spaces is completely determined by its value on the elements of its domain. On the other hand, if $\mathscr{C}$ is a locally persistent category, P an ep metric space, $x, y \in P$, and $\varepsilon \in \mathbf{R}_+$, then there is at most one morphisms $x \to_\varepsilon y$ in $\mathsf{cat}(P)$. This implies that a distance non-increasing functor $\mathscr{C} \to \mathsf{cat}_{\mathsf{interl}}(P)$ is entirely determined by its action on the objects of $\mathscr{C}$. These two observations give a natural bijection between distance non-increasing maps $\mathsf{met}_{\mathsf{interl}}(\mathscr{C}) \to P$ and locally persistent functors $\mathscr{C} \to \mathsf{cat}_{\mathsf{interl}}(P)$, whenever all the morphisms of $\mathscr{C}$ are part of an interleaving. □

Since the adjunctions $\iota \dashv \mathsf{core}$ and $\mathsf{met}_{\mathsf{interl}} \dashv \mathsf{cat}_{\mathsf{interl}}$ go in different directions, they don't compose to an adjunction between **lpCAT** and **epMET**. Nonetheless, there is a counit $c : \mathsf{met} \circ \mathsf{cat} = \mathsf{met}_{\mathsf{interl}} \circ \mathsf{cat}_{\mathsf{interl}} \Rightarrow \mathrm{id}_{\mathbf{epMET}}$, which, for P a metric space, is defined as the distance non-increasing map $\mathsf{met}(\mathsf{cat}(P)) \to P$ that sends each element of P to itself. Note, moreover, that if $d_P(p,q) = r$, then p and q are r-interleaved as objects of $\mathsf{cat}(P)$. This implies the following.

Lemma 4.6.3. *The morphism* $c_P : \mathsf{met}(\mathsf{cat}(P)) \to P$ *is a natural isomorphism of ep metric spaces for every* $P \in \mathbf{epMET}$. □

Although we won't need this, it is important to emphasize the following.

Proposition 4.6.4. *The functor* $\mathsf{cat} : \mathbf{epMET} \to \mathbf{lpCAT}$ *exhibits the category of (large) ep metric spaces as a full subcategory of the category of locally persistent category.*

Proof. Lemma 4.6.3 implies that cat is faithful. To see that cat is full, we use that a functor $\mathsf{cat}(P) \to \mathsf{cat}(Q)$ is completely determined by its action on the objects of $\mathsf{cat}(P)$, as the hom-persistent sets of $\mathsf{cat}(Q)$ are valued in either empty or singleton sets. □

The extension of a distance non-increasing map from a metric space to a locally persistent category works by extending a corresponding locally persistent functor, using a Kan extension. In order to do this, given a distance non-increasing map $P \to \mathsf{met}(\mathscr{C})$ we require a locally persistent functor $\mathsf{cat}(P) \to \mathscr{C}$ representing it. Since c is an isomorphism, we can assign, to each locally persistent functor $F : \mathsf{cat}(P) \to \mathscr{C}$, a distance non-increasing map $\mathsf{met}(F) \circ c_P^{-1} : P \to \mathsf{met}(\mathscr{C})$.

We use this to define the notion of coherent map from a metric space to a locally persistent category.

Definition 4.6.5. *Let P be an ep metric space and let $\mathscr{C}$ be a locally persistent category. We say that a distance non-increasing map $f : P \to \mathsf{met}(\mathscr{C})$ is **coherent** if there exists a locally persistent functor $F : \mathsf{cat}(P) \to \mathscr{C}$ such that $\mathsf{met}(F) \circ c_P^{-1} = f$.*

Before going to the main theorem, let us give some examples that show that interesting problems in the theory of interleaving distances can be phrased as whether certain maps are coherent or as extension problems.

Example 4.6.6. Let $\mathscr{C}$ be a locally persistent category. Then $d_I^{\mathscr{C}}$ reflects interleavings (Definition 4.5.1) if and only if for every $\delta \in \mathbf{R}_+$, every map $2_\delta \to \mathsf{met}(\mathscr{C})$ is coherent.

Example 4.6.7. Let P be an ep metric space. Then d_P is intrinsic if and only if every distance non-increasing map $2_\delta \to P$ can be extended to a distance non-increasing map $[0,\delta] \to P$ (endowing $[0,\delta]$ with the metric induced by $\mathbb{R}$) where the inclusion $2_\delta \to [0,\delta]$ maps a to 0 and b to δ.

Example 4.6.8 (cf. second proof of [BV18, Theorem 4.25]). Let P be an ep metric space. Then d_P is complete if and only if every distance non-increasing map $\{1/2^n\}_{n\geq 0} \to P$ can be extended to a distance non-increasing map $\{1/2^n\}_{n\geq 0} \cup \{0\} \to P$, endowing $\{1/2^n\}_{n\geq 0}$ and $\{1/2^n\}_{n\geq 0} \cup \{0\}$ with the distances induced by $\mathbb{R}$.

We now prove the main theorem. For this we need the following straightforward lemma.

Lemma 4.6.9. *Let $P \to Q$ be a distance preserving map between ep metric spaces. Then, the induced locally persistent functor $\mathsf{cat}(P) \to \mathsf{cat}(Q)$ is fully faithful.* □

Theorem 4.6.10 (cf. [BSN17, Theorem 3.6])**.** *Let $P, Q \in$ **epMet** and $\mathscr{C} \in$ **lpCAT**, let $f : P \to \mathsf{met}(\mathscr{C})$ be a coherent, distance non-increasing map, and let $g : P \to Q$ be a distance preserving map. If $\mathscr{C}$ is complete or cocomplete as an enriched category, then f can be extended along g, as follows*

$$\begin{tikzcd} P \arrow[r, "f"] \arrow[d, "g"'] & \mathsf{met}(\mathscr{C}) \\ Q \arrow[ur, dashed] & \end{tikzcd}$$

Proof. Assume $\mathscr{C}$ is complete (the other case is dual). Since f is coherent, we can form the following diagram of locally persistent categories

$$\begin{tikzcd} \mathsf{cat}(P) \arrow[r, "F"] \arrow[d, "\mathsf{cat}(g)"'] & \mathscr{C} \\ \mathsf{cat}(Q) \arrow[ur, dashed] & \end{tikzcd}$$

where the dotted arrow is the right Kan extension of F along g (Definition 2.4.15), which exists by completeness of $\mathscr{C}$. The triangle commutes strictly, since $\mathsf{cat}(g)$ is full and faithful, by Lemma 4.6.9 and Proposition 2.4.16. After applying met to the diagram, we obtain the desired extension, since, by assumption, we have $\mathsf{met}(F) \circ c_P^{-1} = f$. $\square$

Example 4.6.11. Using Theorem 4.6.10, and the observations in Example 4.6.8 and Example 4.6.7, it follows that the interleaving distance of a complete locally persistent category is intrinsic and complete. This recovers a weak version of Corollary 4.4.3 and Theorem 4.3.1, as here we are assuming that the locally persistent category admits all limits, whereas the aforementioned results require the existence of a specific kind of limit.

4.6.2 Relationship to higher interpolation and extension for persistent modules

The functors met and cat that we described above play the roles of the functors $\bullet^{\mathbf{R}} : \mathbf{CAT} \to \mathbf{epMET}$ and $\bullet\mathbf{R} : \mathbf{epMET} \to \mathbf{CAT}$ of [BSN17]. More precisely, we have that, up to natural isomorphism, the functors $\bullet^{\mathbf{R}}$ and $\bullet\mathbf{R}$ can be factored as follows

$$\begin{tikzcd} \mathbf{CAT} \arrow[r, bend left, "\mathsf{po}"] \arrow[rr, bend left=50, "\bullet^{\mathbf{R}}"] & \mathbf{lpCAT} \arrow[l, bend left, "\mathsf{st}"] \arrow[r, bend left, "\mathsf{met}"] & \mathbf{epMET}, \arrow[l, bend left, "\mathsf{cat}"] \arrow[ll, bend left=50, "\bullet\mathbf{R}"] \end{tikzcd}$$

where $\mathsf{po} : \mathbf{CAT} \to \mathbf{lpCAT}$ maps a category C to its locally persistent category of persistent objects $C^{\mathbf{R}}$, and $\mathsf{st} : \mathbf{lpCAT} \to \mathbf{CAT}$ is a generalization of the space-time construction of [BSN17, Section 2], which we now describe. Given a locally persistent category $\mathscr{C}$, let $\mathsf{st}(\mathscr{C})$ be the category with objects $\mathbf{R} \times \mathsf{obj}(\mathscr{C})$, and such that for every $r, s \in \mathbf{R}$ and $x, y \in \mathscr{C}$ we have

$$\mathrm{Hom}_{\mathsf{st}(\mathscr{C})}\big((r,x),(s,y)\big) = \begin{cases} \mathrm{Hom}_{\mathscr{C}}(x,y)_{s-r}, & \text{if } s-r \geq 0 \\ \emptyset, & \text{otherwise.} \end{cases}$$

In particular, [BSN17, Theorem 3.6] follows from Theorem 4.6.10, above.

As a side remark, we note that cat is a kind of Yoneda embedding.

Remark 4.6.12. The functor cat can also be defined by recalling that the Yoneda embedding $[0,\infty]^{\mathsf{op}} \to \mathbf{Set}^{[0,\infty]}$ is monoidal (Lemma 2.6.8), where the monoidal structure on $[0,\infty]$ is sum and the monoidal structure on $\mathbf{Set}^{[0,\infty]}$ is given by Day convolution. Moreover, any functor $[0,\infty] \to \mathbf{Set}$ can be restricted to a functor $\mathbf{R}_+ \to \mathbf{Set}$, and the restriction operation $\mathbf{Set}^{[0,\infty]} \to \mathbf{Set}^{\mathbf{R}_+}$ is monoidal. The composite induces a functor from the category of $[0,\infty]$-enriched categories to the category of locally persistent categories. Finally, the category of large ep metric spaces (symmetric Lawvere spaces, [Law73]) is a subcategory of the category of $[0,\infty]$-enriched categories (Lawvere metric spaces). This gives a composite

$$\mathbf{epMET} \longrightarrow [0,\infty]^{\mathsf{op}}\text{-}\mathbf{CAT} \longrightarrow \mathbf{Set}^{[0,\infty]}\text{-}\mathbf{CAT} \longrightarrow \mathbf{Set}^{\mathbf{R}_+}\text{-}\mathbf{CAT} = \mathbf{lpCAT},$$

which is naturally isomorphic to cat.

4.7 The observable category of a locally persistent category

In [CCBS14], the *observable category* of persistent modules is defined. Two descriptions of this category are given, one ([CCBS14, Definition 2.3]) is as a direct construction, and the other one ([CCBS14, Corollary 2.13]) is as a quotient of the category of persistent modules by the subcategory of ephemeral persistent modules (i.e. persistent modules all of whose non-identity structure maps are trivial).

The purpose of the observable category is to define a category that is in a sense simpler than the category of persistent modules, but that still has enough information so that important invariants of persistent modules factor through this category. In particular, the property of being q-tame ([CSGO16, Section 2.8]), the undecorated persistent diagram ([CSGO16, Section 1.6]), and the interleaving distance between persistent modules are *observable invariants*, in the sense that they only depend on the image of the relevant persistent modules in the observable category. Moreover, in the observable category, any q-tame persistent module is interval-decomposable ([CCBS14, Corollary 3.8]), a fact that is not true in the category of persistent modules.

In [BP19], the notion of ephemeral persistent module is considered in the case of multi-dimensional persistent modules, and it is used to construct an observable category of multi-dimensional persistent modules, as a quotient of the category of multi-dimensional persistent modules by the category of ephemeral modules. It is proven that, in this generality, the observable category is equivalent to the subcategory of multi-dimensional modules that are sheaves for a convenient topology on the poset $\mathbf{R}^n$.

In this section, we associate an *observable locally persistent category* $\mathscr{C}^\#$ to every locally persistent category $\mathscr{C}$. We show that the interleaving distance is an observable invariant, in the sense that the observable locally persistent category gives rise to the same ep metric space as the original locally persistent category (Proposition 4.7.3). We also extend one of the main results of [BP19], namely, that the observable category of the category of persistent objects of a complete category is equivalent to the subcategory of right continuous persistent objects. We do this for 1-dimensional persistent objects for simplicity, but the same constructions work for higher dimensions.

Recall, from Section 2.6.2, that there is a lax monoidal functor $(-)^\# : \mathbf{Set}^{\mathbf{R}_+} \to \mathbf{Set}^{\mathbf{R}_+}$ given by

$$F^\# = \lim_{r>0} F^r,$$

where F^r is the r-shift to the left of F as in Definition 2.6.1, and a monoidal natural transformation $\eta : \mathrm{id}_{\mathbf{Set}^{\mathbf{R}_+}} \Rightarrow (-)^\#$. As discussed in Section 2.6.2, the functor $(-)^\#$ turns any functor $F : \mathbf{R}_+ \to \mathbf{Set}$ into a right continuous functor in a universal way, in the sense that, if G is right continuous, then morphisms $F \to G$ are in bijection with morphisms $F^\# \to G$, and this bijection is given by precomposition with the morphism $\eta^\#_F : F \to F^\#$.

Since the functor $(-)^\#$ is monoidal, it provides us with a change of enrichment

$$(-)^\# : \mathbf{lpCAT} \to \mathbf{lpCAT}.$$

This functor has the effect of turning the hom-persistent sets of a locally persistent category into right continuous persistent sets. Moreover, by Proposition 2.6.11, the natural transformation $\eta^{\#}$ is monoidal, so we have a natural transformation $\eta^{\#} : \mathsf{id}_{\mathbf{lpCAT}} \Rightarrow (-)^{\#}$ between the functors $\mathsf{id}_{\mathbf{lpCAT}}, (-)^{\#} : \mathbf{lpCAT} \to \mathbf{lpCAT}$. In particular, for every locally persistent category $\mathscr{C}$ we get a locally persistent functor $\eta^{\#}_{\mathscr{C}} : \mathscr{C} \to \mathscr{C}^{\#}$.

Definition 4.7.1. *The **observable locally persistent category** of a locally persistent category $\mathscr{C}$ is defined to be $\mathscr{C}^{\#}$.*

Before proceeding to prove some properties of this construction, we explain its relationship to the original definition of observable category given in [CCBS14]. Since it makes things simpler, we generalize their definition to persistent objects in an arbitrary category C. An **observable morphism** between persistent objects $X, Y \in C^{\mathbf{R}}$ consists of an element of the set $\lim_{r>0} \mathrm{Nat}(X, Y^r)$. This definition is equivalent to [CCBS14, Definition 2.2]. Using this notion of morphism, and the fact that $\mathbf{R}$ is a dense poset, one obtains a well-defined composition, and a category $\mathbf{Ob}_C$, the **observable category of persistent objects** of C.

Note that, for $X, Y \in C^{\mathbf{R}}$, we have

$$\lim_{r>0} \mathrm{Nat}(X, Y^r) = \lim_{r>0} \mathrm{Hom}_{C^{\mathbf{R}}}(X, Y)_r = \mathrm{Hom}_{C^{\mathbf{R}}}(X, Y)^{\#},$$

where all the equalities are by definition. We deduce the following.

Proposition 4.7.2. *For any category C, there is an isomorphism of categories*

$$\mathbf{Ob}_C \cong \left(\left(C^{\mathbf{R}}\right)^{\#} \right)_0. \qquad \square$$

We now prove that the interleaving distance of a locally persistent category is an observable invariant.

Proposition 4.7.3. *Let $\mathscr{C}$ be a locally persistent category and let $x, y \in \mathscr{C}$. Then $d_I^{\mathscr{C}}(x, y) = d_I^{\mathscr{C}^{\#}}(x, y)$.*

Proof. Let $\delta \in \mathbf{R}_+$. On the one hand, we have a locally persistent functor $\eta^{\#}_{\mathscr{C}} : \mathscr{C} \to \mathscr{C}^{\#}$ that is the identity on objects. So, if x and y are δ-interleaved in $\mathscr{C}$, they must be δ-interleaved in $\mathscr{C}^{\#}$.

On the other hand, a δ-interleaving $f : x \;{}_{\delta}\!\!\longleftrightarrow_{\delta}\; y : g$ in $\mathscr{C}^{\#}$, by definition, consists of elements $f \in \left(\mathrm{Hom}_{\mathscr{C}}(x, y)^{\#}\right)_{\delta}$ and $g \in \left(\mathrm{Hom}_{\mathscr{C}}(y, x)^{\#}\right)_{\delta}$ that compose to shifts of the identity appropriately. Let us unfold what this means. Recall that, by definition, we have that

$$\big(\mathrm{Hom}_{\mathscr{C}}(x,y)^{\#}\big)_{\delta} = \lim_{\delta'>\delta} \mathrm{Hom}_{\mathscr{C}}(x,y)_{\delta'}.$$

So, for every $\delta' > \delta$, there is $f_{\delta'} \in \mathrm{Hom}_{\mathscr{C}}(x,y)_{\delta'}$, such that the shift map $\mathbf{S}_{\delta',\delta''}$ maps $f_{\delta'}$ to $f_{\delta''}$, for $\delta < \delta' \leq \delta'' \in \mathbf{R}_{+}$. An analogous discussion applies to g. Let $\delta' > 0$ and consider $f_{\delta'} : x \to_{\delta'} y$ and $g : y \to_{\delta'} x$ as morphisms in the locally persistent category $\mathscr{C}$. By Proposition 2.6.11, we have a commutative diagram

$$\begin{tikzcd}
\mathrm{Hom}_{\mathscr{C}}(y,x) \otimes_{Day} \mathrm{Hom}_{\mathscr{C}}(x,y) \arrow[r, "\circ"] \arrow[d, "\eta^{\#} \otimes_{Day} \eta^{\#}"'] & \mathrm{Hom}_{\mathscr{C}}(x,x) \arrow[d, "\eta^{\#}"] \\
\mathrm{Hom}_{\mathscr{C}}(y,x)^{\#} \otimes_{Day} \mathrm{Hom}_{\mathscr{C}}(x,y)^{\#} \arrow[r, "\circ"] & \mathrm{Hom}_{\mathscr{C}}(x,x)^{\#},
\end{tikzcd}$$

relating the composition in $\mathscr{C}$ to the composition in $\mathscr{C}^{\#}$. Since f and g form an interleaving in $\mathscr{C}^{\#}$, we have that $g_{\delta'} \circ f_{\delta'} = \eta^{\#}(\mathbf{S}_{0,2\delta'}(\mathrm{id}_x))$ in $\mathrm{Hom}_{\mathscr{C}}(x,x)^{\#}$. By Lemma 2.6.13, this implies that for any $\delta'' > \delta'$ we have $g_{\delta''} \circ f_{\delta''} = \mathbf{S}_{0,2\delta''}(\mathrm{id}_x)$. Together with a symmetric argument, this shows that x and y are δ''-interleaved for every $\delta'' > \delta' > \delta$, in $\mathscr{C}$. This is enough to show that $d_I^{\mathscr{C}}(x,y) = \delta$ because δ'' and δ' can be taken to be arbitrarily close to δ, since $\mathbf{R}_{+}$ is dense. □

We conclude this section by characterizing the observable locally persistent category of a category of persistent objects as a locally persistent category of right continuous persistent objects. Recall from Section 2.6.2 that, for any category C, we let $C^{\mathbf{R}}_{\mathrm{right}}$ be the full subcategory of $C^{\mathbf{R}}$ spanned by right continuous persistent objects.

Proposition 4.7.4. *Let C be a complete category. Then we have an equivalence of locally persistent categories*

$$C^{\mathbf{R}}_{\mathrm{right}} \simeq \big(C^{\mathbf{R}}\big)^{\#}$$

given by the composite $C^{\mathbf{R}}_{\mathrm{right}} \to C^{\mathbf{R}} \to \big(C^{\mathbf{R}}\big)^{\#}$, with the first functor being the inclusion, and the second functor being $\eta^{\#}_{C^{\mathbf{R}}}$.

Proof. The result follows directly from the fact that, for $X, Y : \mathbf{R} \to C$, we have

$$\begin{aligned}
\mathrm{Hom}_{(C^{\mathbf{R}})^{\#}}(X, Y) &= \lim_{r>0} \mathrm{Hom}_{C^{\mathbf{R}}}(X, Y)^r \\
&= \lim_{r>0} \mathrm{Hom}_{C^{\mathbf{R}}}(X, Y^r) \\
&\cong \mathrm{Hom}_{C^{\mathbf{R}}}(X, \lim_{r>0} Y^r) \\
&\cong \mathrm{Hom}_{C^{\mathbf{R}}}(X, Y^{\#}).
\end{aligned}$$

□

Chapter 5

Constructing locally persistent categories

As we will see in the examples in Chapter 6, it is often easy to define a locally persistent category directly, by specifying objects, morphisms, composition, and identities, much in the same way that many categories are usually described directly. In this chapter, we provide more principled and systematic ways of constructing locally persistent categories.

Locally persistent categories are categories whose hom-sets are parametrized by the poset $\mathbf{R}_+$. In Section 5.1, we argue that the hom-sets of many categories are more naturally parametrized by posets other than $\mathbf{R}_+$. We then use the change of enrichment construction to construct locally persistent categories from categories whose hom-sets are parametrized by other posets. Our main example is given by locally multi-persistent categories of multi-persistent objects of a locally persistent category. We explain this construction in Section 5.1.2. We also recall the main constructions of [BSS13] which allow one to define a locally persistent category structure on categories of generalized persistent objects.

In Section 5.2, we show that every category with a flow, in the sense of [SMS18], has an associated locally persistent category with the same objects and interleaving distance, thus letting one use the language of locally persistent categories to study the interleaving distance of a category with a flow. We argue that the categorical framework of locally persistent categories is more amenable to abstract reasoning than the framework of categories with a flow, and see categories with a flow as a great source of examples of locally persistent categories.

5.1 Change of enrichment

A locally persistent category $\mathscr{C}$ has, for each pair of objects $x, y \in \mathscr{C}$ and each $\varepsilon \in \mathbf{R}_+$, a set of ε-morphisms, denoted by $\mathrm{Hom}_{\mathscr{C}}(x,y)_\varepsilon$. So, in a precise sense, the collection of all morphisms between x and y is a set parametrized by the poset $\mathbf{R}_+$.

It is often the case that, for a collection of objects $\mathscr{C}$, the morphisms between two objects $x, y \in \mathscr{C}$ is most naturally parametrized by a poset other than $\mathbf{R}_+$.

Definition 5.1.1. *Let Q be a monoidal poset and endow the category* $\mathbf{Set}^Q$ *with the monoidal product given by Day convolution. A* ***locally Q-persistent category*** *is a category enriched in* $\mathbf{Set}^Q$.

One can unfold Definition 5.1.1 and obtain a definition entirely analogous to Definition 3.1.1, the only difference being that instead of $\mathbf{R}_+$ we have Q, and instead of $+$ we have the monoidal product of Q.

Example 5.1.2. Let C be any category and let $n \in \mathbb{N}$. The functor category $C^{\mathbf{R}^n}$ has a natural structure of locally $\mathbf{R}^n_+$-persistent category, where, for $\vec{v} \in \mathbf{R}^n_+$ and $X, Y \in C^{\mathbf{R}^n}$, we have

$$\mathrm{Hom}_{C^{\mathbf{R}^n}}(X,Y)_{\vec{v}} = \mathrm{Nat}(X, Y^{\vec{v}}),$$

where $Y^{\vec{v}}(\vec{w}) = Y(\vec{v}+\vec{w})$ for every $\vec{w} \in \mathbf{R}^n$.

In applications (see, e.g., Section 5.1.2, Section 6.4, and Section 6.5), we construct locally $\mathbf{R}^2_+$-persistent categories. We refer to these as **locally bi-persistent categories**.

In order to define an interleaving distance for a locally Q-persistent category, one can first turn the locally Q-persistent category into a locally persistent category, and then use the usual interleaving distance. A natural way of turning a $\mathbf{Set}^Q$-enriched category into a $\mathbf{Set}^{\mathbf{R}_+}$-enriched category is by constructing a lax monoidal functor $\mathbf{Set}^Q \to \mathbf{Set}^{\mathbf{R}_+}$, and then using the change of enrichment construction (Definition 2.4.2). Although not expressed in the language of locally Q-persistent categories, this is essentially the approach taken in [BSS13].

In [BSS13], two ways of constructing lax monoidal functors of the form $\mathbf{Set}^Q \to \mathbf{Set}^{\mathbf{R}_+}$ are studied. The simplest case is when we already have a lax monoidal functor $\mathbf{R}_+ \to Q$. We go over this construction in Section 5.1.1. The other case is when we have a monoidal functor $Q^{\mathrm{op}} \to \overline{\mathbf{R}}_+{}^{\mathrm{op}} = [0,\infty]^{\mathrm{op}}$. We explain this construction in Section 5.1.3.

In this thesis, the most important examples of locally multi-persistent categories are categories of persistent objects of a locally persistent category. In Section 5.1.2, given a locally persistent category $\mathscr{C}$ and $n \in \mathbb{N}$, we construct a locally $\mathbf{R}_+^{n+1}$-persistent category of functors $\mathbf{R}^n \to \mathscr{C}_0$ that takes into account both the locally persistent category structure of $\mathscr{C}$ as well as the shifts that come from the indexing by $\mathbf{R}^n$.

5.1.1 Superlinear families

Let Q be a monoidal poset and assume given a lax monoidal functor $\mathbf{R}_+ \to Q$. The category $\mathbf{Set}^Q$ is again monoidal, endowing it with Day convolution, and precomposition with $\mathbf{R}_+ \to Q$ provides us with a lax monoidal functor

$$\mathbf{Set}^Q \to \mathbf{Set}^{\mathbf{R}_+}.$$

Example 5.1.3. Let C be any category and let $n \in \mathbb{N}$. Consider the $\mathbf{Set}^{\mathbf{R}^n}$-enrichment of the functor category $C^{\mathbf{R}^n}$ described in Example 5.1.2.

Every $\vec{v} \in \mathbf{R}_+^n$ induces a monoidal functor $\mathbf{R}_+ \to \mathbf{R}^n$ given by mapping ε to $\varepsilon\vec{v}$. The change of enrichment construction then endows $C^{\mathbf{R}^n}$ with a locally persistent category structure $\mathscr{C}$, where, for $\varepsilon \in \mathbf{R}_+$ and $X, Y \in C^{\mathbf{R}^n}$, we have

$$\mathrm{Hom}_{\mathscr{C}}(X,Y)_\varepsilon = \mathrm{Nat}(X, Y^{\varepsilon\vec{v}}),$$

where, as before, $Y^{\varepsilon\vec{v}}(\vec{w}) = Y(\varepsilon\vec{v} + \vec{w})$ for every $\vec{w} \in \mathbf{R}^n$.

Example 5.1.2 is very important, as it provides an interleaving distance for categories of multi-parameter persistent objects. These distances were carefully studied in [Les12]. The same construction allows one to turn any $\mathbf{Set}^{\mathbf{R}_+^n}$-enrichment into a $\mathbf{Set}^{\mathbf{R}_+}$-enrichment, given a vector $\vec{v} \in \mathbf{R}_+^n$. Note that this change of enrichment depends on the choice of vector $\vec{v} \in \mathbf{R}_+^n$, so the interleaving distance we obtain also depends on this vector. A straightforward, but very important property of this construction is the following, which says that, as long as the vector has non-zero coordinates, the induced metric is uniquely defined up to a multiplicative constant; in particular, the topology this metric induces is independent of the choice of vector.

Proposition 5.1.4. *Let $n \in \mathbb{N}$ and let $\mathscr{C}$ be an $\mathbf{R}_+^n$-locally persistent category. Assume that $\vec{v}, \vec{w} \in \mathbf{R}_+^n$ are such that all of their coordinates are strictly positive. Then, the interleaving distance on $\mathscr{C}$ obtained by the change of enrichment using $\vec{v}$ is bi-Lipschitz equivalent to the one obtained by the change of enrichment using $\vec{w}$.*

Proof. Since the coordinates of $\vec{v}$ and $\vec{w}$ are strictly positive, there are $\varepsilon_1, \varepsilon_2 \in \mathbf{R}_+$ such that $\varepsilon_1 \vec{v} \geq \vec{w}$ and such that $\varepsilon_2 \vec{w} \geq \vec{v}$. This means that, for $X, Y \in \mathscr{C}$, if X and Y are δ-interleaved using the locally persistent structure induced by $\vec{v}$, then they are $\varepsilon_2\delta$-interleaved using the locally persistent structure induced by $\vec{w}$. And, conversely, if they are δ-interleaved using the locally persistent structure induced by $\vec{w}$, then they are $\varepsilon_1\delta$-interleaved using the locally persistent structure induced by $\vec{v}$. □

The formalism of change of enrichment allows us to prove useful properties. As an example, we have the following.

Lemma 5.1.5. *Let Q be a monoidal poset and let $F : \mathbf{R}_+ \to Q$ be a strong monoidal functor. Let $\mathscr{C}$ be a locally Q-persistent category and let $\mathscr{C}_F$ be the locally persistent category obtained using the change of enrichment given by $F^* : \mathbf{Set}^Q \to \mathbf{Set}^{\mathbf{R}_+}$. If $\mathscr{C}$ is powered (resp. copowered) by representables, then $\mathscr{C}_F$ is powered (resp. copowered) by representables.* □

Proof. We prove the powering case, the other case being dual. Let $\varepsilon \in \mathbf{R}_+$ and let $X \in \mathscr{C}_F$. The power X^ε in the $\mathbf{Set}^{\mathbf{R}_+}$-enrichment is then given by $X^{F(\varepsilon)}$, where this second powering is in the $\mathbf{Set}^Q$ enrichment. □

We conclude this section by recalling the specific change of enrichment given in [BSS13, Section 2.5]. This construction allows one to endow categories of the form C^P, for P a poset, with a locally persistent category structure.

Let P be a poset, and let Trans_P be the **poset of translations** of P. A **translation** of P is a monotonic map $\Gamma : P \to P$ such that for all $x \in P$ we have $x \leq \Gamma(x)$. The partial order in Trans_P is given by $\Gamma \leq \Delta \in \mathsf{Trans}_P$ if and only if $\Gamma(x) \leq \Delta(x)$ for all $x \in P$. Note that Trans_P is a monoidal poset, with monoidal product given by composition of translations.

An example of a $\mathbf{Set}^{\mathsf{Trans}_P}$-enriched category, studied in [BSS13], is the following category of generalized persistent modules. Let C be a category and consider the functor category C^P. This category has a $\mathbf{Set}^{\mathsf{Trans}_P}$-enrichment, where

$$\mathrm{Hom}(X, Y)_\Gamma = \mathrm{Nat}(X, Y \circ \Gamma).$$

In [BSS13, Section 2.5], lax monoidal functors of the form $\mathbf{R}_+ \to \mathsf{Trans}_P$ are called **superlinear families**.

Remark 5.1.6. Let P be a poset, let C be a category, and let $\mathbf{R}_+ \to \mathsf{Trans}_P$ be a superlinear family. The locally persistent category structure induced on C^P is weakly-powered (Definition 5.2.12) since this structure comes from the flow $\mathbf{R}_+ \to C^P$ given by precomposition with the superlinear family. We can thus consider the strong interleaving distance on C^P, which is the one described in [BSS13, Definition 2.5.1], or the weak interleaving distance (recall the discussion in Section 5.2.2). Note that these two coincide if the superlinear family is a strong monoidal functor.

5.1.2 Persistent objects of a locally persistent category

We now give one of the most important examples of this thesis. Given a locally persistent category $\mathscr{C}$ we define a locally bi-persistent category $\mathscr{C}^{\mathbf{R}}$ whose underlying category is the category of functors $\mathbf{R} \to \mathscr{C}_0$. That is, the objects of $\mathscr{C}^{\mathbf{R}}$ are the persistent objects in the underlying category of $\mathscr{C}$. For the morphisms, the idea is that, for $\varepsilon, \delta \in \mathbf{R}_+$, an (ε,δ)-morphism in $\mathscr{C}^{\mathbf{R}}$ is a natural transformation that shifts the persistence degree by ε and whose components are δ-morphisms of $\mathscr{C}$.

Definition 5.1.7. *Consider the category $\mathscr{C}^{\mathbf{R}}$ that has as objects the (standard) functors $\mathbf{R} \to \mathscr{C}_0$ from the poset $\mathbf{R}$ to the underlying category of $\mathscr{C}$. This category admits a $\mathbf{Set}^{\mathbf{R}_+ \times \mathbf{R}_+}$-enrichment, given as follows. For $X, Y \in \mathscr{C}^{\mathbf{R}}$ and $\varepsilon, \delta \in \mathbf{R}_+$, let*

$$\mathrm{Hom}_{\mathscr{C}^{\mathbf{R}}}(X,Y)_{(\varepsilon,\delta)} = \mathrm{Nat}(X, Y^{\varepsilon})_{\delta},$$

where Y^{ε} is ε-shift to the left of Y and an element $\alpha \in \mathrm{Nat}(X,Y)_{\delta}$ consists of a family $\alpha_r \in \mathrm{Hom}_{\mathscr{C}}(X(r), Y(r))_{\delta}$ of δ-morphisms of $\mathscr{C}$, for $r \in \mathbf{R}$, such that $\varphi^Y_{r,s} \circ \alpha_r = \alpha_s \circ \varphi^X_{r,s} \in \mathrm{Hom}_{\mathscr{C}}(X(r), Y(s))_{\delta}$ for all $r \leq s \in \mathbf{R}_+$.

Given $\vec{v} = (v_1, v_2) \in \mathbf{R}_+ \times \mathbf{R}_+$ let $\mathscr{C}^{\mathbf{R}}_{\vec{v}}$ denote the locally persistent category whose structure is given by the change of enrichment construction using $\vec{v}$, as explained above.

If $(\mathscr{C}, \mathscr{W})$ is a relative locally persistent category, let $\mathscr{W}^{\to}$ denote the class of natural transformations of $(\mathscr{C}_0)^{\mathbf{R}}$ with all of its components in $\mathscr{W}$. This endows $\mathscr{C}^{\mathbf{R}}_{\vec{v}}$ with a relative locally persistent category structure. This construction is well behaved.

Lemma 5.1.8. *Let $\mathscr{C}$ be a locally persistent category. If $\mathscr{C}$ admits copowers or powers by representables, then so does $\mathscr{C}^{\mathbf{R}}_{\vec{v}}$.*

Proof. We prove the powering case, the other case being dual. By Lemma 5.1.5, it is enough to show that $\mathscr{C}^{\mathbf{R}}$ is powered by representables whenever $\mathscr{C}$ is. Let $X \in \mathscr{C}^{\mathbf{R}}$ and let $\varepsilon, \delta \in \mathbf{R}_+$. Then the power $X^{(\varepsilon,\delta)} : \mathbf{R} \to \mathscr{C}_0$ is given by $X^{(\varepsilon,\delta)}(r) = X(r+\varepsilon)^{\delta}$, where the powering on the right hand side of the equality is the powering in $\mathscr{C}$. □

The following will be very useful in examples (Section 6.4).

Proposition 5.1.9. *Let $(\mathscr{C}, \mathscr{W})$ be a relative locally persistent category and let $\vec{v} \in \mathbf{R}_+ \times \mathbf{R}_+$. Assume that $\mathscr{C}_0$ admits pullbacks, products, and sequential limits, and that $\mathscr{C}$ is powered by representables and the powering operation respects pullbacks, products, sequential limits, and morphisms of $\mathscr{W}$. Assume further that $\mathscr{W}$ is closed under sequential limits of $\mathscr{C}_0$. Then, the quotient interleaving distance on $(\mathscr{C}^{\mathbf{R}}_{\vec{v}}, \mathscr{W}^{\to})$ is intrinsic and complete, and it satisfies*

$$\begin{aligned}\left(d_I^{\mathscr{C}^{\mathbf{R}}_{\vec{v}}}\right)_{/\simeq}(X,Y) &= \inf\left\{\delta : \exists X' \simeq X, Y' \simeq Y,\ X' \text{ and } Y' \text{ are } \delta\text{-interleaved}\right\}\\ &= \inf\left\{\delta : \exists \text{ morphisms of } \mathscr{W}^{\to},\ X' \to X \text{ and } Y' \to Y\right.\\ &\qquad\qquad \left.\text{such that } X' \text{ and } Y' \text{ are } \delta\text{-interleaved}\right\}.\end{aligned}$$

Proof. The underlying category of $\mathscr{C}^{\mathbf{R}}$ is the functor category $(\mathscr{C}_0)^{\mathbf{R}}$ so it admits pullbacks, products, and sequential limits since limits are computed pointwise. The locally persistent category $\mathscr{C}^{\mathbf{R}}_{\vec{v}}$ is powered by representables by Lemma 5.1.8. These powers respect pullbacks, products, sequential limits, and morphisms of $\mathscr{W}^{\to}$ since all of these are defined pointwise. Also, $\mathscr{W}^{\to}$ is closed under sequential limits of $(\mathscr{C}^{\mathbf{R}})_0$. It follows that $\mathscr{C}^{\mathbf{R}}$ admits weighted pullbacks (Proposition 3.2.12), weighted sequential limits (Proposition 3.2.15), and terminal midpoints (Proposition 3.2.19), so the quotient interleaving distance is intrinsic (Corollary 4.4.5) and complete (Theorem 4.3.3). The characterization of the quotient interleaving in the statement follows from Theorem 4.1.4. □

Remark 5.1.10. Given $\mathscr{C}$ a locally persistent category and any $n \in \mathbb{N}$, the construction given in Definition 5.1.7 generalizes immediately to endow $(\mathscr{C}_0)^{\mathbf{R}^n}$ with a locally $\mathbf{R}_+^{n+1}$-persistent category structure. It is clear that Proposition 5.1.9 also generalizes to this case.

A simple but powerful observation is the following.

Proposition 5.1.11. *Let $F : \mathscr{C} \to \mathscr{D}$ be a locally persistent functor between relative locally persistent categories that maps acyclic morphisms to acyclic morphisms. Then,* *by applying F objectwise, we obtain a locally bi-persistent functor $F_* : \mathscr{C}^{\mathbf{R}} \to \mathscr{D}^{\mathbf{R}}$ that maps acyclic morphisms to acyclic morphism.* □

5.1.3 Sublinear projections

We start by motivating the discussion with an example based on [BSS13, Section 2.4].

Example 5.1.12. Consider the set of all order-preserving functions $\Gamma : \mathbf{R} \to \mathbf{R}$ such that $\sup_{r\in\mathbf{R}} |\Gamma(r) - r| < \infty$ and $\Gamma(r) \geq r$ for all $r \in \mathbf{R}$. Denote this set by S and note that it is closed under composition. For $\Gamma, \Delta \in S$, say that $\Gamma \leq \Delta$ if $\Gamma(r) \leq \Delta(r)$ for all $r \in \mathbf{R}$. This endows S with the structure of a monoidal poset.

Let C be any category. There is a $\mathbf{Set}^S$-enrichment of $C^{\mathbf{R}}$ given by $\mathrm{Hom}(X, Y)_\Gamma = \mathrm{Nat}(X, Y \circ \Gamma)$. Given $X, Y \in \mathscr{C}^{\mathbf{R}}$ and $\Gamma, \Delta \in S$, we say that X and Y are **(Γ, Δ)-interleaved** if there are natural transformations $X \to Y \circ \Gamma$ and $Y \to X \circ \Delta$ that compose to the structure morphisms $X \to X \circ \Delta \circ \Gamma$ and $Y \to Y \circ \Gamma \circ \Delta$.

Let $F : S \to \mathbf{R}_+$ send Γ to $\sup_{r\in\mathbf{R}} |\Gamma(r) - r|$. For $\varepsilon \in \mathbf{R}_+$, we say that X and Y are ε-interleaved if they are (Γ, Δ)-interleaved with $F(\Gamma), F(\Delta) \leq \varepsilon$. Taking an infimum, we get an interleaving ep metric on $C^{\mathbf{R}}$ ([BSS13, Theorem 2.3.5]).

We will see how to use the map F described in Example 5.1.12 to give a locally persistent structure on $C^{\mathbf{R}}$ such that the interleaving distance of this locally persistent category coincides with the interleaving distance of the example.

Remark 5.1.13. Example 5.1.12 is in fact too simple: It is easy to see that the distance induced by the notion of interleaving in Example 5.1.12 is the usual interleaving distance in $C^{\mathbf{R}}$, since if we let $S_r : \mathbf{R} \to \mathbf{R}$ be given by adding r, then $F(S_r) = r$ and for every $\Gamma : \mathbf{R} \to \mathbf{R}$ such that $F(\Gamma) = r$ we have $\Gamma \leq S_r$.

Nonetheless the example is useful for understanding the constructions in this section. A more interesting and useful example is Example 5.1.17.

Let P be a poset. A **monotone sublinear projection** is given by a lax monoidal functor $\mathsf{Trans}_P^{\mathsf{op}} \to [0,\infty]^{\mathsf{op}}$, where $[0,\infty]$ is a monoidal poset with monoidal product given by addition of real numbers, such that $\infty + r = r + \infty = \infty$ for every $r \in [0,\infty]$.

Monoidal functors of the form $\mathscr{V}^{\mathsf{op}} \to \mathscr{W}^{\mathsf{op}}$ for $\mathscr{V}$ and $\mathscr{W}$ monoidal categories, are usually referred to as **oplax monoidal** functors $\mathscr{V} \to \mathscr{W}$. We now explain how to get a lax monoidal functor $\mathbf{Set}^Q \to \mathbf{Set}^{\mathbf{R}_+}$ out of an oplax monoidal functor $Q \to \overline{\mathbf{R}}_+$.

By precomposition, a functor $F : Q \to \overline{\mathbf{R}}_+$, gives us a functor

$$L : \mathbf{Set}^{\overline{\mathbf{R}}_+} \to \mathbf{Set}^Q.$$

We are seeking a functor going the opposite way so we consider the right adjoint of L which we denote by R. This right adjoint exists since $\mathbf{Set}^{\overline{\mathbf{R}}_+}$ and $\mathbf{Set}^Q$ are locally presentable and L preserves colimits, as they are computed pointwise. Nonetheless, the right adjoint of L is easy to describe concretely.

Lemma 5.1.14. *In the situation above, the right adjoint R can be defined as*

$$R(A)(\varepsilon) = \mathrm{Nat}\,(\mathscr{Y}(\varepsilon)\circ F, A) \cong \coprod_{\substack{q\in Q\\ F(q)\leq\varepsilon}} A(q),$$

for $A \in \mathbf{Set}^Q$, where $\mathscr{Y} : [0,\infty]^{\mathrm{op}} \to \mathbf{Set}^{[0,\infty]}$ is the (co)Yoneda functor (Definition 2.6.5). Moreover, R is lax monoidal.

Proof. Let $B \in \mathbf{Set}^{\overline{\mathbf{R}}_+}$. By the (co)Yoneda lemma, we have

$$B \cong \int^{\varepsilon\in\overline{\mathbf{R}}_+} B(\varepsilon)\times\mathscr{Y}(\varepsilon).$$

We can then compute

$$\begin{aligned}
\mathrm{Nat}(B\circ F, A) &\cong \mathrm{Nat}\left(\left(\int^{\varepsilon\in\overline{\mathbf{R}}_+} B(\varepsilon)\times\mathscr{Y}(\varepsilon)\right)\circ F, A\right)\\
&\cong \mathrm{Nat}\left(\int^{\varepsilon\in\overline{\mathbf{R}}_+} B(\varepsilon)\times(\mathscr{Y}(\varepsilon)\circ F), A\right)\\
&\cong \int_{\varepsilon\in\overline{\mathbf{R}}_+} \mathrm{Nat}\,(B(\varepsilon)\times(\mathscr{Y}(\varepsilon)\circ F), A)\\
&\cong \int_{\varepsilon\in\overline{\mathbf{R}}_+} \mathrm{Hom}_{\mathbf{Set}}\,(B(\varepsilon), \mathrm{Nat}(\mathscr{Y}(\varepsilon)\circ F, A))\\
&\cong \mathrm{Nat}\,(B, \mathrm{Nat}(\mathscr{Y}(-)\circ F, A)) = \mathrm{Nat}(B, R(A)).
\end{aligned}$$

To see that R is lax monoidal, one just uses the adjunction $L \dashv R$ and the fact that L is oplax monoidal. This is a formal argument, and is part of what is often referred to as doctrinal adjunction ([Kel74]). □

We deduce the following.

Proposition 5.1.15 (cf. [BSS13, Section 2.3])**.** *Let P be a poset and let C be a category. Every monotone sublinear projection $F : \mathsf{Trans}_P \to [0,\infty]$ induces a locally persistent category structure on C^P by changing the enrichment of C^P on $\mathbf{Set}^{\mathsf{Trans}_P}$ along the lax* *monoidal functor* $\mathbf{Set}^{\mathsf{Trans}_P} \to \mathbf{Set}^{\mathbf{R}_+}$ *that is the composite of the right adjoint of the oplax monoidal functor* $\mathbf{Set}^{\overline{\mathbf{R}}_+} \to \mathbf{Set}^{\mathsf{Trans}_P}$ *given by precomposition with the monotone sublinear projection* F*, with the restriction functor* $\mathbf{Set}^{\overline{\mathbf{R}}_+} \to \mathbf{Set}^{\mathbf{R}_+}$. □

Remark 5.1.16. The interleaving distance of a locally persistent category obtained using the procedure described in Proposition 5.1.15 coincides with the distance described in [BSS13, Definition 2.3.2]. To see this, let P be a poset and C be a category. Then, by definition of the change of enrichment functor in Lemma 5.1.14, the ε-morphisms from X to Y, objects of the locally persistent category C^P, obtained by using Proposition 5.1.15 are exactly the natural transformations $X \to Y \circ \Gamma$ for $\Gamma \in \mathsf{Trans}_P$ with $F(\Gamma) \leq \varepsilon$. Moreover, composition in the locally persistent structure is defined using the composition in the $\mathbf{Set}^{\mathsf{Trans}_P}$ enrichment, so ε-interleavings in the locally persistent structure correspond to (Γ, Δ)-interleavings such that $F(\Gamma), F(\Delta) \leq \varepsilon$.

A good source of monotone sublinear projections are posets endowed with a metric.

Example 5.1.17 (cf. [BSS13, Section 2.4]). Let P be a poset and let d be a Lawvere metric on P, that is, an ep metric that is not necessarily symmetric. Assume that d satisfies $d(x,y) \leq d(x,z)$ whenever $x \leq y \leq z \in P$. Then, the formula

$$\sup\{d(x, \Gamma(x)) : x \in P\}$$

defines a monotone sublinear projection $\mathsf{Trans}_P \to \overline{\mathbf{R}}_+$.

An interesting instance of this construction is the case of subsets of a fixed metric space M. The subsets of M are ordered by inclusion, and there is a Lawvere metric on the set of subsets of M given by the Hausdorff distance.

5.2 Categories with a flow

Let C be a category and let $\mathrm{End}(C)$ be the category of endomorphisms of C. The objects of $\mathrm{End}(C)$ are functors $C \to C$, and the morphisms are natural transformations. This is a monoidal category, with monoidal product given by composition of endofunctors. In [SMS18], the data of a category C together with a lax monoidal functor $\mathscr{T} : \mathbf{R}_+ \to \mathrm{End}(C)$ is called a **category with a flow**, and to every category with a flow they assign an ep metric, called the interleaving distance, on the collection of objects of C.

It is useful to have some examples in mind. A standard one is any category of persistent objects $C^{\mathbf{R}}$, where the flow is given by shifting to the left. Another source of examples is the following.

Example 5.2.1. Let C be a category and $c \in C$. An $\mathbf{R}_+$**-object structure** on c is given by a monoid morphism $\psi : \mathbf{R}_+ \to \mathrm{Hom}_C(c,c)$. This gives a flow on the slice category C/c that maps $r \in \mathbf{R}_+$ and $f : x \to c$ to $\psi(r) \circ f$.

As a concrete example, one can take $C = \mathbf{Top}$ and $c = \mathbb{R}$ as a topological space, to get a flow on the category of $\mathbb{R}$-filtered topological spaces.

The notion of interleaving makes sense in any category with a flow, so the collection of objects of a category with a flow can be endowed with an ep metric: the interleaving distance.

In this section, we explain how every flow $\mathscr{T} : \mathbf{R}_+ \to \mathrm{End}(C)$ gives rise to a locally persistent category $C_{\mathscr{T}}$ with the same objects as C and much of the same categorical structure of C, in such a way that the interleaving distance of the locally persistent category thus obtained coincides with the interleaving distance of the category with a flow $(C, \mathscr{T})$ in the sense of [SMS18]. The main idea is to set

$$\mathrm{Hom}_{C_{\mathscr{T}}}(x,y)_\varepsilon = \mathrm{Hom}_C(x, \mathscr{T}_\varepsilon(y)).$$

This procedure is useful for a couple of reasons. Firstly, we note that, since categories with a flow are defined as categories together with a lax monoidal action of $\mathbf{R}_+$, many proofs become a bit lengthy, not because they are conceptually complicated, but because pasting diagrams have to be constructed in order to prove some coherences. Examples of this phenomenon are the proof of the triangle inequality of the interleaving distance of a category with a flow ([SMS18, Theorem 2.7]) and the completeness result of Cruz ([Cru19, Theorem 3]). The procedure that assigns a locally persistent category to each category with a flow encapsulates these coherences and uses them once and for all. As a consequence, to prove the triangle inequality for the interleaving distance of a category with a flow, one can use the corresponding fact for the interleaving distance of a locally persistent category (Lemma 3.1.10) which has a very short proof that exactly matches the proof that isomorphisms compose to isomorphisms in any category. Similarly, our completeness result, although not weaker nor stronger than the one of [Cru19] (see the discussion in Section 5.2.4), does not involve any coherences or pasting diagrams (Theorem 4.3.1), and exactly matches the proof that a transfinite composition of isomorphisms is an isomorphism in any category.

Secondly, the results in this section can be interpreted as a way of importing our general metric results to the world of categories with a flow. For example, in Section 5.2.4, we give conditions under which the interleaving distance of a category with a flow is intrinsic and complete, using the results of Chapter 4. So one can think of this procedure as a way of making the arguments in categories with a flow more closely match standard arguments in category theory.

As a final remark, we note that, as we point out in Remark 6.9.6, there are interleaving distances that don't arise naturally from a flow, but that are the interleaving distance of a natural locally persistent structure.

This section is structured as follows. In Section 5.2.1, we explain how to assign a locally persistent category to each category with a flow in a metric preserving way. This is done with a general categorical construction, which can be interpreted as an enriched Kleisli category construction.

In Section 5.2.3, we compare categories with a flow and locally persistent categories and we characterize the locally persistent categories that arise from categories with a flow. The comparison for general flows is more subtle than one may wish, but becomes very simple in the case of categories where the flow is a strong monoidal functor (which is the case in many of the most relevant applications).

In Section 5.2.4, we import some of our general metric results to categories with a flow, and we discuss the relationship between our completeness result and the one proven in [Cru19].

Finally, in Section 5.2.5, we explain how a straightforward generalization of flows, that of Q-flows, allows one to endow a category with the structure of a locally Q-persistent category (Definition 5.1.1).

5.2.1 The enriched Kleisli category construction

We start with the following general procedure, which we call the *enriched Kleisli category construction*. It is hard to know who to attribute it to, but we note that the construction is considered in [Mel].

Let $(\mathcal{V}, \otimes, 1)$ be a monoidal category. A lax monoidal functor $F : \mathcal{V} \to \mathrm{End}(C)$ induces a $\mathbf{Set}^{\mathcal{V}}$-enrichment of C, where the monoidal structure of $\mathbf{Set}^{\mathcal{V}}$ is given by Day convolution. We will denote the enriched category thus obtained by C_F.

We explain this construction; a dual construction gives a $\mathbf{Set}^{\mathcal{V}}$-enrichment of C, for a lax monoidal functor $G : \mathcal{V}^{\mathrm{op}} \to \mathrm{End}(C)$. We start by defining the hom-functors by

$$\mathrm{Hom}_{C_F}(c,d)_a = \mathrm{Hom}_C(c, F(a,d)).$$

The functoriality follows directly from the functoriality of F. To see that this provides an enrichment, we must define identities and composition.

Since F is monoidal, we have a morphism $c \to F(1,c)$ natural in $c \in C$, which, by Yoneda, gives a natural transformation of functors $\mathscr{Y}(1) \to \mathrm{Hom}_C(c, F(-,c)) = \mathrm{Hom}_{C_F}(c,c)$. This gives the identity morphism of $c \in C$ in the $\mathbf{Set}^{\mathcal{V}}$-enrichment of C.

The composite of $f \in \mathrm{Hom}_{C_F}(c,d)_a$ and $g \in \mathrm{Hom}_{C_F}(d,e)_b$ is given by the composite $x \xrightarrow{f} F(a,y) \xrightarrow{F(a,g)} F(a, F(b,z)) \to F(a \otimes b, z)$, where the last morphism comes from the monoidal structure of F. To see that composition is compatible with the functoriality of the hom objects, it is better to describe it as follows. Assume given $c,d,e \in C$. To get a morphism $\mathrm{Hom}(c,d) \otimes \mathrm{Hom}(d,e) \to \mathrm{Hom}(c,e)$ we start by computing the domain, using the Day convolution formula

$$\begin{aligned}
&\big(\mathrm{Hom}_{C_F}(c,d) \otimes \mathrm{Hom}_{C_F}(d,e)\big)(a) \\
\cong &\int^{x,y\in\mathcal{V}} \mathrm{Hom}_{\mathcal{V}}(x \otimes y, a) \times \mathrm{Hom}_{C_F}(c,d)_x \times \mathrm{Hom}_{C_F}(d,e)_y \\
\cong &\int^{x,y\in\mathcal{V}} \mathrm{Hom}_{\mathcal{V}}(x \otimes y, a) \times \mathrm{Hom}_C(c, F(x,d)) \times \mathrm{Hom}_C(d, F(y,e))
\end{aligned}$$

Let $(f,m,n) \in \mathrm{Hom}_{\mathcal{V}}(x \otimes y, a) \times \mathrm{Hom}_C(c, F(x,d)) \times \mathrm{Hom}_C(d, F(y,e))$. This gives $f_* : F(x, F(y,e)) \to F(x \otimes y, e) \to F(a,e)$ by functoriality, and the fact that F is monoidal. We can then form the composite $f_* \circ F(x,n) \circ m : c \to F(a,e)$. The assignment $(f,m,n) \mapsto f_* \circ F(x,n) \circ m$ gives a cowedge $\mathrm{Hom}_{\mathcal{V}}(- \otimes -, a) \times \mathrm{Hom}_C(c, F(-,d)) \times \mathrm{Hom}_C(d, F(-,e)) \to \mathrm{Hom}(c, F(a,e))$, natural in $a \in \mathcal{V}$. This induces a natural transformation $\mathrm{Hom}(c,d) \otimes \mathrm{Hom}(d,e) \to \mathrm{Hom}(c,e)$, by the universal property of the coend.

Unitality and associativity are straightforward, although lengthy: one must use pasting diagrams that are very similar to the ones considered in the proof of [SMS18, Theorem 2.7]. We won't give the details here.

In the next result, we specialize the above construction to the case when $\mathcal{V} = \mathbf{R}_+$, to get an operation $\mathbf{FlCAT} \to \mathbf{lpCAT}$, that assigns a locally persistent category to each category with a flow.

Proposition 5.2.2. *Let $(C, \mathcal{T})$ be a category with a flow. There is a corresponding locally* *persistent category* $C_{\mathscr{T}}$ *with* $\mathsf{obj}(C_{\mathscr{T}}) = \mathsf{obj}(C)$*, and such that for every* $x, y \in C_{\mathscr{T}}$ *we have* $\mathrm{Hom}_{C_{\mathscr{T}}}(x,y)_\delta = \mathrm{Hom}_C(x, \mathscr{T}_\delta(y))$ *and composition is given by mapping*

$$(g, f) \in \mathrm{Hom}_C(y, \mathscr{T}_\delta(z)) \times \mathrm{Hom}_C(x, \mathscr{T}_\varepsilon(y))$$

to the composite

$$x \xrightarrow{f} \mathscr{T}_\varepsilon(y) \xrightarrow{\mathscr{T}_\varepsilon(g)} \mathscr{T}_\varepsilon(\mathscr{T}_\delta(z)) \to \mathscr{T}_{\varepsilon+\delta}(z),$$

where the unlabeled morphism is given by the lax monoidal structure of $\mathscr{T}$*. Identities are given by* $\mathscr{T}_0(x) \in \mathrm{Hom}_C(x, \mathscr{T}_0(x))$. □

5.2.2 Weak and strong interleavings in a category with a flow

We now unfold the definition of flow and consider two possible notions of interleaving in such a category: strong interleavings and weak interleavings. Our goal is to see that weak interleavings in a category with a flow $(C, \mathscr{T})$ correspond to interleavings in its associated locally persistent category $C_{\mathscr{T}}$. We remark that the description given here is the usual description of flow ([SMS18, Definition 2.3]). We also remark that, as will be apparent from the definitions, weak and strong interleavings are equivalent when the flow is a strong flow (Definition 5.2.9), which is often the case in practice.

Let C be a category. We defined a flow as being a lax monoidal functor $\mathscr{T} : \mathbf{R}_+ \to \mathrm{End}(C)$. This is equivalently given by:

- ▷ a functor $\mathscr{T} : \mathbf{R}_+ \to \mathrm{End}(C)$;
- ▷ a natural transformation $u : \mathsf{id}_C \Rightarrow \mathscr{T}_0$;
- ▷ for each $\varepsilon, \delta \in \mathbf{R}_+$, a natural transformation $\mu_{\varepsilon,\delta} : \mathscr{T}_\varepsilon \mathscr{T}_\delta \Rightarrow \mathscr{T}_{\varepsilon+\delta}$;

such that all of the diagrams of the following shapes commute in $\mathrm{End}(C)$:

1.
$$\begin{tikzcd} & \mathscr{T}_\varepsilon \arrow[dl, Rightarrow, "u\mathsf{id}_{\mathscr{T}_\varepsilon}"'] \arrow[dr, Rightarrow, "\mathsf{id}_{\mathscr{T}_\varepsilon}"] & \\ \mathscr{T}_0\mathscr{T}_\varepsilon \arrow[rr, Rightarrow, "\mu_{0,\varepsilon}"] & & \mathscr{T}_\varepsilon \end{tikzcd} \qquad \begin{tikzcd} & \mathscr{T}_\varepsilon \arrow[dl, Rightarrow, "\mathsf{id}_{\mathscr{T}_\varepsilon} u"'] \arrow[dr, Rightarrow, "\mathsf{id}_{\mathscr{T}_\varepsilon}"] & \\ \mathscr{T}_\varepsilon\mathscr{T}_0 \arrow[rr, Rightarrow, "\mu_{\varepsilon,0}"] & & \mathscr{T}_\varepsilon \end{tikzcd} \qquad (5.2.3)$$

2.

$$
\begin{tikzcd}
\mathcal{T}_\varepsilon\mathcal{T}_\delta\mathcal{T}_\gamma \arrow[r, Rightarrow, "\mathrm{id}_{\mathcal{T}_\varepsilon}\mu_{\delta,\gamma}"] \arrow[d, Rightarrow, "\mu_{\varepsilon,\delta}\mathrm{id}_{\mathcal{T}_\delta}"] & \mathcal{T}_\varepsilon\mathcal{T}_{\delta+\gamma} \arrow[d, Rightarrow, "\mu_{\varepsilon,\delta+\gamma}"] \\
\mathcal{T}_{\varepsilon+\delta}\mathcal{T}_\gamma \arrow[r, Rightarrow, "\mu_{\varepsilon+\delta,\gamma}"'] & \mathcal{T}_{\varepsilon+\delta+\gamma}
\end{tikzcd}
\qquad
\begin{tikzcd}
\mathcal{T}_\varepsilon\mathcal{T}_\delta \arrow[r, Rightarrow, "\mu_{\varepsilon,\delta}"] \arrow[d, Rightarrow, "\mathcal{T}_{\varepsilon\le\varepsilon'}\mathcal{T}_{\delta\le\delta'}"] & \mathcal{T}_{\varepsilon+\delta} \arrow[d, Rightarrow, "\mathcal{T}_{\varepsilon+\delta\le\varepsilon'+\delta'}"] \\
\mathcal{T}_{\varepsilon'}\mathcal{T}_{\delta'} \arrow[r, Rightarrow, "\mu_{\varepsilon',\delta'}"'] & \mathcal{T}_{\varepsilon'+\delta'}.
\end{tikzcd}
\tag{5.2.4}
$$

Definition 5.2.5 ([SMS18, Definition 2.6]). *Let $(C,\mathcal{T})$ be a category with a flow, $x, y \in C$, $\delta \in \mathbf{R}_+$, and $f : x \to \mathcal{T}_\delta(y)$ and $g : y \to \mathcal{T}_\delta(x)$. We say that f and g form a* ***weak δ-interleaving*** *between x and y if the following diagram commutes:*

$$
\begin{tikzcd}
\mathcal{T}_0 x \arrow[d] & x \arrow[l] \arrow[drr] & & y \arrow[r] \arrow[dll] & \mathcal{T}_0 y \arrow[d] \\
& \mathcal{T}_\delta x \arrow[drr] & & \mathcal{T}_\delta y \arrow[dll] & \\
\mathcal{T}_{2\delta} x & \mathcal{T}_\delta\mathcal{T}_\delta x \arrow[l] & & \mathcal{T}_\delta\mathcal{T}_\delta y \arrow[r] & \mathcal{T}_{2\delta} y
\end{tikzcd}
$$

where the diagonal morphisms are given by f and g and the functoriality of $\mathcal{T}$, and the remaining morphisms come from the lax monoidal structure of $\mathcal{T}$.

We say that f and g form a ***strong δ-interleaving*** *between x and y if the following diagram commutes:*

$$
\begin{tikzcd}
x \arrow[drr] \arrow[dd, bend right] & & y \arrow[dll] \arrow[dd, bend left] \\
\mathcal{T}_\delta x \arrow[drr] & & \mathcal{T}_\delta y \arrow[dll] \\
\mathcal{T}_\delta\mathcal{T}_\delta x & & \mathcal{T}_\delta\mathcal{T}_\delta y
\end{tikzcd}
$$

where the diagonal morphisms are given by f and g and the functoriality of $\mathcal{T}$, and the curved morphisms come from the lax monoidal structure of $\mathcal{T}$.

In [SMS18], the interleaving distance in a category with a flow is defined using weak interleavings. The following result says that the interleaving distance of a category with a flow coincides with the interleaving distance of its associated locally persistent category. The result follows at once from the description of $C_{\mathcal{T}}$ given in Proposition 5.2.2.

Proposition 5.2.6. *Let $(C,\mathscr{T})$ be a category with a flow, let $x,y \in \mathsf{obj}(C)$, and let $\delta \in \mathbf{R}_+$. Then x and y are weakly δ-interleaved in the category with a flow $(C,\mathscr{T})$ if and only if they are interleaved in the locally persistent category $C_{\mathscr{T}}$.* □

One could now ask: is there a notion of strong interleaving in a locally persistent category? The answer is that there is one, at least when the locally persistent category comes from a category with a flow. We now explain the precise relationship between **FlCAT** and **lpCAT**, and we characterize the image of the enriched Kleisli category construction **FlCAT** → **lpCAT**. The following lemma will help us describe the underlying category of the enriched Kleisli category of a category with a flow.

Lemma 5.2.7. *Let $(C,\mathscr{T})$ be a category with a flow. Then the flow structure on $\mathscr{T}$ induces a monad structure on the functor $\mathscr{T}_0 : C \to C$.*

Proof. Recall the description of flow given in Section 5.2.2. The unit of the monad $\mathscr{T}_0$ is given by the natural transformation $u : \mathsf{id} \Rightarrow \mathscr{T}_0$, and the multiplication is $\mu_{0,0} : \mathscr{T}_0\mathscr{T}_0 \Rightarrow \mathscr{T}_0$. The axioms of monad are then verified using the two triangles of Diagram 5.2.3 and the first square of Diagram 5.2.4, taking $\varepsilon = \delta = \gamma = 0$. □

5.2.3 Categories with a flow vs. locally persistent categories

We now explain some relationships between categories with a flow and locally persistent categories. As most of the proofs of the results in this section are tedious but simple, and largely a matter of careful bookkeeping, we will give fewer details than in other sections, emphasizing only the main points.

The enriched Kleisli category construction gives us the right diagonal functor in the following diagram of categories and functors:

$$\begin{tikzcd}
\mathbf{FlCAT}_{\mathsf{strong}} \arrow[r, hook] \arrow[d, "\sim"'] & \mathbf{FlCAT}_{\mathsf{s.strong}} \arrow[r, hook] \arrow[d, "\sim"] & \mathbf{FlCAT}_{\mathsf{idem}} \arrow[r, hook] \arrow[d] & \mathbf{FlCAT} \arrow[d, two heads, "\text{met. equiv.}"'] \arrow[dr, "\text{met. emb.}"] \arrow[ll, bend right, "\text{met. equiv.}"'] & \\
\mathbf{lpCAT}_{\mathsf{pow}} \arrow[r, hook] & \mathbf{lpCAT}_{\mathsf{w.pow}} \arrow[r, equal] & \mathbf{lpCAT}_{\mathsf{w.pow}} \arrow[r, equal] & \mathbf{lpCAT}_{\mathsf{w.pow}} \arrow[r, hook] & \mathbf{lpCAT}
\end{tikzcd} \tag{5.2.8}$$

We now describe the main constructions that appear in Diagram 5.2.8. For simplicity, we won't describe the category structure of all the categories of the diagram, but just the objects of each of the categories. Similarly, we will define functors only on objects.

The diagonal arrow in Diagram 5.2.8 is labeled as a being "metrically an embedding" because the interleaving distance in a category with a flow is equal to the interleaving distance in its corresponding locally persistent category, by Proposition 5.2.6.

We start with the top row. We can identify different kinds of flow, depending on their strictness.

Definition 5.2.9. *Let $(C, \mathcal{T})$ be a category with a flow. If $\mathcal{T}_0$ is an idempotent monad, we say that $\mathcal{T}$ is a* ***idempotent flow****; if u is a natural isomorphism, we say that $\mathcal{T}$ is an* ***semi-strong flow****; and if u and μ are natural isomorphisms, we say that $\mathcal{T}$ is a* ***strong flow****. We denote the collection of categories with an idempotent, semi-strong, or strong flow by $\mathbf{FlCAT}_{\mathsf{idem}}$, $\mathbf{FlCAT}_{\mathsf{s.strong}}$, or $\mathbf{FlCAT}_{\mathsf{strong}}$ respectively.*

This definition deserves a few remarks. Firstly, to the best of the author's knowledge, all the relevant examples of flow are at least idempotent. The reason to consider idempotent monads is that their Kleisli category is much better behaved than the Kleisli category of an arbitrary monad: it coincides with the Eilenberg–Moore category of the monad and it is complete whenever the original category is. As we have seen, completeness is very relevant when studying metric properties of the interleaving distance. Secondly, strong flows are also referred to as *essentially strict* flows in [Cru19]. We call them strong flows as they are precisely strong monoidal functors $\mathbf{R}_+ \to \mathrm{End}(C)$.

We have now defined all the categories and functors in the first row of Diagram 5.2.8, except for the curved arrow going from right to left. To describe this last functor, we note the following.

Lemma 5.2.10. *Let $(C, \mathcal{T})$ be a category with a flow. Then, the underlying category of the locally persistent category $C_{\mathcal{T}}$ is isomorphic to the Kleisli category of the monad $\mathcal{T}_0$.*

Proof. Both categories have as objects the objects of C. Now, by definition, for $x, y \in C$, a morphism in the Kleisli category of $\mathcal{T}_0$ from x to y is given by a morphism $x \to \mathcal{T}_0(y)$ in C. Moreover, composition in the Kleisli category works in the exact same way as in Proposition 5.2.2, taking $\varepsilon = \delta = 0$. Finally, identities in the Kleisli category are given by the components of the unit $u : \mathrm{id} \Rightarrow \mathcal{T}_0$, as in the underlying category of the locally persistent category $C_{\mathcal{T}}$. $\square$

The following result says that, in a sense, we can always assume that a flow is a semi-strong flow, as long as we are willing to modify the original category a bit.

Proposition 5.2.11. *Let $(C,\mathscr{T})$ be a category with a flow and let C' be the Kleisli category of the monad $\mathscr{T}_0$. Then, the flow $\mathscr{T}$ induces a canonical flow $\mathscr{T}'$ on C'. Moreover, for $x, y \in \mathsf{obj}(C)$, and $\delta > 0$, we have that x and y are weakly (resp. strongly) δ-interleaved in C if and only if they are weakly (resp. strongly) δ-interleaved in C'.*

Proof. If $(C,\mathscr{T})$ is a category with a flow, let $\mathscr{T}'_\varepsilon(x) = \mathscr{T}_\varepsilon(x)$ thought of as a functor $\mathbf{R}_+ \to \mathrm{End}(C')$, with C' the Kleisli category of $\mathscr{T}_0$. This provides us with a well-defined flow since, in C, the natural transformations $\mathsf{id} \Rightarrow \mathscr{T}_\varepsilon$ and $\mathscr{T}_\varepsilon\mathscr{T}_\delta \Rightarrow \mathscr{T}_{\varepsilon+\delta}$ factor as $\mathsf{id} \Rightarrow \mathscr{T}_0 \Rightarrow \mathscr{T}_\varepsilon\mathscr{T}_0 \Rightarrow \mathscr{T}_\varepsilon$ and $\mathscr{T}_\varepsilon\mathscr{T}_\delta \Rightarrow \mathscr{T}_\varepsilon\mathscr{T}_\delta\mathscr{T}_0 \Rightarrow \mathscr{T}_{\varepsilon+\delta}\mathscr{T}_0 \Rightarrow \mathscr{T}_{\varepsilon+\delta}$ respectively.

The condition on the interleavings follows by a routine check. □

Proposition 5.2.11 gives us the curved arrow in the top row of Diagram 5.2.8. We label this arrow as being “metrically an equivalence” because, regardless of us choosing the strong or weak interleaving distance, the construction lets us replace any category with a flow with a category with a semi-strong flow that is (canonically) metrically equivalent to the original one.

We now describe the bottom row of Diagram 5.2.8. We start by describing the categories appearing in the bottom row. To motivate the following definition, recall the definition of being powered by representables Definition 3.2.6, and note that, if $C_{\mathscr{T}}$ is the locally persistent category associated to a category with a flow $(C,\mathscr{T})$, the flow $\mathscr{T}$ seems to provide us with powers by representables, since

$$\mathrm{Hom}_{C_{\mathscr{T}}}(x,y)_\varepsilon = \mathrm{Hom}_C(x,\mathscr{T}_\varepsilon(y)) \cong \mathrm{Hom}_{C_{\mathscr{T}}}(x,\mathscr{T}_\varepsilon(y))_0.$$

One can check that the structure $\mathscr{T}$ provides us with is actually a bit weaker than being powered by representables. We now formalize this notion.

Definition 5.2.12. *Let $\mathscr{C}$ be a locally persistent category. A **weak power structure** on $\mathscr{C}$ consists of, for every $y \in \mathscr{C}$ and $\varepsilon \in \mathbf{R}_+$, an object $w(y,\varepsilon) \in \mathscr{C}$, and for every $x, y \in \mathscr{C}$ and $\varepsilon, \delta \in \mathscr{C}$ a function*

$$\mathsf{wp}_{\varepsilon,\delta} : \mathrm{Hom}_{\mathscr{C}}(x, w(\varepsilon,y))_\delta \to \mathrm{Hom}_{\mathscr{C}}(x,y)_{\varepsilon+\delta},$$

natural in x, y, ε and δ, and such that $\mathsf{wp}_{\varepsilon,0}$ is a bijection.

A weak power structure gives us powers by representables only when the functions $\mathsf{wp}_{\varepsilon,\delta}$ are all bijections.

The collection of locally persistent categories with a weak power structure is denoted by $\mathbf{lpCAT}_{\mathsf{w.pow}}$, and the collection of locally persistent categories that are powered by representables is denoted by $\mathbf{lpCAT}_{\mathsf{pow}}$. We have now described all the constructions involved in the second row, and go on to explain the functors between the first row and the second row of Diagram 5.2.8. All these functors are restrictions of the diagonal functor on the right of the diagram.

If wp is a weak power structure on a locally persistent category $\mathscr{C}$, we get a functor $F^{\mathsf{wp}}_{-}(-) : \mathbf{R}_+ \to \mathrm{End}(\mathscr{C}_0)$ given by $F^{\mathsf{wp}}_{\varepsilon}(y) = w(\varepsilon, y)$, an isomorphism $u^{\mathsf{wp}} : x \to F^{\mathsf{wp}}_0(x)$ for every $x \in \mathscr{C}$, and a morphism $\mu^{\mathsf{wp}} : F^{\mathsf{wp}}_{\varepsilon}\left(F^{\mathsf{wp}}_{\delta}(x)\right) \to F^{\mathsf{wp}}_{\varepsilon+\delta}(x)$. One can easily check that this structure satisfies the definition of flow.

Lemma 5.2.13. *Let $\mathscr{C}$ be a locally persistent category and let* wp *be a weak power structure on $\mathscr{C}$. The functor F^{wp} together with the natural transformations u^{wp} and μ^{wp} constitute a semi-strong flow on $\mathscr{C}_0$.* □

Remark 5.2.14. Lemma 5.2.13 lets us import the definitions of strong interleaving and strong iterleaving distance into locally persistent categories endowed with a weak power structure.

Conversely, we have the following, which is again a routine check.

Lemma 5.2.15. *If $(C, \mathscr{T})$ is a category with a flow, the flow $\mathscr{T}'$ on the underlying category of $C_{\mathscr{T}}$, which is equivalently the Kleisli category of $\mathscr{T}_0$, provides us with a weak power structure. A weak power of $x \in C$ by $\varepsilon \in \mathbf{R}_+$ in $C_{\mathscr{T}}$ is given by $\mathscr{T}_{\varepsilon}(x)$.* □

Lemma 5.2.13 together with Lemma 5.2.15 give the second vertical equivalence in Diagram 5.2.8. Moreover, Lemma 5.2.15 gives the fourth vertical functor in Diagram 5.2.8 which we label as surjective and as being “metrically an equivalence” since it preserves the weak and strong interleaving distances, and since every weakly powered locally persistent category comes from a category with a flow, by Lemma 5.2.13.

Proposition 5.2.16. *Let $(C, \mathscr{T})$ be a category with a flow. If $\mathscr{T}$ is strong, then the locally persistent category $C_{\mathscr{T}}$ is powered by representables. The power of $y \in C$ by $\varepsilon \in \mathbf{R}_+$ is given by $\mathscr{T}_{\varepsilon}(y)$.*

Proof. Let $x \in C$. Then, in $C_{\mathscr{T}}$, we have

$$\mathrm{Hom}_{C_{\mathscr{T}}}(x, y)_{\varepsilon+-} \cong \mathrm{Hom}_C(x, \mathscr{T}_{\varepsilon+-}(y)) \cong \mathrm{Hom}_C(x, \mathscr{T}_{-}\mathscr{T}_{\varepsilon}(y)) \cong \mathrm{Hom}_{C_{\mathscr{T}}}(x, \mathscr{T}_{\varepsilon}(y))$$

as functors $\mathbf{R}_+ \to \mathbf{Set}$. We used the fact that $\mathcal{T}$ is strong in the second isomorphism. Since all of the isomorphisms are natural, this concludes the proof. □

A similar argument proves the following.

Proposition 5.2.17. *Let $\mathscr{C}$ be a locally persistent category that is powered by representables. Then the powering gives rise to a strong flow on its underlying category.* □

Proposition 5.2.16 together with Proposition 5.2.17 give the first vertical equivalence in the diagram, which we label as an equivalence since the data of a locally persistent category that is powered by representables is equivalent to the data of a category with a strong flow, by Proposition 5.2.17 and Proposition 5.2.16. This concludes the description of Diagram 5.2.8.

5.2.4 Metric properties of categories with a flow

We now use our general metric results proven in Chapter 4 to prove metric results about the interleaving distance of a category with a flow.

Proposition 5.2.18. *Let $(C, \mathcal{T})$ be a category with a strong flow.*

1. *If C admits sequential limits and $\mathcal{T}_\varepsilon$ preserves sequential limits for every $\varepsilon \in \mathbf{R}_+$, then the locally persistent category $C_{\mathcal{T}}$ admits weighted sequential limits.*

2. *If C admits binary products and pullbacks and $\mathcal{T}_\varepsilon$ preserves these limits for every $\varepsilon \in \mathbf{R}_+$, then the locally persistent category $C_{\mathcal{T}}$ admits terminal midpoints.*

Proof. In both cases, $C_{\mathcal{T}}$ is powered by representables, by Proposition 5.2.16. Moreover, by assumption, these powers preserve the necessary limits to apply Proposition 3.2.12, for the first claim, and Proposition 3.2.19, for the second, concluding the proof. □

Combining Proposition 5.2.18 with Theorem 4.3.1 and Corollary 4.4.3 we conclude the following.

Theorem 5.2.19. *Let $(C, \mathcal{T})$ be a category with a flow such that $\mathcal{T}$ is strong, C is complete, and $\mathcal{T}_\varepsilon$ preserves limits for all $\varepsilon \in \mathbf{R}_+$. Then the interleaving distance of C is intrinsic and complete.* □

To the best of the author's knowledge, the result about the interleaving distance of categories with a flow being intrinsic is new. On the other hand, a completeness result for categories with a flow has already been proven in [Cru19].

Theorem 5.2.20 ([Cru19, Theorem 4])**.** *Let $(C, \mathscr{T})$ be a category with a flow such that C admits sequential limits and $\mathscr{T}_\varepsilon$ preserves sequential limits for every $\varepsilon \in \mathbf{R}_+$. Then the interleaving distance of C is complete.* □

Although our general completeness theorem (Theorem 4.3.1) is neither stronger nor weaker than the completeness result of [Cru19], since they apply to different objects, the conditions for completeness in Theorem 5.2.19 are a bit stronger than the conditions in [Cru19]. The difference is that we assume that the flow $\mathscr{T}$ is strong, an assumption not present in Theorem 5.2.20.

We finish this section by explaining how to use [Cru19, Theorem 4] to strengthen Theorem 5.2.19.

Proposition 5.2.21. *Let $\mathscr{C}$ be a locally persistent category. Assume that $\mathscr{C}$ admits a weak power structure. If the underlying category of $\mathscr{C}$ admits sequential limits and the weak powers preserve sequential limits, then the interleaving distance of $\mathscr{C}$ is complete.*

Proof. The weak power structure provides us with a flow for $\mathscr{C}_0$ by Lemma 5.2.13. Theorem 5.2.20 then tells us that the interleaving distance of the category with a flow $\mathscr{C}_0$ is complete, which proves the claim, since the interleaving distance given by the flow coincides with the interleaving distance of $\mathscr{C}$, by Proposition 5.2.6. □

5.2.5 Q-flows

In Section 5.1, we defined locally Q-persistent categories, which are categories enriched in $\mathbf{Set}^Q$, for Q a monoidal poset. As argued there, it is often the case that the most natural $\mathbf{Set}^Q$-enrichment of a category is given by a monoidal poset Q different from $\mathbf{R}_+$. The same thing happens with flows.

Definition 5.2.22. *Let Q be a monoidal poset and let C be a category. A **Q-flow** consists of a lax monoidal functor $\mathscr{T} : Q \to \mathrm{End}(C)$.*

By the enriched Kleisli construction (Section 5.2.1), a Q-flow $\mathscr{T}$ on a category C induces a locally Q-persistent structure on C, which we denote by $C_{\mathscr{T}}$.

A simple and interesting example is the following.

Example 5.2.23. Assume given two flows $\mathscr{T}, \mathscr{T}' : \mathbf{R}_+ \to \mathrm{End}(C)$ on a category C. We would like to use both to define an interleaving distance that takes both possible shifts into account. We can do this as long as the flows $\mathscr{T}$ and $\mathscr{T}'$ commute in a coherent way. Concretely, we say that $\mathscr{T}$ and $\mathscr{T}'$ **coherently commute** if there exists a lax monoidal functor $(\mathscr{T}, \mathscr{T}') : \mathbf{R}_+ \times \mathbf{R}_+ \to \mathrm{End}(C)$ and a diagram of monoidal functors and monoidal categories

$$\begin{tikzcd}
\mathbf{R}_+ \arrow[d, "\mathsf{id} \times 0"'] \arrow[dr, "\mathscr{T}"] & \\
\mathbf{R}_+ \times \mathbf{R}_+ \arrow[r, "(\mathscr{T}, \mathscr{T}')"] & \mathrm{End}(C) \\
\mathbf{R}_+ \arrow[u, "0 \times \mathsf{id}"] \arrow[ur, "\mathscr{T}'"'] &
\end{tikzcd}$$

that commutes up to natural isomorphism. We can then use this flow to endow C with a locally bi-persistent structure.

Chapter 6

Applications

6.1 The homotopy interleaving distance of a locally persistent model category

In this section, we study the quotient interleaving distance of a relative locally persistent category whose acyclic morphisms are the weak equivalences of a model structure. Concretely, we are given a model category $\mathscr{M}$ such that its underlying category has the structure of a locally persistent category and we let the class of acyclic morphism of $\mathscr{M}$ be the class of weak equivalences. We denote the quotient interleaving distance obtained using Definition 3.3.6 by $(d_I^{\mathscr{M}})_{/\mathscr{W}}$ or $(d_I^{\mathscr{M}})_{/\simeq}$. We prove that, under mild hypotheses, this distance is intrinsic and complete, and we provide a characterization of the distance as an infimum over interleavings (Theorem 6.1.6).

The main motivation for studying these distances comes from the homotopy interleaving distance, defined by Blumberg and Lesnick in [BL17]. The homotopy interleaving distance is a distance on the category $\mathbf{Top}^{\mathbf{R}}$ of persistent topological spaces and is used to lift the continuity of the Vietoris–Rips filtration, and other invariants of metric spaces, to the homotopy level. More specifically, in the case of the Vietoris–Rips filtration this means that the stability of the persistent homology of the Vietoris–Rips filtration can be proven by first showing that the functor $\mathsf{VR} : \mathbf{Met} \to \mathbf{Top}^{\mathbf{R}}$ is Lipschitz and then using the algebraic stability of barcodes ([CCSGGO09, Theorem 4.4]). This approach is explained in [BL17, Sections 1.2 and 3.2]. We address the stability of Vietoris–Rips and other related filtrations in Section 6.3.

The lift to the homotopy level is useful since it shows that any stable invariant of persistent topological spaces that is homotopy invariant can be used to produce a stable invariant of metric spaces. Another important example of such an invariant is the path components functor π_0, that maps persistent topological spaces to persistent sets.

The *homotopy interleaving distance* between two persistent topological spaces $X, Y \in \mathbf{Top}^{\mathbf{R}}$ is defined as

$$d_{HI}(X,Y) = \inf\big\{\delta \in \mathbf{R}_+ : \exists X' \simeq X \text{ and } Y' \simeq Y \text{ such that } X' \text{ and } Y' \text{ are } \delta\text{-interleaved}\big\}.$$

Note that, with this definition, it is not immediately clear that the homotopy interleaving distance satisfies the triangle inequality ([BL17, Section 4]). We use some of the techniques in [BL17] to show that such a definition yields a metric in very general situations (Theorem 6.1.7), including the case of persistent topological spaces, persistent simplicial sets, and persistent chain complexes (Example 6.1.8). Moreover, this metric coincides with our quotient interleaving distance for locally persistent model categories, so in particular, $d_{HI} = \left(d_I^{\mathbf{Top}^{\mathbf{R}}}\right)_{/\simeq}$. This also lets us apply our metric results for interleaving distances to prove that such distances are intrinsic and complete, so in particular, we prove that the homotopy interleaving distance is intrinsic and complete. Finally, as pointed out in Example 6.1.8, our approach also makes it clear that the spaces of persistent topological spaces and of persistent simplicial sets are equivalent as metric spaces.

In order to apply our results, we first restrict our attention to the case where the acyclic morphisms of $\mathscr{M}$ are taken to be the trivial fibrations. We denote the class of trivial fibrations by $\mathbf{tFib} = \mathbf{Fib} \cap \mathscr{W}$, and consider the quotient of the interleaving distance by $\mathbf{tFib}$, which we denote by $(d_I^{\mathscr{M}})_{/\mathbf{tFib}}$.

We start with a lemma about the stability of trivial fibrations.

Lemma 6.1.1. *Let $\mathscr{M}$ be a locally persistent category with a model structure on its underlying category such that $\mathscr{M}$ is copowered and powered by representables, and such that powers preserve trivial fibrations. Then, trivial fibrations are stable under weighted pullbacks.*

Proof. The category $\mathscr{M}$ is complete, copowered, and powered. We use Lemma 4.1.5. If f is a trivial fibration, then f^{ε} is again a trivial fibration by hypothesis, so the result follows from the fact that, in any model category, trivial fibrations are stable under pullbacks. □

We can now give a more concrete description of $(d_I^{\mathscr{M}})_{/\mathbf{tFib}}$.

Proposition 6.1.2. *Let $\mathscr{M}$ be a locally persistent category with a model structure on its underlying category such that $\mathscr{M}$ is copowered and powered by representables, and such that powers preserve trivial fibrations. Then,*

$$(d_I^{\mathscr{M}})_{/\mathbf{tFib}}(x,y) = \inf\left\{\delta : \exists x' \simeq_{\mathbf{tFib}} x, y' \simeq_{\mathbf{tFib}} y,\ x' \text{ and } y' \text{ are } \delta\text{-interleaved}\right\}.$$

Proof. Lemma 6.1.1 tells us that trivial fibrations are stable under weighted pullbacks. Thus we can apply Theorem 4.1.4, which says that the quotient interleaving distance can be computed as an infimum over interleavings, as in the statement. □

Our next goal is to establish that, under the hypotheses of Proposition 6.1.2, the quotient interleaving distance $(d_I^{\mathscr{M}})_{/\mathbf{tFib}}$ is intrinsic and complete.

Proposition 6.1.3. *Let $\mathscr{M}$ be a locally persistent category with a model structure on its underlying category such that $\mathscr{M}$ is copowered and powered by representables, and such that powers preserve trivial fibrations. Then, $(d_I^{\mathscr{M}})_{/\mathbf{tFib}}$ is intrinsic and complete.*

Proof. This is a simple application of Corollary 4.4.5 and Theorem 4.3.3. The hypotheses of Corollary 4.4.5 are satisfied since $\mathscr{M}$ is complete, copowered and powered by representables, so $\mathscr{M}$ has terminal midpoints by Proposition 3.2.19.

Similarly, the completeness hypotheses of Theorem 4.3.3 are satisfied by the completeness of $\mathscr{M}$ and it being copowered and powered by representables, using Proposition 3.2.15. For the other hypothesis, note that, by Lemma 6.1.1, trivial fibrations are stable under weighted pullbacks and that trivial fibrations are always closed under sequential limits. □

We now use the above propositions to prove analogous theorems about the distance we are really interested in, namely, the interleaving distance of $\mathscr{M}$ quotiented by the equivalence relation given by weak equivalence. We start with a useful lemma.

Lemma 6.1.4. *Let $\mathscr{M}$ be a model category. Any two fibrant and weakly equivalent objects are connected by a zig-zag of trivial fibrations.*

In the following proof we use standard facts about model categories that can be found in, e.g., [Hov07, Section 1.2].

Proof. Let $x \simeq y \in \mathscr{M}$ be fibrant. Let $\alpha' : c' \to x$ be a cofibrant replacement. Since y is fibrant, and $c' \simeq y$, there is a weak equivalence $\beta' : c' \to y$. Consider the morphism $\alpha' \times \beta' : c' \to x \times y$ and factor it as a trivial cofibration $g : c' \to c$ followed by a fibration $f : c \to x \times y$. Since x and y are fibrant, the projections $\pi_x : x \times y \to x$ and $\pi_y : x \times y \to y$ are fibrations. Composing f with the projections to x and y, we get fibrations $\alpha = \pi_x \circ f : c \to x$ and $\beta = \pi_y \circ f : c \to y$. To see that these are trivial fibrations, recall that $\alpha' : c' \to x$ and $\beta' : c' \to y$ are weak equivalence and use the 2-out-of-3 property, noting that $\alpha' = \alpha \circ g$ and $\beta' = \beta \circ g$. □

It follows that, for fibrant objects, the equivalence relation induced by weak equivalences is the same as the equivalence relation induced by trivial fibrations. More specifically, two fibrant objects are connected by a zig-zag of weak equivalences if and only if they are connected by a zig-zag of trivial fibrations.

Lemma 6.1.5. *Let $\mathscr{M}$ be a locally persistent category with a model structure on its underlying category and assume that $\mathscr{M}$ admits a locally persistent fibrant replacement functor $F : \mathscr{M} \to \mathscr{M}$. Then, the functor $F : \mathscr{M} \to \mathscr{M}$ together with the identity functor $\mathsf{id}_{\mathscr{M}} : \mathscr{M} \to \mathscr{M}$ exhibit the ep metric space $(\mathscr{M}, (d_I^{\mathscr{M}})_{/\mathscr{W}})$ as a pseudo retract of $(\mathscr{M}, (d_I^{\mathscr{M}})_{/\mathbf{tFib}})$.*

See Definition 2.2.16 for the notion of pseudo retract of ep metric spaces.

Proof. Both the identity and the fibrant replacement functor are locally persistent. Moreover, if $x \simeq_{\mathscr{W}} y$ in $\mathscr{M}$, then, $Fx \simeq_{\mathbf{tFib}} Fy$ in $\mathscr{M}$ by Lemma 6.1.4, and $x \simeq_{\mathbf{tFib}} y$ in $\mathscr{M}$ clearly implies $x \simeq_{\mathscr{W}} y$ in $\mathscr{M}$. So $\mathsf{id}_{\mathscr{M}}$ and F give well-defined distance non-increasing maps $F : (\mathscr{M}, (d_I^{\mathscr{M}})_{/\mathscr{W}}) \to (\mathscr{M}, (d_I^{\mathscr{M}})_{/\mathbf{tFib}})$ and $\mathsf{id} : (\mathscr{M}, (d_I^{\mathscr{M}})_{/\mathbf{tFib}}) \to (\mathscr{M}, (d_I^{\mathscr{M}})_{/\mathscr{W}})$.

Finally, we must show that for every $x \in \mathscr{M}$, we have $(d_I^{\mathscr{M}})_{/\mathscr{W}}(Fx, x) = 0$. This is clear, since Fx and x are weakly equivalent. □

We are ready to prove the main theorem of this section.

Theorem 6.1.6. *Let $\mathscr{M}$ be a locally persistent category with a model structure on its underlying category such that $\mathscr{M}$ is copowered and powered by representables, and such that powers preserve trivial fibrations. Assume that $\mathscr{M}$ admits a locally persistent fibrant replacement functor $\mathscr{M} \to \mathscr{M}$. Then, $(d_I^{\mathscr{M}})_{/\mathscr{W}}$ is intrinsic and complete and it satisfies*

$$(d_I^{\mathscr{M}})_{/\mathscr{W}}(x, y) = \inf\left\{\delta : \exists x' \simeq x, \exists y' \simeq y,\ x' \text{ and } y' \text{ are } \delta\text{-interleaved}\right\}.$$

Proof. To prove that the distance is intrinsic and complete, it is enough to show that $(d_I^{\mathcal{M}})_{/\mathbf{tFib}}$ is intrinsic and complete, by Lemma 6.1.5 and Lemma 2.2.17. And to show this, we just use Proposition 6.1.3.

Now, for the formula, note that

$$\begin{aligned}(d_I^{\mathcal{M}})_{/\mathcal{W}}(x,y) &= (d_I^{\mathcal{M}})_{/\mathbf{tFib}}(Fx,Fy)\\ &= \inf\{\delta : \exists x' \simeq_{\mathbf{tFib}} Fx, y' \simeq_{\mathbf{tFib}} Fy,\ x' \text{ and } y' \text{ are } \delta\text{-interleaved}\}\\ &= \inf\{\delta : \exists x' \simeq_{\mathcal{W}} x, y' \simeq_{\mathcal{W}} y,\ x' \text{ and } y' \text{ are } \delta\text{-interleaved}\},\end{aligned}$$

where in the first equality we used the pseudo retraction, the second equality is by Proposition 6.1.2, and the final equality follows from the fact that F is locally persistent functor that is homotopy invariant. □

The main application we have in mind is the following. Let $\mathcal{M}$ be a cofibrantly generated model category and consider the functor category $\mathcal{M}^{\mathbf{R}}$ with its projective model structure (Definition 2.5.7). Recall that the projective model structure is characterized by the fact that the weak equivalences and fibrations are defined pointwise.

The category $\mathcal{M}^{\mathbf{R}}$ is locally persistent and copowers and powers by representables are given by shifting to the right and to the left respectively (Example 3.2.7). In particular, these shifts preserve trivial fibrations. Finally, any functorial fibrant replacement $\mathcal{M} \to \mathcal{M}$ provides us with a locally persistent fibrant replacement $\mathcal{M}^{\mathbf{R}} \to \mathcal{M}^{\mathbf{R}}$, so applying Theorem 6.1.6, we deduce the following.

Theorem 6.1.7. *Let $\mathcal{M}$ be a cofibrantly generated model category and let the class of acyclic morphisms of the locally persistent category $\mathcal{M}^{\mathbf{R}}$ be the class of natural transformations that are componentwise weak equivalences. Then, the quotient interleaving distance $(d_I^{\mathcal{M}^{\mathbf{R}}})_{/\simeq}$ is intrinsic and complete, and it satisfies*

$$\left(d_I^{\mathcal{M}^{\mathbf{R}}}\right)_{/\simeq} = \inf\{\delta : \exists x' \simeq x, y' \simeq y,\ x' \text{ and } y' \text{ are } \delta\text{-interleaved}\}. \qquad \square$$

We note that the same theorem holds for categories of the form $\mathcal{M}^{\mathbf{R}^n}$ for $n > 1$.

We conclude the section with a few applications.

Example 6.1.8. Recall from Example 2.5.9, Example 2.5.10, and Example 2.5.11 that there are cofibrantly generated model structures on the categories of chain complexes (over some commutative ring R), simplicial sets, and topological spaces. These give us relative locally persistent categories $\mathbf{Ch}(R)^{\mathbf{R}}$, $\mathbf{sSet}^{\mathbf{R}}$, and $\mathbf{Top}^{\mathbf{R}}$ that satisfy the hypothesis of Theorem 6.1.7. In particular, we get quotient interleaving distances for all of these categories. These distances are intrinsic and complete, and have the form of the homotopy interleaving distance of [BL17].

These locally persistent categories are closely related. For example, one can use the geometric realization and singular complex functors $|-| : \mathbf{sSet} \rightleftarrows \mathbf{Top} : S$ to obtain analogous locally persistent functors $|-| : \mathbf{sSet}^{\mathbf{R}} \rightleftarrows \mathbf{Top}^{\mathbf{R}} : S$ between the corresponding locally persistent categories. Since these functors form a Quillen equivalence, it follows that the ep metric spaces $\left(\mathbf{sSet}^{\mathbf{R}}, \left(d_I^{\mathbf{sSet}^{\mathbf{R}}}\right)_{/\simeq}\right)$ and $\left(\mathbf{Top}^{\mathbf{R}}, \left(d_I^{\mathbf{Top}^{\mathbf{R}}}\right)_{/\simeq}\right)$ are equivalent, in the sense that, after identifying points at distance 0, they become isometric.

Another interesting homotopy invariant locally persistent functor is the chain complex functor $Ch_R : \mathbf{sSet}^{\mathbf{R}} \to \mathbf{Ch}(R)^{\mathbf{R}}$. It follows from Theorem 4.2.2 that this functor is distance non-increasing.

6.2 The Gromov–Hausdorff distance on dissimilarity spaces

In this section, we prove that the Gromov–Hausdorff distance can be interpreted as a quotient interleaving distance and use our results about metric properties of quotient interleaving distances to recover some well-known facts about the Gromov–Hausdorff distance.

In Section 6.2.1, we study the Gromov–Hausdorff distance on the collection of dissimilarity spaces, a very weak version of metric spaces: a dissimilarity space consists of a set X together with a function $X \times X \to [0,\infty]$. Of course, every ep metric space gives rise to a dissimilarity space, so, in Section 6.2.2, we restrict our attention to the Gromov–Hausdorff distance on ep metric spaces and on compact ep metric spaces. Here, we use our results on quotient interleaving distances being complete and geodesic to recover the well-known facts that, when restricted to compact metric spaces, the Gromov–Hausdorff distance is complete and geodesic, and that two compact metric spaces are at distance zero if and only if they are isometric.

In the case of dissimilarity spaces, the hypotheses of our results are verified by abstract considerations, since the locally persistent category of dissimilarity spaces is very well behaved. When proving that the Gromov–Hausdorff distance is geodesic when restricted to compact ep metric spaces, we verify the hypotheses of our results by using well-known constructions that feature in the standard proofs of the fact that the Gromov–Hausdorff distance is geodesic ([INT16], [CM18b]). In this sense, the proof that we give is not new, only the point of view is.

6.2.1 Gromov–Hausdorff distance on dissimilarity spaces

In [Seg16] the notion of *network* is considered. A network consists of a finite set X together with a dissimilarity function $d : X \times X \to \mathbf{R}_+$ such that $d(x,x) = 0$ for all $x \in X$. Networks can be seen as generalized finite metric spaces: they need not satisfy the triangle inequality or symmetry, only reflexivity. The thesis [Seg16] is concerned with algorithmic transformations of networks, and in particular, with clustering algorithms and their stability. In order to formulate stability, the author generalizes the Gromov–Hausdorff distance to networks. This is our starting point. We consider a slightly more general notion of network, similar to the generalization considered in [CM18a], which we call dissimilarity space. Let $\overline{\mathbf{R}}_+ = [0,\infty]$.

Definition 6.2.1. *A* ***dissimilarity space*** *consists of a set X together with a function $d_X : X \times X \to \overline{\mathbf{R}}_+$.*

Example 6.2.2. Of course, any ep metric space is in particular a dissimilarity space. Another important family of dissimilarity spaces is given by merge functions. A **merge function** on a set X is given by a function $m : X \times X \to \overline{\mathbf{R}}_+$ such that, for every $x, y, z \in X$, we have

$$m(x,z) \geq \min(m(x,y), m(y,z)).$$

In Section 6.6.3, we explain how merge functions can be used to encode one-parameter hierarchical clusterings.

The definition of the Gromov–Hausdorff distance using correspondences generalizes to dissimilarity spaces.

Definition 6.2.3. *The* ***Gromov–Hausdorff distance*** *between dissimilarity spaces X and Y is given by*

$$d_{GH}(X,Y) = \frac{1}{2} \inf\left\{\operatorname{dist}(R) : R \subseteq X \times Y \textit{ correspondence}\right\},$$

where, as usual, $\operatorname{dist}(R) = \sup\left\{|d_X(x,x') - d_Y(y,y')| : (x,y),(x',y') \in R\right\}$.

We give a few remarks about this definition. Firstly, the fact that the above distance is in fact an ep metric is proven by composing correspondences, as usual. Secondly, note that, for X and Y dissimilarity spaces and R a correspondence between them, $d_X(x,x)$ can be strictly greater than 0, so the case $|d_X(x,x) - d_Y(y,y)|$ for $(x,y) \in R$ is relevant. Finally, this definition specializes to [Seg16, Equation 5.82] when the dissimilarity spaces happen to be networks and, by Theorem 2.2.27, the definition also specializes to the Gromov–Hausdorff distance on metric spaces.

The goal of this section is to show that the ep metric of Definition 6.2.3 is a quotient interleaving distance, and that our theorems imply that this distance is intrinsic and complete. We start by defining the locally persistent category of dissimilarity spaces.

A morphism of dissimilarity spaces $f : X \to Y$ consists of a function $f : X \to Y$ between the underlying sets, such that $d_X(x,x') \geq d_Y(f(x), f(x'))$ for all $x, x' \in X$. Together with composition of functions and identity functions, this endows the collection of dissimilarity spaces with the structure of a category. We temporarily let $\operatorname{Hom}(X,Y)$ denote the set of morphisms between two dissimilarity spaces X and Y.

Given a dissimilarity space X and $\varepsilon \in \overline{\mathbf{R}}_+$, consider $\varepsilon \cdot X$, the dissimilarity space with the same underlying set as X and distance given by $d_{\varepsilon\cdot X}(x,y) = d_X(x,y) + \varepsilon$. Similarly, given $\varepsilon \in \mathbf{R}_+$, let X^ε be the dissimilarity space with the same underlying set as X and distance given by $d_{X^\varepsilon}(x, y) = \max(0, d_X(x, y) - \varepsilon)$.

Definition 6.2.4. *Let* **Diss** *denote the locally persistent category whose objects are dissimilarity spaces, and such that*

$$\mathrm{Hom}_{\mathbf{Diss}}(X, Y)_\varepsilon = \mathrm{Hom}(\varepsilon \cdot X, Y) = \mathrm{Hom}(X, Y^\varepsilon).$$

Identities and composition are defined in the obvious way. With this locally persistent structure, interleavings have a particularly simple description.

Lemma 6.2.5. *A δ-interleaving between dissimilarity spaces X and Y is given by functions of sets $f : X \to Y$ and $g : Y \to X$ such that f and g are inverse bijections, and such that $|d_X(x, x') - d_Y(f(x), f(x'))| \leq \delta$ for all $x, x' \in P$.* □

Also by construction, we deduce the following.

Lemma 6.2.6. *The locally persistent category* **Diss** *is copowered and powered by representables. Copowers are given by $\varepsilon \cdot X$ and powers are given by X^ε, for $X \in$* **Diss** *and $\varepsilon \in \mathbf{R}_+$.* □

The fact that the locally persistent category **Diss** is copowered and powered by representables is one of the main reasons to work in this category, rather than working directly with metric spaces. The other reason is that the underlying category of **Diss** is complete, and limits have a very concrete description.

Lemma 6.2.7. *The underlying category of* **Diss** *is complete.*

Proof. It is enough to show that **Diss** has arbitrary products and pullbacks. Let $\{X_i\}_{i \in I}$ be a family of dissimilarity spaces. Let

$$X = \prod_{i \in I} X_i \quad \text{and} \quad d_X(\{x_i\}, \{x'_i\}) = \sup_{i \in I} d_{X_i}(x_i, x'_i).$$

It is clear that (X, d_X), together with the natural projections $X \to X_i$, satisfies the universal property of the product.

For pullbacks, let $X \xrightarrow{f} Z \xleftarrow{g} Y$ be a cospan of dissimilarity spaces, and let

$$P = \{(x, y) \in X \times Y : f(x) = g(y)\} \quad \text{and} \quad d_P((x, y), (x', y')) = \max(d_X(x, x'), d_Y(y, y')).$$

It is again straightforward to see that (P, d_P), together with the natural projections $P \to X$ and $P \to Y$, satisfies the universal property of the pullback. □

As usual, combining powers, copowers, and limits of the underlying category, we obtain all the limits that are relevant for the study of the interleaving distance of a locally persistent category.

Lemma 6.2.8. *The locally persistent category* **Diss** *admits weighted pullbacks, weighted sequential limits, and terminal midpoints.*

Proof. Since **Diss** is copowered and powered by representables and is furthermore complete, the claims follow from Proposition 3.2.12, Proposition 3.2.15, and Proposition 3.2.19. □

We now define the acyclic morphisms to endow **Diss** with a relative locally persistent category structure.

Definition 6.2.9. *An **acyclic morphism** between dissimilarity spaces is a surjective and distance preserving morphism.*

We are interested in the quotient interleaving distance $\left(d_I^{\mathbf{Diss}}\right)_{/\simeq}$. We have the following.

Proposition 6.2.10. *Acyclic morphisms between dissimilarity spaces are stable under weighted pullback.*

Proof. The locally persistent category **Diss** is complete and copowered and powered by representables. We use Lemma 4.1.5, so it is enough to show that powers preserve acyclic morphisms and that acyclic morphisms are stable under pullbacks of the underlying category.

By the description of powers (Lemma 6.2.6), it is clear that the power of a surjective map $f : X \to Y$ is surjective. Since $d_X(x,x') = d_Y(f(x), f(x'))$ implies $\max(d_X(x,x') - \varepsilon, 0) = \max(d_Y(f(x), f(x')) - \varepsilon, 0)$, we deduce that powers preserve acyclic morphisms.

By the description of pullbacks in the proof of Lemma 6.2.7, it is clear that a pullback of a surjective map is surjective. To see that a pullback of a distance preserving map is distance preserving, let

$$P = \{(x,y) \in X \times Y : f(x) = g(y)\} \quad \text{and} \quad d_P((x,y),(x',y')) = \max(d_X(x,x'), d_Y(y,y'))$$

be the pullback of a cospan $X \xrightarrow{f} Z \xleftarrow{g} Y$ and assume that g is distance preserving.

Given $(x, y), (x', y') \in P$, we have

$$\begin{aligned} d_P((x,y),(x',y')) &= \max(d_X(x,x'), d_Y(y,y')) \\ &= \max(d_X(x,x'), d_Z(g(y),g(y'))) \\ &= \max(d_X(x,x'), d_Z(f(x),f(x'))), \end{aligned}$$

where in the second equality we used the fact that g is distance preserving and in the third one we used the fact that $(x, y), (x', y') \in P$. Since f is distance non-increasing, we have $d_P((x, y), (x', y')) = d_X(x, x')$, as required. □

We can then deduce the following.

Theorem 6.2.11. *The quotient interleaving distance on dissimilarity spaces is intrinsic and complete and satisfies*

$$\begin{aligned} \left(d_I^{\mathbf{Diss}}\right)_{/\simeq}(X,Y) &= \inf\left\{\delta : \exists X' \simeq X, Y' \simeq Y,\ X' \textit{ and } Y' \textit{ are } \delta\textit{-interleaved}\right\} \\ &= \inf\{\delta : \exists \textit{ acyclic morphisms } X' \to X \textit{ and } Y' \to Y \\ &\qquad\qquad \textit{such that } X' \textit{ and } Y' \textit{ are } \delta\textit{-interleaved}\}. \end{aligned}$$

Proof. The facts that the distance is intrinsic and complete follow from Corollary 4.4.5 and Theorem 4.3.3, using Lemma 6.2.8 and Proposition 6.2.10, and noting that acyclic morphisms are clearly closed under sequential limits.

The description of the quotient interleaving distance follows from Theorem 4.1.4. □

We conclude this section by relating the quotient interleaving distance of **Diss** to the Gromov–Hausdorff distance.

Theorem 6.2.12. *For X and Y dissimilarity spaces we have*

$$\left(d_I^{\mathbf{Diss}}\right)_{/\simeq}(X,Y) = 2d_{GH}(X,Y).$$

Proof. We use the second characterization of the quotient interleaving distance of Theorem 6.2.11.

Assume that $2d_{GH}(X, Y) < \delta$. Then, there is a correspondence $R \subseteq X \times Y$ such that $\mathrm{dist}(R) < \delta$. Consider the ep metric space R^X, with underlying set given by R and metric given by $d_{R^X}((x, y), (x', y')) = d_X(x, x')$. With this definition, the projection

$\pi_X : R^X \to X$ is an acyclic morphism of dissimilarity spaces. Construct, analogously, the metric space R^Y.

Now, consider the bijection $R^X \to R^Y$ given by the identity. Let us see that this function, and its inverse, do not increase the distance more than δ.

$$|d_{R^X}((x,y),(x',y')) - d_{R^Y}((x,y),(x',y'))| = |d_X(x,x') - d_Q(y,y')| \le \mathrm{dist}(R) < \delta.$$

So the identity functions $R^X \to R^Y$ and $R^Y \to R^X$ form a δ-interleaving, by Lemma 6.2.5, and thus $\left(d_I^{\mathbf{Diss}}\right)_{/\simeq}(X,Y) \le \delta$.

For the converse, assume given $\alpha : X' \to X$ and $\beta : Y' \to Y$ acyclic morphisms such that X' and Y' are δ-interleaved. Using Lemma 6.2.5, let the interleaving be given by a bijection $f : X' \to Y'$. Define a correspondence $R \subseteq X \times Y$, where $(x,y) \in R$ if and only if $f(\alpha^{-1}(x)) \cap \beta^{-1}(y) \neq \emptyset$. Since α and β are surjective, this defines a correspondence. Assume $(x,y),(x',y') \in R$. Let $a \in X'$ be such that $\alpha(a) = x$ and $\beta(f(a)) = y$, which exists by construction. Similarly, let $b \in X'$ be such that $\alpha(b) = x'$ and $\beta(f(b)) = y'$. Then

$$|d_P(x,x') - d_Q(y,y')| = |d_{X'}(a,b) - d_{Y'}(f(a),f(b))| \le \delta,$$

so $\mathrm{dist}(R) \le \delta$ and thus $2d_{GH}(X,Y) \le \delta$, concluding the proof. □

6.2.2 The Gromov–Hausdorff distance on metric spaces

We now consider the full locally persistent subcategory **epMet** ⊆ **Diss** of ep metric spaces, and show that its quotient interleaving distance is equal to twice the usual Gromov–Hausdorff distance. We also show that the quotient interleaving distance on ep metric spaces inherits completeness from the quotient interleaving distance of **Diss**. We then recover the facts that, when restricted to compact metric spaces, the Gromov–Hausdorff distance is geodesic and restricts to a non-pseudo distance on the collection of isometry classes of compact metric spaces. Finally, we prove a stability result for functions out of the collection of all ep metric spaces and give an alternative characterization of the acyclic morphisms between ep metric spaces.

Consider the category **epMet** (Definition 2.2.5) of ep metric spaces with the locally persistent category structure given by

$$\mathrm{Hom}_{\mathbf{epMet}}(P,Q)_\varepsilon = \left\{ f : P \to Q \text{ a set map} : \forall\, p,p' \in P, d_P(p,p') + \varepsilon \ge d_Q(f(p),f(p')) \right\}$$

as in Example 3.1.4. Note that, by regarding an ep metric space as a dissimilarity space in the natural way, the locally persistent category **epMet** is a full locally persistent subcategory of the locally persistent category **Diss**. We prove a few properties of this embedding. First, a dissimilarity space is an ep metric space if and only if it is weakly equivalent to one.

Lemma 6.2.13. *If $P \in$ **epMet** and $X \in$ **Diss** are such that $P \simeq X$, then $X \in$ **epMet**. If, in addition, P is a compact ep metric space, then so is X.*

Proof. We start with the first claim. Since, by definition, $P \simeq X$ if and only if they are connected by a zig-zag of acyclic morphisms, by induction, it is enough to show that if $\alpha : P \to X$ is an acyclic morphism, then X is an ep metric space, and that if $\alpha : X \to P$ is an acyclic morphism, then X is an ep metric space. In both cases, we have to check that $d_X(x,x) = 0$ for every $x \in X$, and that d_X satisfies the triangle inequality. In the first case, we use the fact that α is surjective and preserves distances, in the second case we just use the fact that α preserves distances.

The proof of the second claim uses exactly the same strategy as the proof of the first claim. □

The locally persistent subcategory **epMet** $\subseteq$ **Diss** is not closed under powering by representables in **Diss**, as the following example shows. In fact, one can show that **epMet** does not admit powers by representables.

Example 6.2.14. Let P be the metric space with three points $\{a,b,c\}$ such that $d_P(a,b) = 1$, $d_P(b,c) = 1$, and $d_P(a,c) = 2$. If $\varepsilon = 1$, then the dissimilarity space P^ε doesn't satisfy the triangle inequality, since $d_{P^\varepsilon}(a,b) = 0$, $d_{P^\varepsilon}(b,c) = 0$, but $d_{P^\varepsilon}(a,c) = 1$.

Let $\mathbf{epMet}_c$ denote the full locally persistent subcategory of **Diss** spanned by compact ep metric spaces. This is a relative locally persistent category, where the acyclic morphisms are taken to be the 0-morphisms that are acyclic morphisms of **Diss**. Similarly, consider the relative locally persistent category **epMet**.

Proposition 6.2.15. *The locally persistent categories* **epMet** *and* $\mathbf{epMet}_c$ *admit weighted sequential limits of morphisms that are part of an interleaving.*

Proof. Assume given $\varepsilon \in \mathbf{R}_+$ and $\varepsilon_i \in \mathbf{R}_+$ for each $i \in \mathbb{N}$ such that $\sum_i \varepsilon_i = \varepsilon$ and let $\overline{\varepsilon}_i = \varepsilon - \sum_{j<i} \varepsilon_j$. Let

$$\cdots \xrightarrow[\varepsilon_i]{f_i} X_i \xrightarrow[\varepsilon_{i-1}]{f_{i-1}} \cdots \xrightarrow[\varepsilon_1]{f_1} X_1 \xrightarrow[\varepsilon_0]{f_0} X_0 \tag{6.2.16}$$

be a weighted sequential limit diagram of dissimilarity spaces such that, for every i, the morphism f_i is part of an ε_i-interleaving.

We now show the following two facts. If all of the dissimilarity spaces are ep metric spaces, then so is the sequential limit; and if all of the dissimilarity spaces are compact ep metric spaces, then so is the sequential limit.

To prove that **Diss** admits sequential limits, we used Proposition 3.2.15. To satisfy the hypotheses, we proved that **Diss** is copowered and powered by representables and that its underlying category admits sequential limits. By Proposition 3.2.15, the sequential limit of Diagram 6.2.16 is computed by taking the categorical limit of

$$\cdots \xrightarrow{f_i} X_i^{\overline{\varepsilon}_i} \xrightarrow{f_{i-1}} \cdots \xrightarrow{f_1} X_1^{\overline{\varepsilon}_1} \xrightarrow{f_0} X_0^{\overline{\varepsilon}_0} \tag{6.2.17}$$

in the underlying category of **Diss**. By Lemma 6.2.5, all the maps in Diagram 6.2.17 are bijections, so we can assume that all the dissimilarity spaces X_i have the same underlying set X and possibly different metrics $d_i : X \times X \to \overline{\mathbf{R}}_+$.

Then, the limit of the diagram can be taken to have underlying set X and metric d_X given by

$$d_X(x, y) = \lim_i \max\left(0, d_i(x, y) - \overline{\varepsilon_i}\right) = \lim_i d_i(x, y),$$

by the description of powers in **Diss** (Lemma 6.2.6). Here, the limit is just a metric limit, which exist since the sequence $\{d_i(x, y)\}_{i\in\mathbb{N}}$ is Cauchy, as $|d_{i+1}(x, y) - d_i(x, y)| \leq \varepsilon_i \to 0$.

The above implies that the generalized metric on the limit X is a uniform limit of the generalized metrics d_i. So, if all of the generalized metrics satisfy the triangle inequality, so does d_X, and if all of the generalized metrics are compact, so is d_X, concluding the proof. □

Our next goal is to prove that the Gromov–Hausdorff distance is geodesic when restricted to compact ep metric spaces. Although the locally persistent category **Diss** admits terminal midpoints, it is not the case that the locally persistent subcategory **epMet** is closed under this construction. So we prove that **epMet** admits coherent factorizations of interleavings by hand. Although not stated in this language, this construction was first performed in [Stu12] in the case of metric measure spaces, and then specialized to metric spaces in [CM18b].

Proposition 6.2.18. *The locally persistent category* **epMet** *admits coherent factorizations of interleavings. If the interleaving is between two compact ep metric spaces, then* *the factorization can be taken so that every object in the factorization is a compact ep metric space.*

Proof. Assume given an ε-interleaving between ep metric spaces P and Q. Let $f : P \to Q$ be the bijection representing the interleaving. Given $\gamma + \delta = \varepsilon$, define an ep metric space M_γ with underlying set P and metric given by

$$\frac{\delta}{\varepsilon} d_P + \frac{\gamma}{\varepsilon} f^*(d_Q),$$

where $f^*(d_Q)(p, p') = d_Q(f(p), f(p'))$.

The identity function $M_\gamma \to P$ is a γ-interleaving since

$$\left| d_P - \left(\frac{\delta}{\varepsilon} d_P + \frac{\gamma}{\varepsilon} f^*(d_Q) \right) \right| = \left| \frac{\gamma}{\varepsilon} (d_P - f^*(d_Q)) \right| \le \frac{\gamma}{\varepsilon} \varepsilon = \gamma,$$

A similar computation shows that $f : M_\gamma \to Q$ gives a δ-interleaving. To see that these factorizations are coherent, note that the identity function gives functions $M_\gamma \to M'_\gamma$ for every $\gamma \le \gamma' \in [0, \varepsilon]$, and that an analogous computation to the one above shows that this function is a $(\gamma' - \gamma)$-interleaving.

To prove that M_γ is a compact ep metric space when P and Q are, note that a sequence of elements in M_γ has a subsequence that converges in P, which in turn, has a subsequence that converges in Q, and thus, converges in M_γ. □

Lemma 6.2.19. *The quotient interleaving distance on* $\mathbf{epMet}_c$ *reflects quotient interleavings.*

This proof is essentially a rewording of the standard proof that the Gromov–Hausdorff distance is geodesic [CM18b, Theorem 1.2].

Proof. We apply Theorem 4.5.18, so we must check that, for $P, P' \in \mathbf{epMet}_c$, the persistent set of quotient interleavings $\mathbb{QI}(P, P') : \mathbf{R}_+ \to \mathbf{SET}$ is right continuous, and that we can lift it to a q-tame persistent topological space with closed structure maps.

Let us instantiate the definition of the persistent set of quotient interleavings to this case:

$$\begin{aligned}\mathbb{QI}(P, P')_\delta = \Big\{ (Z, Z', u, v, f, g) :\; & Z, Z' \in \mathbf{epMet}_c, \\ & u : Z \to P, v : Z' \to P' \text{ surjective and distance preserving}, \\ & f \text{ and } g \text{ form a } \delta\text{-interleaving between } P \text{ and } P' \Big\}.\end{aligned}$$

The structure morphisms are just inclusions in this case, since, by Lemma 6.2.5, f and g are inverse bijections that don't distort the distances more than δ, and thus, they don't distort the distances more than δ' for any $\delta' \geq \delta$.

We first prove that $\mathbb{QI}(P,P')$ is right continuous. To see this, we note that the structure morphisms are the natural inclusions, and that if a pair of inverse bijections f and g doesn't distort the metric more than δ' for every $\delta' > \delta$, then it doesn't distort the metric more than δ.

To lift $\mathbb{QI}(P,P')$ to a persistent topological space, we follow a classical construction. Let $\mathsf{corr}(P,P')_\delta$ denote the set of all correspondences between P and P' of distortion at most δ. For $\delta' \geq \delta$ we have a natural inclusion $\mathsf{corr}(P,P')_\delta \subseteq \mathsf{corr}(P,P')_{\delta'}$ so $\mathsf{corr}(P,P') : \mathbf{R}_+ \to \mathbf{Set}$ is a persistent set. Now, the Hausdorff distance (Definition 2.2.23) endows the set of subsets of $P \times P'$ with an ep metric. We can then give $\mathsf{corr}(P,P')_\delta$ the subspace topology. We note two facts about this topology. First, in the Hausdorff distance, any set is at distance zero from its closure. Second, by Blaschke's theorem [BBI01, Theorem 7.3.8], the set of subsets of $P \times P'$ endowed with the Hausdorff distance is a compact ep metric space.

Let $R_n \subseteq P \times P'$ be a sequence of closed subsets with limit $R \subseteq P \times P'$, which can be taken to be closed. We make the following two claims: if R_n is a correspondence for each $n \in \mathbb{N}$, then R must be a correspondence, and if the distortion of all the correspondences R_n is bounded by some δ, then the distortion of R is also bounded by δ. The proofs of these claims are elementary (see, e.g., [CM18b, Proof of Proposition 1.1]). We deduce that, when endowed with the topology induced by the Hausdorff metric, the persistent topological space $\mathsf{corr}(P,P') : \mathbf{R}_+ \to \mathbf{Top}$ is q-tame with closed structure morphisms.

We now construct a natural transformation $\mathbb{QI}(P,P') \Rightarrow \mathsf{corr}(P,P')$ with surjective components. This lets us transport the topology on $\mathsf{corr}(P,P')$ to a topology on $\mathbb{QI}(P,P')$ that is q-tame with closed structure maps, finishing the proof.

Given $P \xleftarrow{u} Z \xrightarrow{f} Z' \xrightarrow{v} P'$, we can consider the function $(u, v \circ f) : Z \to P \times P'$, and its image $R \subseteq P \times P'$. If $P \xleftarrow{u} Z \xrightarrow{f} Z' \xrightarrow{v} P'$ is a δ-quotient interleaving, that is, it belongs to $\mathbb{QI}(P,P')_\delta$, then R is a correspondence whose distortion is bounded above by δ. So we have constructed a function $\mathbb{QI}(P,P')_\delta \to \mathsf{corr}(P,P')_\delta$. To conclude the proof, we must show that, for each δ, the function $\mathbb{QI}(P,P')_\delta \to \mathsf{corr}(P,P')_\delta$ is surjective. This follows from the fact that every correspondence R with distortion bounded above by δ gives rise to a δ-quotient interleaving $P \leftarrow R^Z \leftrightarrow R^{Z'} \to P'$ as in the proof of Theorem 6.2.12. □

Theorem 6.2.20. *The quotient interleaving distance of* **epMet** *coincides with twice the Gromov–Hausdorff distance and it is intrinsic and complete.*

The quotient interleaving distance of $\mathbf{epMet}_c$ *coincides with twice the Gromov–Hausdorff distance and it is geodesic and complete. Moreover, if* $P, Q \in \mathbf{epMet}_c$ *are compact (non-pseudo) metric spaces such that* $(d_I^{\mathbf{epMet}_c})_{/\simeq}(X,Y) = 0$*, then they are isometric.*

Proof. By Theorem 6.2.12 and Theorem 6.2.11 we know that the first claim holds for the relative locally persitent category **Diss**. To see that it holds for the locally persistent subcategory **epMet**, we use Lemma 6.2.13 to see that **epMet** is closed under acyclic morphisms, Proposition 6.2.15 to see that it is closed under weighted sequential limits of morphisms that are part of interleavings, and Proposition 6.2.18 to see that it admits coherent factorizations of interleavings.

The same analysis shows that the quotient interleaving distance of $\mathbf{epMet}_c$ coincides with twice the Gromov–Hausdorff distance and is complete. To see that it is geodesic, and that it restricts to a non-pseudo metric on isometry classes, we use Lemma 6.2.19 and Proposition 6.2.18 to satisfy the hypothesis of Theorem 4.5.16. □

We now prove a useful stability result for maps out of the collection of ep metric spaces. This can be regarded as a universal property of the Gromov–Hausdorff distance.

Proposition 6.2.21. *Let* P *be an ep metric space and let* $f : \mathbf{epMet} \to P$ *be a function. Assume that for every* $\varepsilon > 0$ *there exists* $\delta > 0$ *such that if* d_1 *and* d_2 *are two ep metrics on a set* X *such that* $||d_1 - d_2||_\infty \leq \delta$*, then* $d_P(f(X,d_1), f(X,d_2)) \leq \varepsilon$*. Assume further that, if there is a surjective and distance preserving map of ep metric spaces* $X \to Y$*, then* $d_P(f(X), f(Y)) = 0$*. Then,* f *is uniformly continuous with respect to the Gromov–Hausdorff distance and the distance* d_P *on* P*. If above we can take* $\delta = c\varepsilon$*, then* f *is* $2c$*-Lipschitz.*

Proof. This follows from the stability result Theorem 4.2.3, using the fact that the quotient interleaving distance is twice the Gromov–Hausdorff distance (Theorem 6.2.12), and the characterization of interleavings in **epMet** (Lemma 6.2.5). □

We finish this section by giving an alternative description of acyclic morphisms between ep metric spaces.

Given an ep metric space P, consider the equivalence relation on the underlying set of P where $p \sim p'$ if and only if $d_P(p,p') = 0$. By the triangle inequality, if $p \sim q$

and $p' \sim q'$, then $d_P(p,p') = d_P(q,q')$. This means that the quotient set $\overline{P} = P/\sim$ inherits a well-defined distance $d_{\overline{P}}([p],[p']) = d_P(p,p')$. Note that the construction $\overline{(-)}$: **epMet** → **epMet** is a functor between categories (not locally persistent categories), since the morphisms in **epMet** are distance non-increasing.

Definition 6.2.22. *A morphism $f : P \to Q$ between ep metric spaces is a* ***pseudo isometry*** *if $\overline{f} : \overline{P} \to \overline{Q}$ is an isometry.*

Lemma 6.2.23. *For two ep metric spaces P and Q the following are equivalent:*

1. *$\overline{P}$ and $\overline{Q}$ are isometric as metric spaces.*

2. *P and Q are connected by a zig-zag of pseudo isometries.*

3. *P and Q are connected by a zig-zag of acyclic morphisms of dissimilarity spaces.*

Proof. Note that any acyclic morphism of dissimilarity spaces between ep metric spaces is necessarily a pseudo isometry, so (3) implies (2). The fact that (2) implies (1) follows at once from the fact that objects connected by a zig-zag of isomorphisms must be isomorphic.

We now show that that (1) implies (3). Note that the quotient map $P \to \overline{P}$ is distance preserving and surjective, so it is an acyclic morphism of dissimilarity spaces for any P. So, if $\overline{P}$ and $\overline{Q}$ are isomorphic, we have a zig-zag $P \to \overline{P} \cong \overline{Q} \leftarrow Q$, and thus a diagram $P \to \overline{P} \leftarrow Q$ consisting of acyclic morphisms of dissimilarity spaces, concluding the proof. □

6.3 Stability of Vietoris–Rips and related filtrations

In this section, we show that the very simple stability result of Theorem 4.2.2 implies the stability of the Vietoris–Rips filtration. This result was first proven in this generality in [BL17], using the homotopy interleaving distance.

The proof that we give here is essentially a rewording of the proof in [Mé17], only using the language of locally persistent categories. The point is to show that, once this language is set up, the proof can be split into two orthogonal parts: showing that Vietoris–Rips is a locally persistent functor, and showing that Vietoris–Rips preserves acyclic morphisms. In Section 6.4, we give another application of this methodology, namely, we prove the stability of a parametrized Vietoris–Rips filtration that maps persistent metric spaces to bi-persistent topological spaces.

Let **sCpx** denote the category of simplicial complexes. An object of **sCpx** consists of a set X together with a family of finite and non-empty subsets S_X of X such that if $\sigma \in S_X$ and $\tau \subseteq \sigma$ with $\tau \neq \emptyset$, then $\tau \in S_X$. Given (X, S_X) and (Y, S_Y) simplicial complexes, a morphism from X to Y consists of a function of sets $f : X \to Y$ such that, for every $\sigma \in S_X$, we have $f(\sigma) \in S_Y$.

Definition 6.3.1. *Let P be a metric space. Define the* ***Vietoris–Rips*** *filtration of P as the following persistent simplicial complex $\mathsf{VR_c}(X) : \mathbf{R} \to \mathbf{sCpx}$. For $r \in \mathbf{R}$, let $\mathsf{VR_c}(X)(r) = (X, S_X(r))$ with*

$$S_X(r) = \left\{ \{x_0, \ldots, x_n\} : d(x_i, x_j) \leq r \textit{ for all } 0 \leq i, j \leq n \right\}.$$

The structure maps are the natural inclusions.

Note that the definition of the Vietoris–Rips filtration doesn't make any use of the triangle inequality of P. In fact, exactly the same definition works for P a dissimilarity space (Definition 6.2.1).

In order to show that the Vietoris–Rips filtration is stable, we extend it to a locally persistent functor.

Lemma 6.3.2. *The Vietoris–Rips filtration $\mathsf{VR_s} : \mathbf{Diss} \to \mathbf{sCpx}^{\mathbf{R}}$ is a locally persistent functor.*

Proof. We have defined $\mathsf{VR_s}$ on objects. We must now show that, given $\varepsilon \in \mathbf{R}_+$ and a morphism $f \in \mathrm{Hom}_{\mathbf{Diss}}(X, Y)_\varepsilon$ between dissimilarity spaces X and Y, there is an induced morphism $\mathrm{Hom}_{\mathbf{sCpx}^{\mathbf{R}}}(X,Y)_\varepsilon$, and that this mapping respects identities and composition.

Let $f \in \mathrm{Hom}_{\mathbf{Diss}}(X,Y)_\varepsilon$ be an ε-morphism between dissimilarity spaces X and Y. Given $r \in \mathbf{R}$, let $\mathsf{VR_s}(f)(r) : \mathsf{VR_s}(X)(r) \to \mathsf{VR_s}(Y)(r+\varepsilon)$ be the map of simplicial complexes that sends $x \in X$ to $f(x) \in Y$. To see that it is simplicial, we must show that if $\sigma \subseteq \mathsf{VR_s}(X)(r)$ is a simplex, then $f(\sigma) \subseteq \mathsf{VR_s}(X)(r+\varepsilon)$ is also a simplex. This is because, if $\sigma = \{x_0,\dots,x_n\}$, then $d_X(x_i,x_j) \leq r$ for all $0 \leq i,j \leq n$, and thus $d_Y(f(x_i),f(x_j)) \leq r+\varepsilon$ for all $0 \leq i,j \leq n$, since $f \in \mathrm{Hom}_{\mathbf{Diss}}(X,Y)_\varepsilon$.

The fact that this mapping preserves identities and composition is evident. □

Recall that there is a geometric realization functor $|-| : \mathbf{sCpx} \to \mathbf{Top}$ (see, e.g., [Spa12, Chapter 3, Section 1]). We will need the following version of Quillen's Theorem A.

Lemma 6.3.3. *Let $f : X \to Y$ be a map of simplicial complexes such that f is surjective on the underlying sets, and such that $\sigma \subseteq X$ is a simplex of X if and only if $f(\sigma)$ is a simplex of Y. Then $|f| : |X| \to |Y|$ is a weak equivalence of topological spaces.*

Proof. There are at least to ways to prove this. One option is to choose arbitrary orders on the underlying sets of X and Y in such a way that f preserves the orders. These orders induce two simplicial sets X_s and Y_s and a simplicial map $f_s : X_s \to Y_s$ as follows. We let the n-simplices of X_s be given by lists $(x_0,\dots,x_n)$ such that $\{x_0,\dots,x_n\} \in S_X$ and such that $x_i \leq x_{i+1}$. The simplicial set Y_s is defined analogously, and the simplicial map f_s is defined in the only possible way, after prescribing that $f_s((x)) = (f(x))$ for every $x \in X$. It is straightforward to see that the preimage of every simplex of Y_s along f_s is contractible, by assumption, and thus that f_s is a weak equivalence of simplicial sets. This implies that $|f_s| : |X_s| \to |Y_s|$ is a weak equivalence of topological spaces, but $|f_s| : |X_s| \to |Y_s|$ is equal to $|f| : |X| \to |Y|$, after identifying $|X|$ with $|X_s|$ and $|Y|$ with $|Y_s|$.

The other option is to let P_X be the poset S_X. The poset P_X can be seen as a category, and thus as a simplicial set. The geometric realization of P_X is naturally homeomorphic to the realization of X: the poset P_X corresponds to the barycentric subdivision of X. Performing the same construction with Y, we get a diagram of topological spaces

$$\begin{tikzcd} |P_X| \arrow[r, "|P_f|"] \arrow[d, "\simeq"'] & |P_Y| \arrow[d, "\simeq"] \\ |X| \arrow[r, "|f|"] & |Y| \end{tikzcd}$$

in which the vertical maps are homeomorphisms. Now, by assumption, and Quillen's theorem A ([Qui73]), the top horizontal map is a weak equivalence, and thus, $|f|$ is a weak equivalence. □

We can compose the Vietoris–Rips filtration with the geometric realization functor to get a locally persistent functor

$$\mathsf{VR} = |-| \circ \mathsf{VR}_\mathsf{s} : \mathbf{Diss} \to \mathbf{Top}^{\mathbf{R}}.$$

Proposition 6.3.4. *The locally persistent functor* $\mathsf{VR} : \mathbf{Diss} \to \mathbf{Top}^{\mathbf{R}}$ *maps acyclic morphisms to weak equivalences.*

Proof. The acyclic morphisms of **Diss** are the surjective and distance preserving maps. Note that, if $f : X \to Y$ is a surjective and distance preserving map between dissimilarity spaces, then, for every $r \in \mathbf{R}$, the morphism of simplicial complexes

$$\mathsf{VR}_\mathsf{s}(f) : \mathsf{VR}_\mathsf{s}(X)(r) \to \mathsf{VR}_\mathsf{s}(Y)(r)$$

is surjective on underlying sets and has the property that $\sigma \subseteq X$ is a simplex of $\mathsf{VR}_\mathsf{s}(X)(r)$ if and only if $f(\sigma) \subseteq Y$ is a simplex of $\mathsf{VR}(Y)(r)$, since f is distance preserving. It then follows from Lemma 6.3.3 that

$$\mathsf{VR}(f)(r) = |\mathsf{VR}(f)(r)| : |\mathsf{VR}(X)(r)| \to |\mathsf{VR}(Y)(r)|$$

is a weak equivalence of topological spaces. □

We then have the following.

Theorem 6.3.5 (cf. [BL17])**.** *The mapping* $\mathsf{VR} : \mathbf{Diss} \to \mathbf{Top}^{\mathbf{R}}$ *is* 2*-Lipschitz with respect to the Gromov–Hausdorff distance and the homotopy interleaving distance.*

Proof. By Theorem 4.2.2, we have that VR is 1-Lipschitz with respect to the quotient interleaving distance on **Diss** and the quotient interleaving distance on $\mathbf{Top}^{\mathbf{R}}$, using

Proposition 6.3.4 and Lemma 6.3.2. But the quotient interleaving distance on **Diss** is twice the Gromov–Hausdorff distance (Theorem 6.2.12), and the quotient interleaving distance on $\mathbf{Top}^{\mathbf{R}}$ is the homotopy interleaving distance (Theorem 6.1.7). □

Alternatively, to prove Theorem 6.3.5 in the case of ep metric spaces, one can use the universal property of the Gromov–Hausdorff distance (Proposition 6.2.21). This universal property can actually be used to prove the stability of many related invariants of metric spaces. These include, for example, the Čech filtration, the persistent homology of the filtrations given in [Cho19a] (called nerve functors there), as well as the filtrations introduced in [CCMSW17]. The methodology is always the same: one shows that the invariant maps strict interleavings to close-by invariants, and acyclic morphisms to invariants at distance 0.

6.4 The Gromov–Hausdorff-interleaving distance on persistent dissimilarity spaces

One of the advantages of having a categorical framework to define distances, such as the one developed in this thesis, is compositionality. In this context, this refers to the process of combining distances on simple objects to get a distance on a class of more structured objects. An example of this is the formation of a locally persistent category $\mathscr{C}^{\mathbf{R}}$ of persistent objects of a locally persistent category $\mathscr{C}$, described in Section 5.1.2. This locally persistent category structure takes into account the shifts of the objects of $\mathscr{C}^{\mathbf{R}}$ as well as the "pointwise" locally persistent structure of $\mathscr{C}$. In this section, we give three applications of this construction.

In Section 6.4.2, we use the quotient interleaving distance on the category of persistent dissimilarity spaces to generalize the Gromov–Hausdorff distance on filtered metric spaces considered in [CCSGMO09] and [CM10c]. In Section 6.4.3, we prove the stability of a parametrized version of the Vietoris–Rips filtration, that maps persistent metric spaces to bi-persistent topological spaces. In Section 6.4.4, we use the quotient interleaving distance on a relative locally persistent category category of bi-persistent dissimilarity spaces to generalize the λ-slack interleaving distance on dynamic metric spaces introduced in [KM20].

The purpose of these examples is not the generalization of the distances in itself, but to show how our framework allows one to easily combine distances such as the Gromov–Hausdorff distance and the interleaving distance in order to compare objects that have both a metric structure and a persistent structure.

6.4.1 Persistent dissimilarity spaces

In this section, we show how to use the theory developed in this thesis to define a distance on persistent dissimilarity spaces and, in particular, on persistent metric spaces that takes into account metric perturbations and persistence perturbations.

Recall from Section 6.2.1 that there is a locally persistent category of dissimilarity spaces that generalizes the locally persistent category of metric spaces. We now specialize Definition 5.1.7 to dissimilarity spaces and define a locally bi-persistent category of persistent dissimilarity spaces. The idea is that, for $\varepsilon, \delta \in \mathbf{R}_+$, an (ε, δ)-morphism is a morphism that shifts the persistence degree by ε and that is allowed to increase the metric by at most δ.

Definition 6.4.1. *Consider the category* $\mathbf{Diss}^{\mathbf{R}}$ *the has as objects the standard functors* $\mathbf{R} \to \mathbf{Diss}_0$ *from the poset* $\mathbf{R}$ *to the category of dissimilarity spaces with distance non-increasing maps. This category admits a* $\mathbf{Set}^{\mathbf{R}_+ \times \mathbf{R}_+}$*-enrichment, given as follows. For* $X, Y \in \mathbf{Diss}^{\mathbf{R}}$ *and* $\varepsilon, \delta \in \mathbf{R}_+$*, let*

$$\mathrm{Hom}_{\mathbf{Diss}^{\mathbf{R}}}(X,Y)_{(\varepsilon,\delta)} = \mathrm{Nat}(X, Y^{(\varepsilon,\delta)}),$$

where $Y^{(\varepsilon,\delta)}(r)$ *is the dissimilarity space with underlying set* $Y(\varepsilon + r)$ *and metric given by*

$$\max(d_{Y(\varepsilon+r)}(y,y') - \delta, 0),$$

for $y, y' \in Y$*.*

From now on, fix $\vec{v} \in \mathbf{R}_+ \times \mathbf{R}_+$. As explained in Section 5.1.2, we get a locally persistent category $\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}$ where an ε-morphism is given by an $(\varepsilon\vec{v})$-morphism in the locally bi-persistent category $\mathbf{Diss}^{\mathbf{R}}$. Concretely, if $\vec{v} = (v_1, v_2)$, for $X, Y \in \mathbf{Diss}^{\mathbf{R}}$ and $\varepsilon \in \mathbf{R}_+$, we have

$$\mathrm{Hom}_{\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}}(X,Y)_{\varepsilon} = \mathrm{Nat}(X, Y^{(\varepsilon v_1, \varepsilon v_2)}).$$

As also explained in Section 5.1.2, the locally persistent category $\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}$ inherits a class of acyclic morphisms from $\mathbf{Diss}$, which endows $\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}$ with the structure of a relative locally persistent category.

Definition 6.4.2. *Let* $X, Y \in \mathbf{Diss}^{\mathbf{R}}_{\vec{v}}$ *and let* $f : X \to Y$ *be a* 0*-morphism. We say that* f *is an* ***acyclic morphism*** *if all of its components are surjective and distance preserving.*

Definition 6.4.3. *The quotient interleaving distance on* $\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}$ *is called the* ***Gromov–Hausdorff-interleaving distance****.*

We remark that the Gromov–Hausdorff-interleaving distance actually depends on the vector $\vec{v}$, but, by Proposition 5.1.4, any two choices of such $\vec{v}$ with strictly positive coordinates will yield bi-Lipschitz equivalent distances.

The Gromov–Hausdorff-interleaving distance is well behaved.

Proposition 6.4.4. *The Gromov–Hausdorff-interleaving distance on* $\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}$ *is intrinsic* *and complete, and satisfies:*

$$\left(d_I^{\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}}\right)_{/\simeq}(X,Y) = \inf\left\{\delta : \exists X' \simeq X, Y' \simeq Y,\ X' \textit{ and } Y' \textit{ are } \delta\textit{-interleaved}\right\}$$
$$= \inf\left\{\delta : \exists \textit{ acyclic morphisms } X' \to X \textit{ and } Y' \to Y \textit{ such that } X' \textit{ and } Y' \textit{ are } \delta\textit{-interleaved}\right\}.$$

Proof. This follows at once from Proposition 5.1.9. □

Remark 6.4.5. As explained in Remark 5.1.10, for every $n \in \mathbb{N}$, we can construct a locally $\mathbf{R}_+^{n+1}$-persistent category $\mathbf{Diss}^{\mathbf{R}^n}$ of multi-persistent dissimilarity spaces. Moreover, given $\vec{v} \in \mathbf{R}_+^{n+1}$ we obtain a relative locally persistent category $\mathbf{Diss}^{\mathbf{R}^n}_{\vec{v}}$. We also refer to the quotient interleaving distance of this relative locally persistent category as the Gromov–Hausdorff-interleaving distance. As is clear, Proposition 6.4.4 also holds for this distance.

6.4.2 Filtered metric spaces

We now show that the Gromov–Hausdorff-interleaving distance generalizes the distance on finite filtered metric spaces used in [CM10c].

In [CM10c], multi-parameter hierarchical clustering algorithms are studied. The algorithms they study take as input a finite filtered metric space. A **finite filtered metric space** consists of a finite metric space (X, d_X) together with a filtering function $f_X : X \to \mathbf{R}$, which is not required to satisfy any assumptions. Let **ffMet** denote the collection of all finite filtered metric spaces. Let $X, Y \in$ **ffMet**, and let $R \subseteq X \times Y$ be a correspondence between the underlying sets. The **filtered distortion** of R is given by

$$\mathrm{dist}_f(R) = \max\left(\mathrm{dist}(R), ||\pi_X \circ f_X - \pi_Y \circ f_Y||_\infty\right),$$

where $\mathrm{dist}(R)$ denotes the distortion of R as a correspondence between metric spaces (as in Definition 2.2.26). The following distance between finite filtered metric spaces is given in [CM10c, Definition 2]:

$$\mathbf{D}(X,Y) = \inf_{\substack{R \subseteq X \times Y \\ \text{correspondence}}} \mathrm{dist}_f(R).$$

We now explain how the metric $\mathbf{D}$ can be interpreted as a quotient interleaving distance. Given $X \in \mathbf{ffMet}$, let $f_X^{-1} \in \mathbf{Diss}^{\mathbf{R}}$ be given by

$$f_X^{-1}(r) = f_X^{-1}((-\infty, r)) \in \mathbf{Diss},$$

for every $r \in \mathbf{R}$. The structure maps are given by the natural inclusions $f^{-1}((-\infty, r)) \subseteq f^{-1}((-\infty, r'))$ for $r \leq r' \in \mathbf{R}$. This gives a mapping $\mathbf{ffMet} \to \mathbf{Diss}^{\mathbf{R}}$, which lets us interpret the collection of finite filtered metric spaces as a full locally persistent subcategory of $\mathbf{Diss}^{\mathbf{R}}$.

Theorem 6.4.6. *Let $\vec{v} = (1,1)$. Let $X, Y \in \mathbf{ffMet}$. Then*

$$\left(d_I^{\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}}\right)_{/\simeq}(f_X^{-1}, f_Y^{-1}) = \mathbf{D}(X, Y).$$

We use the same methodology as in the proof of Theorem 6.2.12. For simplicity, we write $\mathbf{Diss}^{\mathbf{R}}$ instead of $\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}$.

Proof. Let $R \subseteq X \times Y$ be a correspondence such that $\mathrm{dist}_f(R) \leq \delta$. Consider the finite pseudo metric space R^X with underlying set R and where $d_{R^X}((x,y),(x',y')) = d_X(x,x')$. Define a function $f_{R^X} : R^X \to \mathbf{R}$ by $f_{R^X} = f_X \circ \pi_X$. Consider then the persistent dissimilarity space given by $f_{R^X}^{-1} : \mathbf{R} \to \mathbf{Diss}$. Define, analogously, a persistent dissimilarity space $f_{R^Y}^{-1} : \mathbf{R} \to \mathbf{Diss}$. The projections $\pi_X : R^X \to X$ and $\pi_Y : R^Y \to Y$ induce acyclic morphisms $f_{R^X}^{-1} \to f_X^{-1}$ and $f_{R^Y}^{-1} \to f_Y^{-1}$ of $\mathbf{Diss}^{\mathbf{R}}$. Moreover, as in the proof of Theorem 6.2.12, the identity $R^X \to R^Y$ induces a δ-interleaving between $f_{R^X}^{-1}$ and $f_{R^Y}^{-1}$, so $\left(d_I^{\mathbf{Diss}^{\mathbf{R}}}\right)_{/\simeq}(f_X^{-1}, f_Y^{-1}) \leq \delta$.

The other direction is more interesting. Assume that $\left(d_I^{\mathbf{Diss}^{\mathbf{R}}}\right)_{/\simeq}(f_X^{-1}, f_Y^{-1}) < \delta$. By the second characterization of the quotient interleaving distance in Proposition 6.4.4 there exist persistent dissimilarity spaces S and T, acyclic morphisms $\alpha : S \to f_X^{-1}$ and $\beta : T \to f_Y^{-1}$, and a δ-interleaving between S and T, given by natural transformations φ and ψ. Let S' be the colimit of S, seen as a functor $S : \mathbf{R} \to \mathbf{Diss}_0$. Similarly, let T' be the colimit of T. By the interleaving, there is a bijection $\varphi' : S' \to T'$ that doesn't distort the metric more than δ, that is, φ' is a δ-interleaving of dissimilarity spaces. There are also surjective and distance preserving functions $\alpha' : S' \to X$ and $\beta' : T' \to Y$, which are just the colimit of the acyclic morphisms $S \to f_X^{-1}$ and $T \to f_Y^{-1}$ respectivelly.

Define the relation $R \subseteq X \times Y$ as follows. We have $(x, y) \in R$ if and only if $|f_X(x) - f_Y(y)| \leq \delta$ and $\varphi'(\alpha'^{-1}(x)) \cap \beta'^{-1}(y) \neq \emptyset$. Assume for the moment that R is a correspondence. By the same argument in the proof of Theorem 6.2.12, the distortion of R is bounded by δ, so the filtered distortion of R is also bounded above by δ, by construction of R.

It only remains to be shown that R is a correspondence. We show that every $x \in X$ is related to some $y \in Y$, a symmetric argument finishes the proof. To see this, let $x \in X$ and let $r = f_X(x)$. So $x \in f_X^{-1}(r)$ and $x \notin f_X^{-1}(r')$ for every $r' < r$. It is enough to show that there is $y \in f_X^{-1}(r+\delta)$ such that $(x,y) \in R$ and such that $y \notin f_X^{-1}(r')$ for every $r' < r - \delta$.

Given $x' \in S(r)$, which exists, since the components of α are surjective, consider $\varphi(x') \in T(r+\varepsilon)$. We now show that $(x, \beta(\varphi(x'))) \in R$, concluding the proof. To prove this, we must show that $f_Y(\beta(\varphi(x'))) \geq r - \varepsilon$. If this is not the case, then there is $y'' \in T(r')$ with $r' < r - \varepsilon$ such that y'' maps to $\varphi(x')$ under the structure map $T(r') \to T(r+\varepsilon)$. By the interleaving, we must have $\alpha(\psi(y'')) = x$, but then $\psi(y'') \in S(r'+\varepsilon)$ and $r' + \varepsilon < r$. This is a contradiction, since if $x' \in S(r)$ is such that $\alpha'(x') = x$ then there is no $x'' \in S(r')$ for $r' < r$ that maps to it under the structure map $S(r') \to S(r)$, as $x \notin f_X^{-1}(r')$. □

6.4.3 Parametrized Vietoris–Rips

The Gromov–Hausdorff-interleaving distance is a good abstraction since it lets us lift stable invariants of metric spaces to stable invariants of persistent metric spaces. In this section, we show how this works in the case of the Vietoris–Rips filtration.

Since $\mathsf{VR} : \mathbf{Diss} \to \mathbf{Top}^{\mathbf{R}}$ is a locally persistent functor, it induces a locally bi-persistent functor $\mathsf{VR}_* : \mathbf{Diss}^{\mathbf{R}} \to \mathbf{Top}^{\mathbf{R}^2}$, that is, a $\mathbf{Set}^{\mathbf{R}^2}$-enriched functor between $\mathbf{Set}^{\mathbf{R}^2}$-enriched categories. Choosing $\vec{v} = (v_1, v_2) \in \mathbf{R}_+ \times \mathbf{R}_+$, we get a locally persistent functor $\mathsf{VR}_* : \mathbf{Diss}^{\mathbf{R}}_{\vec{v}} \to \mathbf{Top}^{\mathbf{R}^2}_{\vec{v}}$.

Since the acyclic morphisms of $\mathbf{Diss}^{\mathbf{R}}_{\vec{v}}$ and $\mathbf{Top}^{\mathbf{R}^2}_{\vec{v}}$ are the natural transformations all of whose components are acyclic morphisms, it follows from Proposition 6.3.4 that VR_* preserves acyclic morphisms. From Proposition 5.1.11, we deduce the following.

Theorem 6.4.7. *The mapping* $\mathsf{VR}_* : \mathbf{Diss}^{\mathbf{R}}_{\vec{v}} \to \mathbf{Top}^{\mathbf{R}^2}_{\vec{v}}$ *that applies the Vietoris–Rips filtration componentwise to a persistent dissimilarity space is* 1*-Lipschitz.* □

6.4.4 Dynamic metric spaces

We now give a high-level explanation of how to use quotient interleaving distances to recover the λ-slack interleaving distance on dynamic metric spaces introduced in

[KM20].

In [KM20], dynamic metric spaces are studied. A **dynamic metric space** consists of a finite set X together with a function $d_X : \mathbb{R} \times X \times X \to [0,\infty)$ such that

- ▷ for every $r \in \mathbb{R}$, the function $d_X(t) : X \times X \to [0,\infty)$ is a pseudo metric;
- ▷ for every $x \neq x' \in X$, the function $d_X(-)(x,x') : \mathbb{R} \to [0,\infty)$ is not identically 0 and continuous.

A distance between dynamic metric spaces is also introduced, called the λ-slack interleaving distance ([KM20, Definition 2.10]). This distance takes distortion in time as well as distortion of the metric into account, and is defined using tripods, which are a simple generalization of correspondences. We will not describe this metric here, but we will provide an equivalent characterization of this metric, using the machinery of this thesis.

We instantiate the multi-dimensional case of Definition 6.4.3 (Remark 6.4.5) to the case of tri-persistent dissimilarity spaces. We do this by using the isomorphism of posets $\mathbf{R} \cong \mathbf{R}^{\mathrm{op}}$ that maps r to $-r$.

Definition 6.4.8. *Consider the category* $\mathbf{Diss}^{\mathbf{R}^{\mathrm{op}} \times \mathbf{R}}$ *that has as objects the standard functors* $\mathbf{R}^{\mathrm{op}} \times \mathbf{R} \to \mathbf{Diss}_0$ *from the poset* $\mathbf{R}^{\mathrm{op}} \times \mathbf{R}$ *to the category of dissimilarity spaces with distance non-increasing maps. This category admits a* $\mathbf{Set}^{\mathbf{R}_+ \times \mathbf{R}_+ \times \mathbf{R}_+}$*-enrichment, given as follows. For* $X, Y \in \mathbf{Diss}^{\mathbf{R}^{\mathrm{op}} \times \mathbf{R}}$ *and* $\varepsilon_1, \varepsilon_2, \delta \in \mathbf{R}_+$*, let*

$$\mathrm{Hom}_{\mathbf{Diss}^{\mathbf{R}}}(X,Y)_{(\varepsilon_1,\varepsilon_2,\delta)} = \mathrm{Nat}(X, Y^{(\varepsilon_1,\varepsilon_2,\delta)}),$$

where $Y^{(\varepsilon_1,\varepsilon_2,\delta)}((t_1,t_2))$ *is the dissimilarity space with underlying set* $Y(t_1-\varepsilon_1, t_2+\varepsilon_2)$ *and metric given by*

$$\max(d_{Y(t_1-\varepsilon_1,t_2+\varepsilon_2)}(y,y') - \delta, 0),$$

for $y, y' \in Y$.

As in Section 6.4, we fix $\vec{v} = (v_1, v_2, v_3) \in \mathbf{R}_+ \times \mathbf{R}_+ \times \mathbf{R}_+$ and we let the Gromov–Hausdorff-interleaving distance on $\mathbf{Diss}^{\mathbf{R}^{\mathrm{op}} \times \mathbf{R}}$ be the quotient interleaving distance of $\mathbf{Diss}^{\mathbf{R}^{\mathrm{op}} \times \mathbf{R}}_{\vec{v}}$, where the acyclic morphisms are the componentwise acyclic morphisms. We now explain how to use the Gromov–Hausdorff-interleaving distance to compare dynamic metric spaces. We start with one of the main constructions of [KM20].

Let **DMS** denote the collection of all dynamic metric spaces. We can interpret **DMS** as a full locally persistent subcategory of $\mathbf{Diss}^{\mathbf{R}^{\mathrm{op}} \times \mathbf{R}}$ using the following construction.

Given $(X, d_X) \in \mathbf{DMS}$, let $F(X) \in \mathbf{Diss}^{\mathbf{R}^{\mathrm{op}} \times \mathbf{R}}$ be defined as follows. If $t_1 > t_2$, let

$$F(X)(t_1, t_2) = \emptyset,$$

and if $t_1 \le t_2$, let

$$F(X)(t_1, t_2) = \left(X, \bigvee_{t \in [t_1, t_2]} d_X(t)(-,-) \right) \in \mathbf{Diss}$$

where $\vee_{t \in [t_1,t_2]} d_X(t)(-,-) : X \times X \to [0,\infty]$ denotes the function that, at $x, x' \in X$, takes the value

$$\inf_{t \in [t_1, t_2]} d_X(t)(x, x').$$

Using the characterization of the quotient interleaving distance of $\mathbf{Diss}_{\vec{v}}^{\mathbf{R}^{\mathrm{op}} \times \mathbf{R}}$ as an infimum over quotient interleavings (as in Proposition 6.4.4), one can show that the λ-slack interleaving distance between dynamic metric spaces is equivalent to the quotient interleaving distance $\left(d_I^{\mathbf{Diss}_{\vec{v}}^{\mathbf{R}^{\mathrm{op}} \times \mathbf{R}}} \right)_{/\simeq}$ restricted to dynamic metric spaces. More specifically, the λ-slack interleaving distance is obtained by taking $\vec{v} = (1, 1, \lambda)$.

6.5 Stability of the kernel density filtration

In this section, we define a filtration on compact metric probability spaces and show that it is stable with respect to the Gromov–Hausdorff–Prokhorov metric. This filtration is a generalization of the degree-Rips bi-filtration ([LW15]), for which a stability theorem related to the one presented here has been established in unpublished work of Blumberg and Lesnick. The kernel density filtration and its stability appears in [RS20] and is joint work with Alexander Rolle.

Fixing a sufficiently well-behaved kernel K (Definition 6.5.5), we assign, to each compact metric probability space X, the bi-filtration of X that maps s, k to

$$\left\{x \in X : \int_{x' \in X} K\left(\frac{d_X(x,x')}{s}\right) \mathrm{d}\mu_X \geq k\right\}.$$

This filtration is formally defined in Section 6.5.2, where we interpret it as a functor

$$\mathsf{KDF}(X) : \mathbf{R} \times \mathbf{R}^{\mathrm{op}} \to \mathbf{Diss}.$$

The category $\mathbf{Diss}^{\mathbf{R}\times\mathbf{R}^{\mathrm{op}}}$ has a relative locally persistent category structure (Remark 6.4.5). We will prove the following.

Theorem 6.5.1. *The mapping* KDF*, from compact metric probability spaces to* $\mathbf{Diss}^{\mathbf{R}\times\mathbf{R}^{\mathrm{op}}}$*, is uniformly continuous with respect to the Gromov–Hausdorff–Prokhorov distance and the quotient interleaving distance. If* KDF *is constructed using the uniform kernel, then it is* 2*-Lipschitz.*

As a direct consequence of Theorem 6.5.1 and Theorem 6.4.7, we get the following.

Corollary 6.5.2. *For X a compact metric probability space, we have an associated three-parameter persistent module* $H_n \circ \mathsf{VR} \circ \mathsf{KDF}(X) : \mathbf{R} \times \mathbf{R} \times \mathbf{R}^{\mathrm{op}} \to \mathbf{Vec}_k$ *obtained by taking n-th homology with coefficients in a field k of the Vietoris–Rips filtration applied objectwise to the filtration* $\mathsf{KDF}(X) : \mathbf{R} \times \mathbf{R}^{\mathrm{op}} \to \mathbf{Diss}$*. This construction is uniformly continuous with respect to the Gromov–Hausdorff–Prokhorov distance and the interleaving distance on three-parameter persistent vector spaces, using the direction vector* $\vec{v} = (1,1,1)$*. If* KDF *is defined using the uniform kernel, then the above construction is* 2*-Lipschitz.* □

In order to prove the theorem, we extend the filtration to a more general class of objects: weighted dissimilarity spaces.

6.5.1 The Gromov–Hausdorff–Prokhorov metric

We start by recalling the definition of the Gromov–Hausdorff–Prokhorov metric.

Definition 6.5.3. *Let μ, ν be Borel probability measures on a metric space Z. The* ***Prokhorov distance*** *between μ and ν is*

$$d_P(\mu,\nu) = \inf\left\{\varepsilon > 0 : \mu(A) \leq \nu(A^\varepsilon) + \varepsilon \text{ and } \nu(A) \leq \mu(A^\varepsilon) + \varepsilon \text{ for all closed sets } A \subseteq Z\right\},$$

where for a subset $A \subseteq Z$ and $\varepsilon \geq 0$, we let $A^\varepsilon = \{z \in Z : \exists a \in A, d_Z(a,z) < \varepsilon\}$.

Definition 6.5.4. *Let $(X,\mu_X),(Y,\mu_Y)$ be compact metric probability spaces. The* ***Gromov–Hausdorff–Prokhorov*** *distance between (X,μ_X) and (Y,μ_Y) is*

$$d_{GHP}(X,Y) = \inf_{i,j}\left(\max(d_H^Z(i(X), j(Y)),\ d_P(i_*\mu_X, j_*\mu_Y))\right),$$

where the infimum is taken over all isometric embeddings $i : X \to Z$ and $j : Y \to Z$ into a common metric space Z.

Say that two metric probability spaces (X,μ_X) and (Y,μ_Y) are isometry-equivalent if there is a bijective isometry $\psi : X \to Y$ such that $\psi_*(\mu_X) = \mu_Y$. Then, the Gromov–Hausdorff–Prokhorov distance is a metric on the set of isometry-equivalence classes of compact metric probability spaces; see, e.g., [Mie09].

6.5.2 The kernel density filtration

We now formally define the kernel density filtration. For this, we restrict our attention to a general well-behaved class of kernels.

Definition 6.5.5. *A* ***kernel*** *is a non-increasing function $K : \mathbf{R}_+ \to \mathbf{R}_+$ that is continuous from the right and such that $0 < \int_0^\infty K(r)\,\mathrm{d}r < \infty$.*

Note that, in particular, $K(0) > 0$ and $\lim_{r\to\infty} K(r) = 0$.

Example 6.5.6. Many kernels used for density estimation are kernels in the above sense. We will be particularly interested in $K = \mathbf{1}_{\{r<1\}} : \mathbf{R}_+ \to \mathbf{R}_+$, with $K(x) = 1$ if $x < 1$ and $K(x) = 0$ otherwise. We refer to this as the **uniform kernel**.

Definition 6.5.7. *Let K be a kernel, and let X be a metric probability space. Define the* ***local density estimate*** *of a point $x \in X$ at scale $s > 0$ as*

$$\left(\mu_X * K_s\right)(x) := \int_{x' \in X} K\left(\frac{d_X(x, x')}{s}\right) \mathrm{d}\mu_X.$$

Given $s, k \in \mathbf{R}$, let $X_{[s,k]} \subseteq X$ be the sub-metric space given by

$$X_{[s,k]} = \begin{cases} \left\{x \in X \,:\, \left(\mu_X * K_s\right)(x) \geq k\right\}, & \text{if } s, k > 0 \\ X, & \text{if } k \leq 0 \\ \emptyset, & \text{if } s \leq 0 \text{ and } k > 0. \end{cases}$$

Note that, since K is non-increasing, we have $X_{[s,k]} \subseteq X_{[s',k']} \subseteq X$ whenever $s' \geq s$ and $k' \leq k$. This forms a 2-parameter filtration of X, which we call the ***kernel density filtration*** *of X.*

The following lemma will be useful when proving the stability of the kernel density filtration.

Lemma 6.5.8. *Let K be a kernel, and let X be a metric probability space. Let $K^{-1} : \mathbf{R}_{>0} \to \mathbf{R}_+$ be defined as $K^{-1}(t) = \min\{u : K(u) \leq t\}$. Then K^{-1} is a non-increasing function with compact support, and we have, for every $x \in X$,*

$$\left(\mu_X * K_s\right)(x) = \int_0^\infty \mu_X\left(B(x, sK^{-1}(r))\right) \mathrm{d}r.$$

Proof. Since $K(r) \to 0$ as $r \to \infty$, for every $t > 0$ the set $\{u : K(u) \leq t\}$ is non-empty. Moreover, K is continuous from the right, so the set has a minimum, and thus K^{-1} is well-defined. The fact that K^{-1} is non-increasing is clear, and the fact that it has compact support follows from the fact that K is bounded.

To prove the statement about $\left(\mu_X * K_s\right)$, we need the following fact about K^{-1}: for every $s, t \in \mathbf{R}_+$ we have $K^{-1}(t) > s$ if and only if $t < K(s)$. We prove this now. Having $t < K(s)$ is equivalent to s not belonging to the set $\{u : K(u) \leq t\}$, which in turn is equivalent to s being strictly less than any u such that $K(u) \leq t$. This last statement is equivalent to $s < \min\{u : K(u) \le t\} = K^{-1}(t)$. We finish the proof by computing

$$\begin{aligned}\int_{x' \in X} K\left(\frac{d(x,x')}{s}\right) \mathrm{d}\mu_X &= \int_{x' \in X} \int_0^\infty \mathbf{1}_{\left\{r < K\left(\frac{d(x,x')}{s}\right)\right\}} \mathrm{d}r \, \mathrm{d}\mu_X \\ &= \int_{x' \in X} \int_0^\infty \mathbf{1}_{\{d(x,x') < sK^{-1}(r)\}} \mathrm{d}r \, \mathrm{d}\mu_X \\ &= \int_0^\infty \int_{x' \in X} \mathbf{1}_{\{d(x,x') < sK^{-1}(r)\}} \mathrm{d}\mu_X \, \mathrm{d}r \\ &= \int_0^\infty \mu_X\left(B(x, sK^{-1}(r)\right) \mathrm{d}r.\end{aligned}$$

□

6.5.3 Weighted dissimilarity spaces and stability of the kernel density filtration

In order to prove the stability of the kernel density filtration, it is convenient to generalize it so that its domain category becomes a relative locally persistent category. The idea is to first prove the stability with respect to the quotient interleaving distances, and then relate the quotient interleaving distance of the domain to the Gromov–Hausdorff–Prokhorov distance.

Definition 6.5.9. *A* ***weighted dissimilarity space*** *consists of a dissimilarity space* (X, d_X) *together with a function* $M_X : X \times \mathbf{R}_+ \to [0,\infty)$ *such that, for* $x \in X$ *and* $r \le s \in \mathbf{R}_+$*, we have* $M_X(x,r) \le M_X(x,s)$*.*

For X a dissimilarity space, the number $M_X(x,r) \in [0,\infty)$ should be interpreted as the measure of the ball of radius r centered at x.

Example 6.5.10. Any metric probability space (X, d_X, μ_X) can be seen as a weighted dissimilarity space. In order to do this, we define

$$M_X(x,r) = \mu_X(B(x,r)).$$

We now define a locally bi-persistent category of weighted dissimilarity spaces. The idea is that, for $\varepsilon, \delta \in \mathbf{R}_+$, an (ε,δ)-morphism between weighted dissimilarity spaces is a function that doesn't increase the metric more than ε and that doesn't decrease the measure more than δ.

Definition 6.5.11. *Let* **wDiss** *be the locally bi-persistent category that has as objects all weighted dissimilarity spaces and has morphisms given as follows. Let* X *and* Y *be* *weighted dissimilarity spaces and let $\varepsilon, \delta \in \mathbf{R}_+$. Define*

$$\begin{aligned}\mathrm{Hom}_{\mathbf{wDiss}}(X,Y)_{(\varepsilon,\delta)} = \big\{ f : X \to Y \textit{ function of sets } : \\ \forall x, x' \in X, d_X(x,x') + \varepsilon \geq d_Y(f(x), f(x')) \\ \forall x \in X, \forall r \geq 0, M_X(x,r) \leq M_Y(f(x), r+\varepsilon) + \delta \big\},\end{aligned}$$

with composition and identities given by composition of functions and identity functions.

Example 6.5.12. Distance non-increasing and measure preserving maps $f : X \to Y$ between metric probability spaces are $(0,0)$-morphisms in the above locally bi-persistent structure.

Lemma 6.5.13. *The locally bi-persistent category* **wDiss** *is copowered and powered by representables.*

Proof. For $Y \in$ **wDiss** we let $Y^{(\varepsilon,\delta)}$ have the same underlying set as Y and distance and measure given by

$$\begin{aligned}d_{Y^{(\varepsilon,\delta)}}(y,y') &= \max\big(d_Y(y,y') - \varepsilon, 0\big), \\ M_{Y^{(\varepsilon,\delta)}}(y,r) &= M_Y(y, r+\varepsilon) + \delta\end{aligned}$$

respectively. For $X \in$ **wDiss** we let $(\varepsilon,\delta)\cdot X$ have the same underlying set as X, distance and measure given by

$$\begin{aligned}d_{(\varepsilon,\delta)\cdot X}(x,x') &= d_X(x,x') + \varepsilon, \\ M_{(\varepsilon,\delta)\cdot X}(x,r) &= \max\big(M'_X(x,r) - \delta, 0\big)\end{aligned}$$

respectively, where $M'_X(x,r) = M_X(x, r-\varepsilon)$ if $r - \varepsilon \geq 0$ and $M'_X(x,r) = 0$ otherwise. □

We define a relative locally persistent category structure on **wDiss** by letting the acyclic morphisms be the 0-morphisms that are surjective and that preserve both the measure and the metric.

Definition 6.5.14. *A* 0*-morphism between weighted dissimilarity spaces $f : X \to Y$ is an* ***acyclic morphism*** *if it is surjective, we have $d_X(x,x') = d_Y(f(x), f(x'))$ for every $x, x' \in X$, and we have $M_X(x,r) = M_Y(f(x), r)$ for every $x \in X$ and $r \in \mathbf{R}_+$.*

Lemma 6.5.15. *The underlying category of* **wDiss** *admits pullbacks, and acyclic morphisms are stable under pullback.*

Proof. Let $X \xrightarrow{f} Z \xleftarrow{g} Y$ be a cospan in the underlying category of **wDiss**. The pullback has as underlying set $P = \{(x,y) \in X \times Y : f(x) = g(y)\}$ and distance and measure given by

$$\begin{aligned} d_P((x,y),(x',y')) &= \max(d_X(x,x'), d_Y(y,y')), \\ M_P((x,y),r) &= \min(M_X(x,r), M_Y(y,r)). \end{aligned}$$

respectively. The universal property follows at once from the fact that the underlying set of P is the pullback of the corresponding cospan of the underlying sets. Note that the underlying dissimilarity space of P is the pullback of the corresponding cospan of dissimilarity spaces, as in the proof of Lemma 6.2.7.

Now, if g is surjective, distance preserving, and measure preserving, we have that $\pi_X : P \to X$ is also surjective and distance preserving, by Proposition 6.2.10. To see that it is also measure preserving, we compute

$$\begin{aligned} M_P((x,y),r) &= \min(M_X(x,r), M_Y(y,r)) \\ &= \min(M_X(x,r), M_Z(g(y),r)) \\ &= \min(M_X(x,r), M_Z(f(x),r)), \end{aligned}$$

where in the second equality we used the fact that g is measure preserving and in the third equality we used the fact that $(x,y) \in P$. Since f is measure non-decreasing, we have $M_P((x,y),r) = M_X(x,r)$, as required. □

The following is straightforward, using the characterization of powers (Lemma 6.5.13).

Lemma 6.5.16. *Let $\varepsilon, \delta \in \mathbf{R}_+$ and let $f : X \to Y$ be an acyclic morphism between weighted dissimilarity spaces. Then $f^{(\varepsilon,\delta)} : X^{(\varepsilon,\delta)} \to Y^{(\varepsilon,\delta)}$ is an acyclic morphism.* □

We let **wDiss** be the relative locally persistent category whose ε-morphisms are the $(\varepsilon,\varepsilon)$-morphisms in the above locally bi-persistent structure, and whose acyclic morphisms are the acyclic morphisms defined above.

Lemma 6.5.17. *Acyclic morphisms of* **wDiss** *are stable under weighted pullback.*

Proof. The locally persistent category **wDiss** admits weighted pullbacks □

We define a relative locally persistent category structure on $\mathbf{Diss}^{\mathbf{R}\times\mathbf{R}^{\mathrm{op}}}$ by letting it be a locally persistent category of persistent objects of **Diss**, and letting the acyclic morphisms be the natural transformations all of whose components are surjective and distance-preserving, as we did in Section 6.4.4. We choose $\vec{v} = (1,1,1)$. We refer to the quotient interleaving distance on $\mathbf{Diss}_{\vec{v}}^{\mathbf{R}\times\mathbf{R}^{\mathrm{op}}}$ as the Gromov–Hausdorff-interleaving distance.

From now on, we fix a kernel K. Thanks to the formula in Lemma 6.5.8, the definition of the kernel density filtration extends to dissimilarity spaces, since the integral in Lemma 6.5.8 is always defined for dissimilarity spaces, as it is an integral of a monotonic function. So we have a mapping $\mathsf{KDF} : \mathbf{wDiss} \to \mathbf{Diss}_{(1,1)}^{\mathbf{R}\times\mathbf{R}^{\mathrm{op}}}$ from weighted dissimilarity spaces to bi-persistent dissimilarity spaces given by

$$\mathsf{KDF}(X)(s,k) = \begin{cases} \{x \in X \,:\, (M_X * K_s)(x) \geq k\}, & \text{if } s,k>0 \\ X, & \text{if } k \leq 0 \\ \emptyset, & \text{if } s \leq 0 \text{ and } k>0. \end{cases}$$

where

$$(M_X * K_s)(x) := \int_0^\infty M_X\left(x, sK^{-1}(r)\right) \mathrm{d}r.$$

The following is the key lemma in proving that KDF is stable. In order to state it concisely, we need the following definition.

Definition 6.5.18. *A weighted dissimilarity space Y is* ***bounded*** *with constant M if there exists $M \geq 0$ such that $M_Y(y,r) \leq M$ for all $y \in Y$ and $r \geq 0$. The collection of all weighted dissimilarity spaces bounded with constant M is denoted by* $\mathbf{wDiss}_M$.

Example 6.5.19. The weighted dissimilarity space associated to any metric probability space is bounded with constant 1.

Lemma 6.5.20. *Let K be a kernel and let $r' \in (0, K(0))$. Let $f : X \to Y$ be an ε-morphism between weighted dissimilarity spaces and assume that Y is bounded with constant M. Let $x \in X$. Then*

$$(M_X * K_s)(x) \leq \left(M_Y * K_{s+\varepsilon_s}\right)(f(x)) + \varepsilon_k,$$

for $\varepsilon_s = \frac{\varepsilon}{K^{-1}(r')}$ and $\varepsilon_k = MK(0)\left(\frac{K(0)}{r'} - 1\right) + K(0)\varepsilon$.

Proof. We know that $(M_X * K_s)(x) = \int_0^{K(0)} M_X\left(x, sK^{-1}(r)\right) \mathrm{d}r$, since, if $r > K(0)$, then $K^{-1}(r) = 0$. Note that, by the assumption that f is an ε-morphism, we have that, for any radius $R \geq 0$,

$$M_X(x,R) \leq M_Y\left(f(x), R+\varepsilon\right) + \varepsilon$$

so we can bound the local density estimate of x as follows.

$$(M_X * K_s)(x) \leq \int_0^{K(0)} M_Y(f(x), sK^{-1}(r)+\varepsilon)+\varepsilon \, \mathrm{d}r = \int_0^{K(0)} M_Y(y, sK^{-1}(r)+\varepsilon) \, \mathrm{d}r + K(0)\varepsilon.$$

Since K^{-1} is non-increasing, and $r' < K(0)$, it follows that $K^{-1}(rr'/K(0)) \geq K^{-1}(r)$ for every $r \geq 0$. Moreover, for any $0 \leq r \leq K(0)$, we have $K^{-1}(rr'/K(0)) \geq K^{-1}(r')$. These two considerations imply that, for $0 \leq r \leq K(0)$, we have

$$sK^{-1}(r) + \varepsilon \leq \left(s + \varepsilon / K^{-1}(r')\right) K^{-1}\left(rr'/K(0)\right).$$

Combining this with the above bound for the local density estimate of x we get

$$\begin{aligned}
(M_X * K_s)(x) &\leq \int_0^{K(0)} M_Y\left(f(x), \left(s+\varepsilon/K^{-1}(r')\right) K^{-1}\left(rr'/K(0)\right)\right) \, \mathrm{d}r + K(0)\varepsilon \\
&= \frac{K(0)}{r'} \int_0^{r'} M_Y\left(f(x), \left(s+\varepsilon/K^{-1}(r')\right) K^{-1}(r)\right) \, \mathrm{d}r + K(0)\varepsilon \\
&\leq \frac{K(0)}{r'} \left(M_Y * K_{(s+\varepsilon/K^{-1}(r'))}\right)(f(x)) + K(0)\varepsilon.
\end{aligned}$$

Finally, note that, for $0 \leq a \leq M < \infty$ and $c \geq 1$, we have $ca \leq a + M(c-1)$. Since Y is bounded, there is $M \geq 0$ with $M_Y(y,r) \leq M$ for all $y \in Y$ and $r \in \mathbf{R}_+$. This implies that any local density estimate of Y is bounded by $M \cdot K(0)$. So we have that

$$(M_X * K_s)(x) \leq \left(M_Y * K_{(s+\varepsilon)/K^{-1}(r'))}\right)(f(x)) + MK(0)\left(\frac{K(0)}{r'} - 1\right) + K(0)\varepsilon,$$

as required. □

We now extend KDF to a functor.

Lemma 6.5.21. *Let $M \geq 0$. The mapping $\mathsf{KDF} : \mathbf{wDiss}_M \to \mathbf{Diss}^{\mathbf{R}\times\mathbf{R}^{\mathrm{op}}}$ extends to a functor between the underlying categories. Moreover, this functor maps acyclic morphisms to acyclic morphisms.*

Proof. Let $f : X \to Y$ be a 0-morphism between bounded weighted dissimilarity spaces. If we take $\varepsilon = 0$ in Lemma 6.5.20 and we let $r' \to K(0)$ we see that if $x \in X_{[s,k]}$

then $f(x) \in Y_{[s,k]}$. This gives us the functor in the statement.

We now prove that the functor maps acyclic morphisms to acyclic morphisms, concluding the proof. If $f : X \to Y$ is an acyclic morphism of weighted dissimilarity spaces, then the induced functions $X_{[s,k]} \to Y_{[s,k]}$ are distance preserving for every $s, k \geq 0$. Moreover, if $y \in Y_{[s,k]}$, let $x \in X$ such that $f(x) = y$. Since f is measure preserving, the local density estimate of x is equal to the local density estimate of y, so $x \in X_{[s,k]}$, and thus $X_{[s,k]} \to Y_{[s,k]}$ is also surjective. □

We can now prove the stability of KDF with respect to the quotient interleaving distances.

Theorem 6.5.22. *Let $M \geq 0$. The kernel density filtration* KDF : $\mathbf{wDiss}_M \to \mathbf{Diss}^{\mathbf{R}\times\mathbf{R}^{\mathrm{op}}}$ *is uniformly continuous with respect to the quotient interleaving distance and the Gromov–Hausdorff-interelaving distance. If* KDF *is constructed using the uniform kernel, then it is* 1*-Lipschitz.*

Proof. By Lemma 6.5.20, KDF maps ε-interleavings to ε'-interleavings, where

$$\varepsilon' = \max\left(\frac{\varepsilon}{K^{-1}(r')},\ MK(0)\left(\frac{K(0)}{r'} - 1\right) + K(0)\varepsilon\right).$$

Note that ε' can be made arbitrarily small by first taking r' sufficiently close to $K(0)$ and then choosing ε. So KDF is uniformly continuous with respect to the interleaving distances.

Moreover, KDF maps weakly equivalent objects to weakly equivalent objects, so it is uniformly continuous with respect to the quotient interleaving distances, by Theorem 4.2.3, using the fact that the acyclic morphisms of $\mathbf{Diss}^{\mathbf{R}\times\mathbf{R}^{\mathrm{op}}}$ are stable under weighted pullback.

For the last statement, note that, if K is the uniform kernel, then $K(0) = 1$ and $K^{-1} = K$. So, letting $r' \to 1$, we see that we can take $\varepsilon' = \varepsilon$. □

We now relate the quotient interleaving distance on **wDiss** to the Gromov–Hausdorff–Prokhorov distance. This is the last ingredient in the proof that KDF is stable.

Theorem 6.5.23. *Let X and Y be compact metric probability spaces. Then, we have*

$$\left(d_I^{\mathbf{wDiss}_{(1,1)}}\right)_{/\simeq}(X,Y) \leq 2 d_{GHP}(X,Y).$$

Proof. Assume that there is a metric space Z and embeddings $i : X \to Z$ and $j : Y \to Z$ such that

$$d_P(i_*(\mu_X), j_*(\mu_Y)) < \varepsilon \text{ and } d_H^Z(i(X), i(Y)) < \varepsilon.$$

Let $R \subseteq X \times Y$ be the correspondence such that $(x, y) \in R$ if and only if $d_Z(i(x), j(y)) < \varepsilon$. This is a correspondence since $d_H^Z(i(X), i(Y)) < \varepsilon$. Consider the weighted dissimilarity space R^X defined as follows. The underlying set of R^X is simply R. For $(x, y), (x', y') \in R^X$, we have $d_{R^X}((x, y), (x', y')) = d_X(x, x')$. Finally, for $(x, y) \in R^X$ and $r \geq 0$, we let $\mathscr{M}_{R^X}((x, y), r) = \mu_X(B(x, r))$. With this definitions, the map $\pi_X : R^X \to X$ is an acyclic morphism of **wDiss**.

Define R^Y in an analogous way. It is then enough to show that the identity map $R^X \to R^Y$ is part of a 2ε-interleaving. In order to show this, we must show that it doesn't increase the metric more than 2ε and that, for every $(x, y) \in R$ and every $r \geq 0$, we have

$$\mu_X(B(x, r)) \leq \mu_Y(B(y, r + 2\varepsilon)) + 2\varepsilon.$$

The first statement follows from the definition of the correspondence R. For the second one, we compute

$$\begin{aligned}
\mu_X(B(x, r)) &= i_*(\mu_X)(B(i(x), r)) \\
&\leq j_*(\mu_Y)(B(i(x), r)^{\varepsilon}) + \varepsilon \\
&\leq j_*(\mu_Y)(B(i(x), r + \varepsilon)) + \varepsilon \\
&\leq j_*(\mu_Y)(B(j(y), r + 2\varepsilon)) + \varepsilon \\
&= \mu_Y(B(y, r + 2\varepsilon)) + \varepsilon,
\end{aligned}$$

using the fact that $d_P(i_*(\mu_X), j_*(\mu_Y)) < \varepsilon$: □

The stability of KDF follows.

Proof of Theorem 6.5.1. This follows at once from Theorem 6.5.22 and Theorem 6.5.23. □

6.6 The correspondence-interleaving distance on hierarchical clusterings

In this section, we propose a distance to compare hierarchical clusterings, the correspondence-interleaving distance, and show that it can be interpreted as a quotient interleaving distance, which lets us deduce that the distance has good metric properties. We also explain in what way the correspondence-interleaving distance is a generalization of the distances considered in [CM10a] and [EBW15].

The main definitions of this section appear in [RS20] and are joint work with Alexander Rolle.

In Section 6.6.1, we define the notion of hierarchical clustering and the correspondence-interleaving distance as they appear in [RS20]. In Section 6.6.2, we show that this distance is a quotient interleaving distance and prove some metric properties of this distance, using the theory developed in this thesis. In Section 6.6.3, we explain how the correspondence-interleaving distance is a generalization of other distances between hierarchical clusterings that have been considered in the literature.

6.6.1 The correspondence-interleaving distance

Definition 6.6.1. *Let X be a set. A **clustering** of X is a subpartition of X, that is, a set of non-empty, disjoint subsets of X. The elements of a clustering are called **clusters**.*

Definition 6.6.2. *Let X be a set. The **poset of clusterings** of X, denoted $\mathsf{C}(X)$, is the poset whose elements are the clusterings of X, and where $S \preceq T \in \mathsf{C}(X)$ if, for each cluster $A \in S$, there is a (necessarily unique) cluster $B \in T$ such that $A \subseteq B$.*

Let $\mathbf{R}_{>0} = ((0,\infty), \leq)$.

Definition 6.6.3. *Let X be a set. A **covariant hierarchical clustering** of X is an order-preserving map $H : \mathbf{R}_{>0} \to \mathsf{C}(X)$. A **contravariant hierarchical clustering** of X is an order-preserving map $H : \mathbf{R}_{>0}^{\mathrm{op}} \to \mathsf{C}(X)$.*

An important motivating example is the following:

Example 6.6.4. If $f : \mathbb{R}^d \to \mathbb{R}$ is a probability density function, and $X = \mathrm{supp}(f)$, then there is a contravariant hierarchical clustering $H(f)$ of X, where, for $r > 0$, $H(f)(r)$ is the set of connected components of $\{x \in X : f(x) \geq r\}$.

A well-known covariant hierarchical clustering is given by the single-linkage hierarchical clustering algorithm:

Example 6.6.5. Let X be a metric space. We denote by $\mathsf{SL}(X)$ the **single-linkage** covariant hierarchical clustering of X, where, for $r > 0$, $\mathsf{SL}(X)(r)$ is the partition of X defined by the smallest equivalence relation $\sim_r$ on X with $x \sim_r y$ if $d_X(x, y) \le r$.

We can consider multi-parameter hierarchical clusterings that are covariant in some parameters, and contravariant in others.

Definition 6.6.6. *Let $n \ge 1$, and let $\vec{v} \in \{-1,1\}^{\times n}$. Let $\mathbf{R}^{\vec{v}}_{>0}$ be the product poset $\mathbf{R}^{\vec{v}}_{>0} = R_1 \times \cdots \times R_n$, where*

$$R_i = \begin{cases} \mathbf{R}_{>0} & \text{if } v_i = 1 \\ \mathbf{R}^{\mathsf{op}}_{>0} & \text{if } v_i = -1\,. \end{cases}$$

Definition 6.6.7. *Let X be a set, let $n \ge 1$, and let $\vec{v} \in \{-1,1\}^{\times n}$. A $\vec{v}$-**hierarchical clustering** of X is a map of posets $H : \mathbf{R}^{\vec{v}}_{>0} \to \mathsf{C}(X)$.*

Notation 6.6.8. Let $\vec{v} \in \{-1,1\}^{\times n}$. We write $\vec{\varepsilon} = (\varepsilon_1, \ldots, \varepsilon_n) \ge 0$ if $\varepsilon_i \ge 0$ for $1 \le i \le n$. For $\vec{r} = (r_1, \ldots, r_n) \in \mathbf{R}^{\vec{v}}_{>0}$, we write $\vec{r} + \vec{v}\vec{\varepsilon}$ for $(r_1 + v_1\varepsilon_1, \ldots, r_n + v_n\varepsilon_n)$ and we write $\vec{r} + \vec{v}\vec{\varepsilon} > 0$ if $r_i + v_i\varepsilon_i > 0$ for all $1 \le i \le n$.

Definition 6.6.9. *Let H and E be $\vec{v}$-hierarchical clusterings of a set X, and let $\vec{\varepsilon} \ge 0$. We say that H and E are $\vec{\varepsilon}$-**interleaved** if, for all $\vec{r} \in \mathbf{R}^{\vec{v}}_{>0}$ such that $\vec{r} + \vec{v}\vec{\varepsilon} > 0$, we have the following relations in $\mathsf{C}(X)$:*

$$H(\vec{r}) \preceq E(\vec{r} + \vec{v}\vec{\varepsilon}) \quad \text{and} \quad E(\vec{r}) \preceq H(\vec{r} + \vec{v}\vec{\varepsilon})\,.$$

As an example, we have the following stability result.

Proposition 6.6.10. *Let $f, g : \mathbb{R}^d \to \mathbb{R}_{\ge 0}$ be probability density functions with the same support. If $||f - g||_\infty < \varepsilon$, then $H(f)$ and $H(g)$ are ε-interleaved.*

Proof. For every $r \ge \varepsilon$, we have $\{f \ge r\} \subseteq \{g \ge r - \varepsilon\}$, and $\{g \ge r\} \subseteq \{f \ge r - \varepsilon\}$. This implies that, after taking connected components, every connected component of $\{f \ge r\}$ is included in a connected component of $\{g \ge r - \varepsilon\}$, and that every connected component of $\{g \ge r\}$ is included in a connected component of $\{f \ge r - \varepsilon\}$. □

To compare hierarchical clusterings of different sets, we use correspondences.

If $\psi : Y \to X$ is a function between sets, and $S = \{C_i\}_{i \in I}$ is a clustering of X, then $\psi^*(S) = \{\psi^{-1}(C_i)\}_{i \in I}$ is a clustering of Y. This defines a map of posets $\psi^* : \mathsf{C}(X) \to$

$\mathsf{C}(Y)$. So if H is a $\vec{v}$-hierarchical clustering of X, then there is a $\vec{v}$-hierarchical clustering $\psi^*(H)$ of Y, defined by the composition

$$\psi^*(H) : \mathbf{R}^{\vec{v}}_{>0} \xrightarrow{H} \mathsf{C}(X) \xrightarrow{\psi^*} \mathsf{C}(Y)\,.$$

Definition 6.6.11. *Let H and E be $\vec{v}$-hierarchical clustering of sets X and Y respectively, and let $R \subseteq X \times Y$ be a correspondence. Let $\vec{\varepsilon} \geq 0$. We say that H and E are $\vec{\varepsilon}$-**interleaved with respect to** R if $\pi_X^*(H)$ and $\pi_Y^*(E)$ are $\vec{\varepsilon}$-interleaved as $\vec{v}$-hierarchical clusterings of R.*

Definition 6.6.12. *Let H and E be $\vec{v}$-hierarchical clustering of sets X and Y respectively. Define the **correspondence-interleaving distance***

$$\begin{aligned} d_{CI}(H,E) = \inf\{\varepsilon \geq 0 :\ & \textit{there is a correspondence } R \subseteq X \times Y \\ & \textit{such that } H, E \textit{ are } (\varepsilon,\ldots,\varepsilon)\textit{-interleaved with respect to } R\}\,. \end{aligned}$$

6.6.2 The locally persistent category of hierarchical clusterings

We now interpret the correspondence-interleaving distance as a quotient interleaving distance and use this to prove some metric properties of this distance.

Let $n \geq 1$ and let $\vec{v} \in \{-1,1\}^{\times n}$. Let $\mathbf{hCl}^{\vec{v}}$ be the following locally $\mathbf{R}^n_+$-persistent category. An object of $\mathbf{hCl}^{\vec{v}}$ consist of a set X together with a $\vec{v}$-hierarchical clustering H_X on X. Given (X, H_X) and (Y, H_Y) objects of $\mathbf{hCl}^{\vec{v}}$, an $\vec{\varepsilon}$-morphism from (X, H_X) to (Y, H_Y), for $\vec{\varepsilon} \in \mathbf{R}^n_+$, consists of a function of sets $\psi : X \to Y$ such that, for every $\vec{r} \in \mathbf{R}^{\vec{v}}_{>0}$ such that $\vec{r} + \vec{v}\vec{\varepsilon} > 0$, we have that $H_X(\vec{r}) \leq \psi^*(H_Y)(\vec{r} + \vec{v}\vec{\varepsilon})$. Composition and identities are given by composition of functions and identity functions.

In order to get a locally persistent category, we perform a change of enrichment, as in Section 5.1. For simplicity, we choose the vector $(1,\ldots,1) \in \mathbf{R}^n_+$. So, from now on, $\mathbf{hCl}^{\vec{v}}$ denotes the locally persistent category where a δ-morphism is a $(\delta,\ldots,\delta)$-morphism in the $\mathbf{Set}^{\mathbf{R}^n_+}$-enrichment of $\mathbf{hCl}^{\vec{v}}$ described above.

An acyclic morphism $\psi : (X, H_X) \to (Y, H_Y)$ consists of a 0-morphism $\psi : X \to Y$ such that ψ is surjective, and such that $\psi^*(H_Y) = H_X$. This endows $\mathbf{hCl}^{\vec{v}}$ with the structure of a relative locally persistent category.

This relative locally persistent structure is well behaved.

Lemma 6.6.13. *Let $n \geq 1$ and let $\vec{v} \in \{-1,1\}^{\times n}$. The relative locally persistent category $\mathbf{hCl}^{\vec{v}}$ is copowered and powered by representables, and its underlying category admits pullbacks, binary products, and sequential limits. In particular, $\mathbf{hCl}^{\vec{v}}$ admits weighted pullbacks, weighted sequential limits, and terminal midpoints. Moreover, acyclic morphisms are closed under weighted pullbacks.*

We describe the construction of the limits, but we don't check the universal properties.

Proof. Let $X = (X, H_X) \in \mathbf{hCl}^{\vec{v}}$ and let $\varepsilon \in \mathbf{R}_+$. Let $X^\varepsilon = (X, H_X^\varepsilon)$ be the $\vec{v}$-hierarchical clustering with X as underlying set, and such that

$$H_X^\varepsilon(\vec{r}) = \begin{cases} \{X\} & \text{if } \vec{r} + \vec{v}\vec{\varepsilon} \not> 0 \\ H_X(\vec{r} + \vec{v}\vec{\varepsilon}) & \text{if } \vec{r} + \vec{v}\vec{\varepsilon} > 0. \end{cases}$$

By definition of the locally persistent structure of $\mathbf{hCl}^{\vec{v}}$, we have

$$\mathrm{Hom}_{\mathbf{hCl}^{\vec{v}}}(X, Y)_{\varepsilon+-} \cong \mathrm{Hom}_{\mathbf{hCl}^{\vec{v}}}(X, Y^\varepsilon),$$

isomorphism of functors, natural in X, Y, and ε, so $\mathbf{hCl}^{\vec{v}}$ admits powers by representables. Analogously, we define $\varepsilon \cdot X = (X, \varepsilon \cdot H_X)$ by

$$\varepsilon \cdot H_X(\vec{r}) = \begin{cases} \emptyset & \text{if } \vec{r} - \vec{v}\vec{\varepsilon} \not> 0 \\ H_X(\vec{r} - \vec{v}\vec{\varepsilon}) & \text{if } \vec{r} - \vec{v}\vec{\varepsilon} > 0, \end{cases}$$

which shows that $\mathbf{hCl}^{\vec{v}}$ is copowered by representables.

To define products and pullbacks, we first define two operations on hierarchical clusterings. Let X be a set and let $H, E : \mathbf{R}^{\vec{v}}_{>0} \to \mathsf{C}(X)$ be $\vec{v}$-hierarchical clusterings of X. The product of H and E, denoted by $H \times E : \mathbf{R}^{\vec{v}}_{>0} \to \mathsf{C}(X)$ is the $\vec{v}$-hierarchical clustering of X whose clusters at $\vec{r} \in \mathbf{R}^{\vec{v}}_{>0}$ are $\{C \times D : C \in H(\vec{r}), D \in E(\vec{r})\}$. If $Y \subseteq X$ is any subset, let $H|_Y : \mathbf{R}^{\vec{v}}_{>0} \to \mathsf{C}(Y)$ be the $\vec{v}$-hierarchical clustering of Y whose clusters at $\vec{r} \in \mathbf{R}^{\vec{v}}_{>0}$ are $\{C \cap Y : C \in H(\vec{r}) \text{ s.t. } C \cap Y \neq \emptyset\}$.

Let $X, Y \in \mathbf{hCl}^{\vec{v}}$. Let $X \times Y := (X \times Y, \pi_X^*(H_X) \times \pi_Y^*(H_Y))$. This is a product in the underlying category of $\mathbf{hCl}^{\vec{v}}$. Pullbacks are defined in a similar way. Given a cospan $X \to Z \leftarrow Y$ of $\vec{v}$-hierarchical clusterings, let P be the pullback of the cospan given by the underlying sets of X, Y, and Z, seen as a subset of $X \times Y$. Define a hierarchical clustering on P by $H_P = (H_{X \times Y})|_P$. This provides a pullback for the original cospan.

For sequential limits, we define another operation between hierarchical clusterings. Let $\{H_i\}$ be a family of $\vec{v}$-hierarchical clusterings of a fixed set X. Let $\bigcap_i H_i : \mathbf{R}^{\vec{v}}_{>0} \to \mathsf{C}(X)$ be the $\vec{v}$-hierarchical clustering of X whose clusters at $\vec{r} \in \mathbf{R}^{\vec{v}}_{>0}$ are $\{\bigcap_i C_i : C_i \in H_i(\vec{r}) \text{ s.t. } \bigcap_i C_i \neq \emptyset\}$.

Let

$$\cdots \to X_{i+1} \to X_i \to \cdots \to X_0$$

be a sequential diagram in the underlying category of $\mathbf{hCl}^{\vec{v}}$. Consider the hierarchical clustering (X, H_X) whose underlying set is the sequential limit of the sequential diagram of the underlying sets, and such that $H_X = \bigcap_i \pi^*_{X_i}(H_{X_i})$. This provides a sequential limit in the underlying category of $\mathbf{hCl}^{\vec{v}}$.

Using Proposition 3.2.12, Proposition 3.2.19, and Proposition 3.2.15 we conclude that $\mathbf{hCl}^{\vec{v}}$ admits weighted pullbacks, weighted sequential limits, and terminal midpoints

To check that acyclic morphisms are closed under weighted limits, it is enough to check that they are closed under powers by representables and under pullbacks in the underlying category, by Proposition 3.2.12. Both facts follow directly from the construction of powers and pullbacks. $\square$

As usual, we deduce the following.

Theorem 6.6.14. *Let $n \geq 1$ and let $\vec{v} \in \{-1,1\}^{\times n}$. The quotient interleaving distance on $\mathbf{hCl}^{\vec{v}}$ is intrinsic and complete and satisfies*

$$\begin{aligned}\left(d_I^{\mathbf{hCl}^{\vec{v}}}\right)_{/\simeq}(X,Y) &= \inf\{\delta : \exists X' \simeq X, Y' \simeq Y,\ X' \text{ and } Y' \text{ are } \delta\text{-interleaved}\}\\ &= \inf\{\delta : \exists \text{ acyclic morphisms } X' \to X \text{ and } Y' \to Y\\ &\qquad\qquad \text{such that } X' \text{ and } Y' \text{ are } \delta\text{-interleaved}\}.\end{aligned}$$

Proof. The facts that the distance is intrinsic and complete follow from Corollary 4.4.5 and Theorem 4.3.3, using Lemma 6.6.13 and noting that acyclic morphisms are clearly closed under sequential limits.

The description of the quotient interleaving distance follows from Theorem 4.1.4. $\square$

We conclude by showing that $\left(d_I^{\mathbf{hCl}^{\vec{v}}}\right)_{/\simeq}$ coincides with the correspondence-interleaving distance.

Theorem 6.6.15. *Let $n \geq 1$ and let $\vec{v} \in \{-1,1\}^{\times n}$. Then, for all $X, Y \in \mathbf{hCl}^{\vec{v}}$, we have*

$$\left(d_I^{\mathbf{hCl}^{\vec{v}}}\right)_{/\simeq}(X,Y) = d_{CI}(X,Y).$$

Proof. We use the second characterization of the quotient interleaving distance of Theorem 6.6.14.

Given a correspondence $R \subseteq X \times Y$ such that $\pi_X^*(H_X)$ and $\pi_Y^*(H_Y)$ are δ-interleaved, we have acyclic morphisms $\pi_X : (R, \pi_X^*(H_X)) \to (X, H_X)$ and $\pi_Y : (R, \pi_Y^*(H_Y)) \to (Y, H_Y)$, and a δ-interleaving between $(R, \pi_X^*(H_X))$ and $(R, \pi_Y^*(H_Y))$ given by the identity function of R. So $\left(d_I^{\mathbf{hCl}^{\vec{v}}}\right)_{/\simeq}(X,Y) \leq \delta$.

Conversely, given acyclic morphisms $\alpha : Z \to X$ and $\beta : W \to Y$ and a δ-interleaving between Z and W, we see that the interleaving gives us a bijection $\psi : Z \to W$ between the underlying sets, and thus we get a function $(\alpha, \beta \circ \psi) : Z \to X \times Y$, whose image is a correspondence $R \subseteq X \times Y$. This correspondence shows that $d_{CI}(X,Y) \leq \delta$, as needed. □

6.6.3 Comparison to previous distances on hierarchical clusterings

Using correspondences and interleaving distances to compare hierarchical clusterings has already been considered in the literature. We now briefly explain the relationship between the correspondence-interleaving distance and distances considered in [CM10a] and [EBW15].

In their work on the stability of the single-linkage hierarchical clustering algorithm [CM10b], Carlsson and Mémoli compare dendrograms (which are covariant hierarchical clusterings with some tameness conditions) (X_1, D_1) and (X_2, D_2) on finite sets X_1 and X_2 by associating to them ultra metric spaces (X_1, d_{D_1}) and (X_2, d_{D_2}) in a standard way and using the Gromov–Hausdorff distance to compare these ultra metric spaces. This was explained in Section 6.7.2. By unfolding the definitions, one sees that $d_{CI}(D_1, D_2) = 2\, d_{GH}\big((X_1, d_{D_1}), (X_2, d_{D_2})\big)$. So, the correspondence-interleaving distance is a generalization of the Carlsson-Mémoli distance on dendrograms.

The merge distortion metric of Eldridge, Belkin, and Wang [EBW15] is also closely related to the correspondence-interleaving distance.

Definition 6.6.16. *A **merge function** on a set X consists of a function $m : X \times X \to [0,\infty]$ that is symmetric and satisfies $m(x,z) \geq \min(m(x,y), m(y,z))$ for all $x,y,z \in X$.*

To every hierarchical clustering H on a set X one can assign a merge function η_H on X by letting $\eta_H(x,y) = \sup\{r > 0 \, : \, \exists C \in H(r), x, y \in C\}$ for every $x, y \in X$. On the other hand, given $m : X \times X \to [0,\infty]$ a merge function, and $r > 0$, let $X_{[r]} = \{x \in X \, : \, m(x,x) \geq r\} \subseteq X$, and define an equivalence relation on $X_{[r]}$ where $x \sim_{[r]} y$ if and only if $m(x,y) \geq r$. This defines a hierarchical clustering $H(m)$ of X where, for $r > 0$, $H(m)(r) = X_{[r]} \, / \sim_{[r]}$.

Definition 6.6.17. *Let X be a set. A* ***cluster tree*** *of X is given by a family $\mathcal{T}$ of subsets of X with the property that whenever A and B are distinct elements of $\mathcal{T}$, then one of the following is true: $A \cap B = \emptyset$, $A \subseteq B$, or $B \subseteq A$.*

Given sets X_1 and X_2 and a correspondence $R \subseteq X_1 \times X_2$, as well as cluster trees $\mathcal{T}_1$ and $\mathcal{T}_2$ on X_1 and X_2 respectively, together with functions $h_1 : X_1 \to \mathbb{R}_{\geq 0}$ and $h_2 : X_2 \to \mathbb{R}_{\geq 0}$, Eldridge, Belkin, and Wang define the *merge distortion* $d_R((\mathcal{T}_1, h_1), (\mathcal{T}_2, h_2))$. Furthermore, to each pair $(\mathcal{T}, h)$, they define an associated merge function [EBW15, Definition 6], which yields a hierarchical clustering $H(\mathcal{T}, h)$. By unrolling the definitions, one sees that

$$d_R((\mathcal{T}_1, h_1), (\mathcal{T}_2, h_2)) = \inf\{\varepsilon : H(\mathcal{T}_1, h_1) \text{ and } H(\mathcal{T}_2, h_2) \text{ are } \varepsilon\text{-interleaved w.r.t. } R\}.$$

So the infimum over all correspondences between X_1 and X_2 of the merge distortion with respect to that correspondence is equal to the correspondence-interleaving distance between the hierarchical clusterings $H(\mathcal{T}_1, h_1)$ and $H(\mathcal{T}_2, h_2)$.

6.7 The interleaving distance on persistent sets

In this section, we study the interleaving distance on the locally persistent category $\mathbf{Set}^{\mathbf{R}}$, and on a better behaved subcategory, which we define below. Our starting point is the following straightforward corollary of our theory.

Theorem 6.7.1. *The interleaving distance on the locally persistent category* $\mathbf{Set}^{\mathbf{R}}$ *is intrinsic and complete.*

Proof. The category **Set** is complete, so we just apply Corollary 4.4.4 and Corollary 4.3.2. □

Although an explicit counterexample may be difficult to construct, it doesn't seem to be the case that the interleaving distance on $\mathbf{Set}^{\mathbf{R}}$ is geodesic. In order to get a geodesic distance, we restrict our attention to a better behaved subcategory. Recall that, given $X \in \mathbf{Set}^{\mathbf{R}}$ and $r \leq s \in \mathbf{R}$, we let $\varphi^X_{r,s} : X(r) \to X(s)$ denote the structure morphism of X. Recall from Definition 4.5.4 that an object $X \in \mathbf{Set}^{\mathbf{R}}$ is q-tame if the image of $\varphi^X_{r,s}$ is a finite set whenever $r < s$. Recall from Definition 2.6.9 that an object $X \in \mathbf{Set}^{\mathbf{R}}$ is right continuous if the canonical function $X(r) \to \lim X_{>r}$ is a bijection. Here $X_{>r} : \{r' \in \mathbf{R} : r' > r\} \to \mathbf{Set}$ is the restriction of the functor X.

We consider the locally persistent category of q-tame, right continuous persistent sets, which we denote by $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$, and the locally persistent category of q-tame persistent sets $\mathbf{Set}^{\mathbf{R}}_{\mathsf{tame}}$. The main result that we prove is the following.

Theorem 6.7.2. *The interleaving distances on* $\mathbf{Set}^{\mathbf{R}}_{\mathsf{tame}}$ *and* $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$ *are geodesic and complete. Moreover, if* $X, Y \in \mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$ *and* $d_I(X,Y) = 0$*, then* X *and* Y *are isomorphic.*

The locally persistent category $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$ contains many interesting subcategories: the category of dendrograms in the sense of [CM10b] and the category of complete ultra metric spaces ([Ack13]) are both full subcategories of (the underlying category of) $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$.

The reason why the interleaving distance is geodesic when restricted to q-tame, right continuous persistent sets is that, in a sense, these persistent sets behave like compact metric spaces. In Section 6.7.2, we make this statement precise and discuss some connections between dendrograms and ultra metric spaces.

6.7.1 Interleaving distance on q-tame persistent sets

The following results are the main building blocks in showing that the interleaving distance on $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$ is geodesic and complete.

Proposition 6.7.3. *The locally persistent category* $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$ *is copowered and powered by representables, and admits binary products and pullbacks. In particular,* $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$ *admits terminal midpoints.*

Proof. The locally persistent category $\mathbf{Set}^{\mathbf{R}}$ is copowered and powered by representables and these are given by shifts (Example 3.2.7). Clearly, shifts preserve the properties of being q-tame and right continuous, so $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$ is copowered and powered by representables.

Since limits commute with limits, a product or pullback of right continuous persistent sets must be right continuous. Moreover, since limits in $\mathbf{Set}^{\mathbf{R}}$ are computed pointwise and finite limits of finite sets are finite sets, we conclude that a finite limit of q-tame objects of $\mathbf{Set}^{\mathbf{R}}$ must be q-tame.

The last claim then follows from Proposition 3.2.19. □

Proposition 6.7.4. *The locally persistent category* $\mathbf{Set}^{\mathbf{R}}_{\mathsf{righ,tame}}$ *admits weighted sequential limits of morphisms that are part of an interleaving.*

Proof. We know that $\mathbf{Set}^{\mathbf{R}}$ admits weighted sequential limits, by Corollary 4.3.2. We must show that if all the objects in a weighted sequential limit are q-tame and right continuous, and all the morphisms in the diagram are part of an interleaving, then the weighted sequential limit is also q-tame and right continuous.

By Proposition 3.2.15, the weighted sequential limit is computed as a (categorical) sequential limit of q-tame and right continuous objects. This implies that the limit is right continuous, by the fact that limits commute with limits. The fact that the limit is q-tame requires just a bit more work.

By Proposition 3.2.14, the weighted sequential limit of a diagram where each morphism is part of an interleaving is δ-interleaved with some object of the diagram for arbitrarily small $\delta \in \mathbf{R}_+$. In particular, the weighted sequential limit of the diagram is δ-interleaved with a q-tame persistent set for arbitrarily small δ. This implies that the limit must be q-tame too, since every non-identity structure map of the limit factors through a non-identity structure map of a q-tame persistent set. □

Our next goat is to prove that the interleaving distance on q-tame and right continuous persistent sets is geodesic. We will be using Theorem 4.5.14. We start by constructing an enrichment of $\mathbf{Set}^{\mathbf{R}}_{\text{righ,tame}}$ in persistent topological spaces.

For every $n \in \mathbb{N}$, let $\mathbf{n} = [-n, n] \subseteq \mathbf{R}$ denote the corresponding subinterval. Given $X \in \mathbf{Set}^{\mathbf{R}}$, consider $X|_{\mathbf{n}} : \mathbf{n} \to \mathbf{Set}$ the restriction to the interval $\mathbf{n}$. Each natural transformation in $\mathrm{Hom}_{\mathbf{Set}^{\mathbf{R}}}(X,Y)_\varepsilon = \mathrm{Nat}(X, Y^\varepsilon)$ induces a natural transformation $\mathrm{Nat}(X|_{\mathbf{n}}, Y^\varepsilon|_{\mathbf{n}})$, by restriction. Further restrictions let us form a sequential diagram of sets

$$\cdots \to \mathrm{Nat}(X|_{\mathbf{n+1}}, Y^\varepsilon|_{\mathbf{n+1}}) \to \mathrm{Nat}(X|_{\mathbf{n}}, Y^\varepsilon|_{\mathbf{n}}) \to \cdots \to \mathrm{Nat}(X|_{\mathbf{1}}, Y^\varepsilon|_{\mathbf{1}}) \to \mathrm{Nat}(X|_{\mathbf{0}}, Y^\varepsilon|_{\mathbf{0}}).$$

This construction is clearly natural in X, Y, and ε. Since a natural transformation is defined by its components, the limit of the above diagram is naturally isomorphic to the set $\mathrm{Hom}_{\mathbf{Set}^{\mathbf{R}}}(X,Y)_\varepsilon$. Let the topology of $\mathrm{Hom}_{\mathbf{Set}^{\mathbf{R}}}(X,Y)_\varepsilon$ be the sequential limit topology, where each set in the sequential diagram is given the discrete topology. By naturality, this provides an enrichment in $\mathbf{Top}^{\mathbf{R}_+}$ for the locally persistent category $\mathbf{Set}^{\mathbf{R}}$, and thus for the locally persistent category $\mathbf{Set}^{\mathbf{R}}_{\text{righ,tame}}$.

Theorem 6.7.5. *The interleaving distance on q-tame, right continuous persistent sets is complete and geodesic. Moreover, if two q-tame and right continuous persistent sets are at interleaving distance* 0*, then they are isomorphic.*

Proof. Completeness follows from Proposition 6.7.4 and Theorem 4.3.1. Assuming that the interleaving distance on $\mathbf{Set}^{\mathbf{R}}_{\text{righ,tame}}$ reflects interleavings (Definition 4.5.1), the fact that this distance is geodesic follows from Theorem 4.5.2 and Proposition 6.7.3, and the second claim follows immediately. So it remains to be shown that the interleaving distance of $\mathbf{Set}^{\mathbf{R}}_{\text{righ,tame}}$ reflects interleavings. In order to do this, we use Theorem 4.5.14.

Consider the $\mathbf{Top}^{\mathbf{R}_+}$-enrichment of $\mathbf{Set}^{\mathbf{R}}_{\text{righ,tame}}$ constructed above. Applying Theorem 4.5.14, it is enough to show that $\mathrm{Hom}_{\mathbf{Set}^{\mathbf{R}}}(X,Y)$ is a q-tame, right continuous persistent topological space, whenever X and Y are q-tame, right continuous persistent sets, and that $\mathrm{Hom}_{\mathbf{Set}^{\mathbf{R}}}(X,Y)_\varepsilon$ is T_1 for every ε.

The fact that $\mathrm{Hom}_{\mathbf{Set}^{\mathbf{R}}}(X,Y)_\varepsilon$ is T_1 is clear, since its topology is a sequential limit of discrete topologies. Right continuity of $\mathrm{Hom}_{\mathbf{Set}^{\mathbf{R}}}(X,Y)$ follows directly from Proposition 2.6.10.

We conclude the proof by showing that $\mathrm{Hom}_{\mathbf{Set}^{\mathbf{R}}}(X,Y)$ is a q-tame persistent topological space whenever X and Y are q-tame persistent sets. Let $\varepsilon < \delta$, we must show that the image of the map

$$\mathrm{Nat}(X, Y^{\varepsilon}) \to \mathrm{Nat}(X, Y^{\delta})$$

given by postcomposition with $Y^{\varepsilon} \to Y^{\delta}$ is a compact set. We will show that it is a sequential limit of finite sets. The image of $\mathrm{Nat}(X, Y^{\varepsilon}) \to \mathrm{Nat}(X, Y^{\delta})$ is isomorphic to the sequential limit of the images of the maps $\mathrm{Nat}(X, Y^{\varepsilon}) \to \mathrm{Nat}(X|_{\mathbf{n}}, Y^{\delta}|_{\mathbf{n}})$ for $n \in \mathbb{N}$. It is then enough to show that the images of these maps are finite sets.

Fix $n \in \mathbb{N}$ and consider a natural transformation $f \in \mathrm{Nat}(X, Y^{\varepsilon})$. The image of f in $\mathrm{Nat}(X|_{\mathbf{n}}, Y^{\delta}|_{\mathbf{n}})$ is given by the natural transformation whose component $r \in \mathbf{n}$ is given by

$$\varphi^{Y}_{r+\varepsilon, r+\delta} \circ f_r = f_{r+\delta-\varepsilon} \circ \varphi^{X}_{r, r+\delta-\varepsilon}. \tag{6.7.6}$$

Denote the image of f by g. We will show that there are finitely many possible natural transformations $g \in \mathrm{Nat}(X|_{\mathbf{n}}, Y^{\delta}|_{\mathbf{n}})$ whose components are of the form given in Eq. (6.7.6) for some $f \in \mathrm{Nat}(X, Y^{\varepsilon})$.

Again, fix $n \in \mathbb{N}$ and consider a natural transformation $f \in \mathrm{Nat}(X, Y^{\varepsilon})$. Let $\gamma = (\delta - \varepsilon)/2$ and let $M \in \mathbb{N}$ be the smallest natural number such that $-n + \gamma(M-2) \geq n$. For each $0 \leq i \leq M$, consider the map $f_{-n+\gamma i} : X(-n+\gamma i) \to Y(-n+\gamma i+\varepsilon)$. Since f is a natural transformation, for $r \in [-n, n]$, if we let $0 \leq i \leq M$ be the smallest natural number such that $-n + \gamma i \geq r$, we have that the r-component $g_r : X(r) \to Y(r+\delta)$ of the corresponding natural transformation $g \in \mathrm{Nat}(X|_{\mathbf{n}}, Y^{\delta}|_{\mathbf{n}})$ is equal to

$$\varphi^{Y}_{-n+\gamma(i+1)+\varepsilon, r+\delta} \circ f_{-n+\gamma(i+1)} \circ \varphi^{X}_{-n+\gamma i, -n+\gamma(i+1)} \circ \varphi^{X}_{r, -n+\gamma i}.$$

This shows that the natural transformation g is completely determined by the functions

$$f'_i = f_{-n+\gamma(i+1)} \circ \varphi^{X}_{-n+\gamma i, -n+\gamma(i+1)} : X(-n+\gamma i) \to Y(-n+\gamma(i+1)+\varepsilon) \tag{6.7.7}$$

for $0 \leq i \leq M$. Since there are finitely many i such that $0 \leq i \leq M$ is then enough to show that there are only finitely many possible functions f'_i of the form given in Eq. (6.7.7) for some $f \in \mathrm{Nat}(X, Y^{\varepsilon})$.

Fix $n \in \mathbb{N}$ and consider a natural transformation $f \in \mathrm{Nat}(X, Y^{\varepsilon})$. Fix i such that $0 \leq i \leq M$. The function f'_i defined in Eq. (6.7.7) is completely determined by the value of $f_{-n+\gamma(i+1)}$ on the image of $\varphi^{X}_{-n+\gamma i, -n+\gamma(i+1)}$, which is a finite set by the tameness of

X. Moreover, by naturality of f, we have

$$f'_i = \varphi^Y_{-n+\gamma i+\varepsilon, -n+\gamma(i+1)+\varepsilon} \circ f_{-n+\gamma i},$$

so the image of f'_i is finite, by tameness of Y. Together, these last two facts say that f'_i is completely determined by assigning each of the finitely many elements of the image of $\varphi^X_{-n+\gamma i, -n+\gamma(i+1)}$ to one of the finitely many elements of the image of $\varphi^Y_{-n+\gamma i+\varepsilon, -n+\gamma(i+1)+\varepsilon}$, so there are finitely many possible functions f'_i of the form given in Eq. (6.7.7), concluding the proof. □

Proof of Theorem 6.7.2. By Theorem 6.7.5, it is enough to show that the inclusion of $\mathbf{Set}^{\mathbf{R}}_{\text{righ,tame}}$ into $\mathbf{Set}^{\mathbf{R}}_{\text{tame}}$ induces an isometry with respect to the interleaving distances.

To see this, note that, by Proposition 2.6.12, for any $X \in \mathbf{Set}^{\mathbf{R}}_{\text{tame}}$, we have that $d_I(X, X^\#) = 0$. Since $X^\#$ is right continuous, it is enough to show that $X^\#$ is q-tame as well. This is true since X and $X^\#$ are ε-interleaved for arbitrarily small $\varepsilon > 0$. □

6.7.2 Persistent sets and ultra metric spaces

In this section, we give precise meaning to the statement that q-tame and right continuous persistent sets are like compact metric spaces. We start with a bit of context.

In [CM10b], the stability of hierarchical clustering algorithms is studied. The input of a hierarchical clustering algorithm is taken to be a finite metric space X and the output is taken to be a dendrogram on X.

Definition 6.7.8. *A **dendrogram** on a finite set X is given by a function $\Theta : [0,\infty) \to \mathsf{partitions}(X)$ such that:*

- $\Theta(0)$ *is the discrete partition of X;*
- *there is t_0 such that $\Theta(t_0)$ is the codiscrete partition of X;*
- *if $r \leq s$, then $\Theta(r)$ refines $\Theta(s)$;*
- *for every r there is an $\varepsilon > 0$ such that $\Theta(r) = \Theta(r+\varepsilon)$.*

In order to get a metric between dendrograms, they embed the collection of dendrograms into the category of ultra metric spaces and use the Gromov–Hausdorff distance between metric spaces.

Definition 6.7.9. *An ultra ep metric d_P on a set P is an ep metric such that for every $p, p', p'' \in P$ we have $d_P(p, p'') \leq \max(d_P(p, p'), d_P(p, p''))$.*

It is easy to see that, given a dendrogram Θ on a finite set X, the following defines an ultra metric on X:

$$d_\Theta(x, x') = \inf\left\{ t \in [0, \infty) : \ x \text{ and } x' \text{ belong to the same equivalence class of } \Theta(t) \right\}.$$

For any finite set X, this construction gives a bijection between ultra metrics on X and dendrograms on X ([CM10b, Theorem 9]).

Now, a dendrogram Θ on X also gives rise to a persistent set $\Theta : \mathbf{R}_+ \to \mathbf{Set}$, as follows. Note that $\Theta : [0, \infty) \to \mathsf{partitions}(X)$ can be regarded as a persistent partition of X, but a partition of X is just a set of subsets of X, so there is a forgetful functor $\mathsf{partitions}(X) \to \mathbf{Set}$ that forgets that the subsets happen to be subsets of X. Then, as a persistent set, $\Theta : \mathbf{R}_+ \to \mathbf{Set}$ is just the composite of Θ with this forgetful functor.

This construction embeds the collection of all dendrograms of finite sets into the category of persistent sets. With not much more work, one can see that, combining this construction with a generalization of the construction above that allows for possibly infinite ultra metric spaces, one can embed the category of ultra metric spaces (with distance non-increasing maps between them) into the category of persistent sets. This result is not new, and in fact this perspective allows one to generalize ultra metric spaces to Γ-valued ultra metrics, for Γ a complete lattice ([PCR96], [PCR97]).

We now explain a stronger connection between persistent sets and ultra metric spaces, established in [Ack13], and we strengthen this connection.

An **extended ultra metric space** is an ep ultra metric space such that only equal points are at distance 0 from each other. Recall that $X \in \mathbf{Set}^{\overline{\mathbf{R}}_+}$ is right continuous if, for every $r \in \overline{\mathbf{R}}_+$, the natural map $X(r) \to \lim_{r' > r} X(r')$ is a bijection. We say that $X \in \mathbf{Set}^{\overline{\mathbf{R}}_+}$ is **separated** if the natural map $X(r) \to \lim_{r' > r} X(r')$ is an injection. Of course, every right continuous persistent set is separated. Finally, we say that $X \in \mathbf{Set}^{\overline{\mathbf{R}}_+}$ is **flabby** if all of its structure morphisms are surjective.

In [Ack13], the following equivalence of categories is proven. We state the theorem for the lattice $\Gamma = \overline{\mathbf{R}}_+$, but the theorem in the paper works for any complete lattice.

Theorem 6.7.10 ([Ack13]). *There is an equivalence of categories between the category of extended ultra metric spaces and distance non-increasing maps and the category of flabby and separated objects of $\mathbf{Set}^{\overline{\mathbf{R}}_+}$. The equivalence is given by mapping X to $\mathsf{SL}(X)$, the single-linkage clustering of X, taken as a persistent set.*

Moreover, an extended ultra metric space is complete if and only if its corresponding persistent set is right continuous. □

We give a few remarks. Firstly, the notion of completeness in the theorem is the one of Definition 2.2.9. Secondly, although the theorem is not stated in exactly the same language as in [Ack13], it is straightforward to do the translation by recalling that, as discussed in Remark 2.6.14, right continuous persistent objects are sheaves for the canonical topology associated to the frame $\overline{\mathbf{R}}_+$. Thirdly, the fact that the equivalence is given by the single-linkage construction is evident from [Ack13, Definition 3.13]. Finally, as explained above, the single-linkage of an ultra metric space X is the persistent set $\mathsf{SL}(X) : \overline{\mathbf{R}}_+ \to \mathbf{Set}$ such that $\mathsf{SL}(X)(r) = X/\sim_r$ where $x \sim_r y$ if and only if $d_X(x,y) \leq r$.

We can strengthen the above theorem further by classifying the compact extended ultra metric spaces, using q-tameness. This result gives some insight into why q-tameness allows us to extract a coherent family out of a non-coherent one (Proposition 4.5.9). The notion of compactness in the following result is the one of Definition 2.2.11.

Theorem 6.7.11. *Under the correspondence of Theorem 6.7.10, an extended ultra metric space is totally bounded if and only if its corresponding persistent set is q-tame. In particular, an extended ultra metric space is compact if and only if its corresponding persistent set is q-tame and right continuous.*

Proof. Let us start with the first statement. Note that a flabby persistent set Y is q-tame if and only if, for every $r > 0$, we have that $Y(r)$ is finite. So let X be an extended ultra metric space. On the one hand, if X is totally bounded, given $r > 0$, we can cover X with finitely many open balls of radius r, so $\mathsf{SL}(X)(r)$ has finite cardinality. On the other hand, if for $r > 0$ we have that $\mathsf{SL}(X)(r)$ has finite cardinality, choose finitely many x_i such that each equivalence class of $\mathsf{SL}(X)(r)$ is represented by one of the points x_i. Since the metric of X is ultra, it follows that $X \subseteq \bigcup_i B(x_i, \varepsilon)$, as required.

The second claim follows at once from the first one and Lemma 2.2.12. □

Informally, we conclude that right continuity is like completeness, and q-tameness is like total boundedness.

6.8 Mémoli's $d_{\mathscr{F}}$ distance on finite filtered simplicial complexes

For X a set, let $\mathscr{P}(X)$ denote the set of subsets of X. A **finite filtered simplicial complex** (X, F_X) consists of a finite set X together with a function $F_X : \mathscr{P}(X) \to \mathbf{R}$ that respects inclusions. That is, if $\sigma \subset \tau \in \mathscr{P}(X)$, then $F_X(\sigma) \leq F_X(\tau)$.

Given a finite filtered simplicial complex (X, F_X) and a surjective function $f : Z \to X$ from a finite set Z, we get an induced finite filtered simplicial complex $(Z, f^*(F_X))$, where $f^*(F_X)(\sigma) = F_X(f(\sigma))$ for every subset $\sigma \subseteq Z$.

In [Mé17], the following ep metric between finite filtered simplicial complexes is defined:

$$d_{\mathscr{F}}(X, Y) = \inf \left\{ \max_{\sigma \subseteq Z} \left| f^*(F_X)(\sigma) - g^*(F_Y)(\sigma) \right| \right\},$$

where the infimum is taken over all finite sets Z and surjective functions $f : Z \to X$ and $g : Z \to Y$.

It is claimed in [Mé17] that this distance is geodesic. This is justified by constructing a path of length at most δ between X and Y, given a finite set Z and surjective functions $f : Z \to X$ and $g : Z \to Y$ such that $\max_{\sigma \subseteq Z} \left| f^*(F_X)(\sigma) - g^*(F_Y)(\sigma) \right| \leq \delta$. We note that it is not explicitly justified why, if $d_{\mathscr{F}}(X, Y) = \delta$, then there exist Z, f, and g as above (in the language of [Mé17], why there is a minimizing tripod). Nonetheless, it is not too hard to fill this gap.

In this section, we give a relative locally persistent category structure to the collection of finite filtered simplicial complexes and we show that its quotient interleaving distance is geodesic and that if two elements are at distance 0, then they are weakly equivalent. We also show that this distance coincides with $d_{\mathscr{F}}$, thus providing a proof that $d_{\mathscr{F}}$ is geodesic. It is interesting to note that the paths we construct for the geodesics, which are obtained from general arguments (Corollary 6.8.3), are not the same as the paths given in [Mé17, Section 6.1].

Finally, we show that Mémoli's distance and the homotopy interleaving distance induce different metrics on the collection of finite filtered simplicial complexes (Remark 6.8.9).

The collection of finite filtered simplicial complexes can be endowed with a locally persistent category structure where an ε-morphism from (X, F_X) to (Y, F_Y) is given by a function $f : X \to Y$ such that, for every subset $\sigma \subseteq X$, we have $F_X(\sigma) \leq F_Y(f(\sigma)) + \varepsilon$. We denote this locally persistent category by **ffsCpx**.

An **acyclic morphism** $f : (X, F_X) \to (Y, F_Y)$ is given by a 0-morphism such that $f : X \to Y$ is surjective and such that $F_X(\sigma) = F_Y(f(\sigma))$.

In **ffsCpx**, it is easy to construct copowers and powers by representables.

Lemma 6.8.1. *The locally persistent category* **ffsCpx** *is copowered and powered by representables.*

Proof. Let $\varepsilon \in \mathbf{R}_+$ and let $X, Y \in$ **ffsCpx**. Let the underlying set of $\varepsilon \cdot X \in$ **ffsCpx** be X and let $F_{\varepsilon \cdot X}(\sigma) = F_X(\sigma) - \varepsilon$ for every $\sigma \subseteq X$. Similarly, let the underlying set of $Y^\varepsilon \in$ **ffsCpx** be Y and let $F_{Y^\varepsilon}(\tau) = F_Y(\tau) + \varepsilon$ for every $\tau \subseteq Y$. It is then clear that, for every $\delta \in \mathbf{R}_+$, we have

$$\mathrm{Hom}_{\mathbf{ffsCpx}}(X, Y)_{\varepsilon+\delta} \cong \mathrm{Hom}_{\mathbf{ffsCpx}}(X, Y^\varepsilon)_\delta \cong \mathrm{Hom}_{\mathbf{ffsCpx}}(\varepsilon \cdot X, Y)_\delta,$$

and that these isomorphisms are natural in X, Y, δ, and ε. □

We can also construct binary products and pullbacks.

Lemma 6.8.2. *The underlying category of* **ffsCpx** *admits binary products and pullbacks.*

Proof. Given $X, Y \in$ **ffsCpx**, their product $X \times Y \in$ **ffsCpx** has as underlying set the product of the underlying sets of X and Y, and $F_{X\times Y}(\sigma) = \max(F_X(\pi_X(\sigma)), F_Y(\pi_Y(\sigma)))$ for every $\sigma \subseteq X \times Y$. The fact that this is a categorical product is clear, as the universal property can be verified directly.

Pullbacks are similar. Let $X \to Z \leftarrow Y$ be a cospan in **ffsCpx**. Define $P \in$ **ffsCpx** with underlying set the pullback of the cospan formed by the underlying sets of X, Z, and Y, and $F_P(\sigma) = \max(F_X(\pi_X(\sigma)), F_Y(\pi_Y(\sigma)))$, for every $\sigma \subseteq P$. The universal property is easy to verify. □

Proposition 3.2.12 and Proposition 3.2.19 then allow us to conclude the following.

Corollary 6.8.3. *The locally persistent category* **ffsCpx** *admits weighted pullbacks and terminal midpoints.* □

Acyclic morphisms behave well with respect to weighted pullbacks, as the following result shows.

Lemma 6.8.4. *Acyclic morphisms in* **ffsCpx** *are stable under weighted pullback.*

Proof. By Lemma 4.1.5, it is enough to show that acyclic morphisms are stable under pullbacks and under powering by representables, and this is clear by the construction of pullbacks and powers. □

This allows us to characterize the quotient interleaving distance, prove that it is geodesic, and prove that the quotient interleaving distance coincides with $d_{\mathscr{F}}$.

Theorem 6.8.5. *For $X, Y \in$ **ffsCpx** we have*

$$\begin{aligned}\left(d_I^{\mathbf{ffsCpx}}\right)_{/\simeq}(X,Y) &= \inf\left\{\delta : \exists X' \simeq X, Y' \simeq Y,\ X' \text{ and } Y' \text{ are } \delta\text{-interleaved}\right\}\\ &= \inf\{\delta : \exists \text{ acyclic morphisms } X' \to X \text{ and } Y' \to Y\\ &\qquad\qquad \text{such that } X' \text{ and } Y' \text{ are } \delta\text{-interleaved}\}.\end{aligned}$$

Proof. This follows directly from Theorem 4.1.4, using Lemma 6.8.4. □

Lemma 6.8.6. *The distance $(d_I^{\mathbf{ffsCpx}})_{/\simeq}$ reflects quotient interleavings.*

Proof. We apply Theorem 4.5.18, so we must check that, for $X, X' \in$ **ffsCpx**, the persistent set of quotient interleavings $\mathbb{QI}(X, X') : \mathbf{R}_+ \to \mathbf{SET}$ is right continuous, and that we can lift it to a q-tame persistent topological space with closed structure maps.

Let us instantiate the definition of the persistent set of quotient interleavings to this case:

$$\begin{aligned}\mathbb{QI}(X,Y)_\delta = \Big\{(Z, Z', u, v, f, g) : {} & Z, Z' \in \mathbf{epMet}_c,\\ & u : Z \to X, v : Z' \to X' \text{ are acyclic morphisms,}\\ & f \text{ and } g \text{ form a } \delta\text{-interleaving between } X \text{ and } X' \Big\}.\end{aligned}$$

The structure morphisms are just inclusions in this case, since, by definition of the locally persistent category structure of **ffsCpx**, the functions f and g are inverse bijections between the underlying sets of Z and Z' such that, for every $\sigma \subseteq Z$, we have $|F_Z(\sigma) - F_{Z'}(f(\sigma))| \le \delta$.

We first prove that $\mathbb{QI}(X, X')$ is right continuous. This follows from the fact that the structure morphisms are the natural inclusions, and, if a pair of inverse bijections f and g between Z and Z' satisfy $|F_Z(\sigma) - F_{Z'}(f(\sigma))| \le \delta'$ for every $\delta' > \delta$, then they satisfy the analogous condition for δ.

To lift $\mathbb{QI}(X, X')$ to a persistent topological space, we map into a simpler persistent topological space and pull back the topology. Let $\mathbb{QI}_b(X, X')_\delta$ be the subset of

$\mathbb{QI}(X,X')_\delta$ of δ-quotient interleavings $X \leftarrow Z \leftrightarrow Z' \rightarrow X'$ such that the underlying subset of Z is equal to the underlying subset of Z', and both are a subset of $X \times X'$. For each $\delta \in \mathbf{R}_+$, we can endow $\mathbb{QI}_b(X,X')_\delta$ with the discrete topology, which makes $\mathbb{QI}_b(X,X') : \mathbf{R}_+ \to \mathbf{Top}$ into a q-tame persistent topological space with closed structure morphisms, since this persistent set takes values in finite sets.

Given a δ-quotient interleaving $X \xleftarrow{u} Z \xrightarrow{f} Z' \xrightarrow{v} X'$ we have a set map $(u, v\circ f)$: $Z \to X \times X'$, and its image gives us a subset $S \subseteq X \times X'$. Together with the projections to X and X', we get a diagram of sets $X \xleftarrow{\pi_X} S \xrightarrow{\mathrm{id}} S \xrightarrow{\pi_{X'}} X'$. Now, we can use π_X and $\pi_{X'}$ to endow S with two structures of finite filtered simplicial complex, by pulling back the structures from X and X' respectively, and thus obtain a diagram of finite filtered simplicial complexes $X \xleftarrow{\pi_X} S \xrightarrow{h} S' \xrightarrow{\pi_{X'}} X'$. By construction, h is part of a δ-interleaving and π_X and $\pi_{X'}$ are acyclic morphisms of finite filtered simplicial complexes, so $X \xleftarrow{\pi_X} S \xrightarrow{h} S' \xrightarrow{\pi_{X'}} X'$ is an element of $\mathbb{QI}_b(X,X')_\delta$. This provides us with a natural transformation $\mathbb{QI}(X,X') \Rightarrow \mathbb{QI}_b(X,X')$, where naturality follows from the fact that the structure morphisms of both persistent sets are the natural inclusions.

The components of the above natural transformation are surjective since the composite $\mathbb{QI}_b(X,X') \Rightarrow \mathbb{QI}(X,X') \Rightarrow \mathbb{QI}_b(X,X')$ is the identity. We can then pull back the topology on $\mathbb{QI}_b(X,X')$ to get a persistent topological space structure $\mathbb{QI}(X,X')$: $\mathbf{R}_+ \to \mathbf{TOP}$ that is q-tame and such that all of the structure maps are closed, concluding the proof. □

We can now prove that $(d_I^{\mathbf{ffsCpx}})_{/\simeq}$ is geodesic.

Theorem 6.8.7. *The distance $(d_I^{\mathbf{ffsCpx}})_{/\simeq}(X,Y)$ is geodesic and if $(d_I^{\mathbf{ffsCpx}})_{/\simeq}(X,Y) = 0$ then $X \simeq Y$.*

Proof. This follows from Theorem 4.5.16, using Lemma 6.8.6 and Corollary 6.8.3 to satisfy the hypotheses. □

We now show that $(d_I^{\mathbf{ffsCpx}})_{/\simeq}$ coincides with the distance $d_{\mathscr{F}}$ presented in [Mé17].

Proposition 6.8.8. *We have $(d_I^{\mathbf{ffsCpx}})_{/\simeq}(X,Y) = d_{\mathscr{F}}(X,Y)$.*

We use the same methodology as in the proof of Theorem 6.2.12.

Proof. We use the second characterization of the quotient interleaving distance of Theorem 6.8.5.

Given $f : Z \to X$ and $g : Z \to Y$ such that

$$\max_{\sigma \subseteq Z} \left| f^*(F_X)(\sigma) - g^*(F_Y)(\sigma) \right| \leq \delta$$

we have $(Z, f^*(F_X)) \simeq X$ and $(Z, g^*(F_Y)) \simeq Y$, and $(Z, f^*(F_X))$ and $(Z, g^*(F_Y))$ δ-interleaved. So $(d_I^{\mathbf{ffsCpx}})_{/\simeq}(X,Y) \leq d_{\mathscr{F}}(X,Y)$.

Going the other way, given $\alpha : X' \to X$ and $\beta : Y' \to Y$ acyclic morphisms such that X' and Y' are δ-interleaved, we let Z be the underlying set of X', which is in bijection with the underlying set of Y' under a bijection $\gamma : X' \to Y'$ that represents the interleaving between Y' and X'. We moreover define functions $Z \to X$ and $Z \to Y$ by α and $\beta \circ \gamma$ respectively. It follows that $d_{\mathscr{F}}(X,Y) \leq (d_I^{\mathbf{ffsCpx}})_{/\simeq}(X,Y)$, as required. □

We conclude this section by proving that Mémoli's distance does not coincide with the homotopy interleaving distance in general.

Remark 6.8.9. There is a locally persistent functor

$$\mathsf{R} : \mathbf{ffsCpx} \to \mathbf{Top}^{\mathbf{R}}$$

given by applying geometric realization. This locally persistent functor maps acyclic morphisms to weak equivalences. To see this, we apply Lemma 6.3.3. This implies that, for $X, Y \in \mathbf{ffsCpx}$, we have

$$d_{\mathscr{F}}(X,Y) \geq d_{HI}(\mathsf{R}(X), \mathsf{R}(Y)).$$

We now show that, in general, the distances do not agree on finite filtered simplicial complexes. Consider, on the one hand, the set $X = \{a, b\}$. Let $r \geq 0$ and consider the filtration of the simplicial complex $\{\{a\},\{b\},\{a,b\}\}$ given by $F_X(\{a\}) = 0$, $F_X(\{b\}) = r$, and $F_X(\{a,b\}) = r$. Consider, on the other hand, the set $Y = \{c\}$ and the filtration of the simplicial complex $\{\{c\}\}$ given by $F_Y(\{c\}) = 0$. Both these filtered simplicial complexes are empty before 0 and pointwise contractible after 0, and thus $d_{HI}(\mathsf{R}(X), \mathsf{R}(Y)) = 0$.

Since Y is a singleton, for any set Z there is exactly one set map $Z \to Y$. By inspection, this fact implies that $d_{\mathscr{F}}(X,Y) = r$. Since r is arbitrary, we see that $d_{\mathscr{F}}$ can be arbitrarily larger than d_{HI}, and thus that $d_{\mathscr{F}}$ and d_{HI} are non-equivalent metrics on the collection of finite filtered simplicial complexes.

6.9 Towards distances on persistence diagrams as interleaving distances

In this short section, we recover the Wasserstein distances on persistence diagrams as interleaving distances of locally persistent categories. We conclude the section by outlining future work in this direction.

Definition 6.9.1. *Let $q \in [1,\infty]$. The metric space H_q is defined as the subspace*

$$\mathsf{H}_q = \left\{(x,y) \in [-\infty,\infty]^2 \,:\, x \leq y\right\} \subseteq [-\infty,\infty]^2$$

with distance

$$d_{\mathsf{H}_q}\left((x,y),(x',y')\right) = \left(|x-x'|^q + |y-y'|^q\right)^{1/q}$$

if $q \in [1,\infty)$ and

$$d_{\mathsf{H}_q}\left((x,y),(x',y')\right) = \max\left(|x-x'|,|y-y'|\right)$$

if $q=\infty$. Here $|x-x'|$ denotes the distance between x and x' in $[-\infty,\infty]$.

*Let $\Delta \subseteq \mathsf{H}_q$ be the set of points of the form $(x,x) \in \mathsf{H}_q$. We refer to Δ as the **diagonal**.*

Informally, a persistence diagram is usually defined to be a multiset of points of H_q that has countably many off-diagonal points, and such that each point in the diagonal has countably infinite multiplicity. There are many ways in which this can be formalized; we prefer the following.

Definition 6.9.2. *A **persistence diagram** consists of a set X together with a function $\psi_X : X \to \mathsf{H}_q$ such that $\psi_X^{-1}(\mathsf{H}_q \setminus \Delta)$ is countable and such that the restriction $\psi_X|_{\psi_X^{-1}(\Delta)} : \psi_X^{-1}(\Delta) \to \Delta$ has countably infinite fibers.*

We will usually denote a persistence diagram (X,ψ_X) by its underlying set X. Let I be a set and let $\{x_i\}_{i\in I}$ be a collection of elements of $[0,\infty]$. Define their sum $\sum_{i\in I} x_i \in [0,\infty]$ as

$$\sum_{i\in I} x_i := \sup_{\substack{J\subseteq I\\ J \text{ finite}}} \sum_{i\in J} x_i.$$

Definition 6.9.3. *Let $p,q \in [1,\infty]$, let X and Y be persistence diagrams, and let $f : X \to Y$ be a function of sets. The **p-distortion** of f is defined as*

$$\mathrm{dist}^p(f) = \left(\sum_{x\in X} d_{\mathsf{H}_q}\left(\psi_X(x),\psi_Y(f(x))\right)^p\right)^{1/p}.$$

for $p \in [1,\infty)$, and as

$$\mathrm{dist}^p(f) = \sup_{x \in X} d_{\mathsf{H}_q}\left(\psi_X(x), \psi_Y(f(x))\right).$$

for $p = \infty$.

Definition 6.9.4. *Let $p \in [1,\infty]$. The locally persistent category of persistence diagrams, denoted by $\mathsf{PD}^{\mathsf{p}}_{\mathsf{q}}$, is the locally persistent category whose objects are persistence diagrams and whose morphisms are given by*

$$\mathrm{Hom}_{\mathsf{PD}^{\mathsf{p}}_{\mathsf{q}}}(X,Y)_\varepsilon = \left\{f : X \to Y \text{ injective function of sets} : \mathrm{dist}^p(f) \le \varepsilon\right\},$$

for $X, Y \in \mathsf{PD}^{\mathsf{p}}_{\mathsf{q}}$. Composition and identities are given by composition of functions and identity functions, respectively.

The restriction to injective functions is so that composition is well-defined, in the sense that, for $f : X \to Y$ and $g : Y \to Z$ functions between persistence diagrams, and $p \in [1,\infty]$, we have $\mathrm{dist}^p(g \circ f) \le \mathrm{dist}^p(g) + \mathrm{dist}^p(f)$.

The following result is immediate from the definitions. For a definition of the $\ell^p[\ell^q]$ matching (or Wasserstein) distance see, e.g., [Cho19b].

Proposition 6.9.5. *Let $p, q \in [1,\infty]$. The interleaving distance of $\mathsf{PD}^{\mathsf{p}}_{\mathsf{q}}$ coincides with the $\ell^p[\ell^q]$ matching (or Wasserstein) distance on persistence diagrams.* □

Remark 6.9.6. We note that, when $p \neq \infty$ or $q \neq \infty$, the locally persistent category of persistence diagrams does not arise as a category with a flow in any natural way. When $p = q = \infty$, there is a category with a flow whose objects are persistence diagrams (or barcodes), and whose interleaving distance is the bottleneck distance; in fact, this category with a flow is a functor category of the form $C^{\mathbf{R}}$. This structure and its usage to formulate the induced matching theorem, a refinement of the algebraic stability theorem, is the subject of [BL20].

Two important directions of work remain to be explored. One direction includes finding possibly larger locally persistent categories of persistence diagrams in which we can apply our theorems to deduce that the interleaving distance is complete and geodesic. For now, we refer the reader to [Cho19b] for the study of geodesics in spaces of persistence diagrams, and to [MMH11] for the completeness of these spaces.

The other direction includes connecting these locally persistent categories of persistence diagrams with locally persistent categories of persistent vector spaces in order to obtain a categorical proof of the algebraic stability theorem ([CCSGGO09]). We believe that this can be done by rephrasing the main result of [BL14] as a theorem about locally persistent categories.

Bibliography

[Ack13] N. Ackerman. "Completeness in generalized ultrametric spaces". In: *P-Adic Numbers, Ultrametric Analysis, and Applications* 5 (Apr. 2013).

[AFL03] M. Amico, P. Frosini, and C. Landi. "Optimal matching between reduced size functions". In: *DISMI, Univ. di Modena e Reggio Emilia, Italy, Technical report no.* 35 (Jan. 2003).

[BBI01] Dmitri Burago, Yuri Burago, and Sergei Ivanov. "A Course in Metric Geometry". In: *Graduate Studies in Math.* 33 (Jan. 2001).

[BG18] Nicolas Berkouk and Grégory Ginot. *A derived isometry theorem for constructible sheaves on* $\mathbb{R}$. 2018. arXiv: 1805.09694 [math.AT].

[BGMP14] Andrew J. Blumberg, Itamar Gal, Michael A. Mandell, and Matthew Pancia. "Robust Statistics, Hypothesis Testing, and Confidence Intervals for Persistent Homology on Metric Measure Spaces". In: *Found. Comput. Math.* 14.4 (Aug. 2014), 745–789. ISSN: 1615-3375.

[BL14] Ulrich Bauer and Michael Lesnick. "Induced Matchings of Barcodes and the Algebraic Stability of Persistence". In: *Proceedings of the Thirtieth Annual Symposium on Computational Geometry*. SOCG'14. Kyoto, Japan: Association for Computing Machinery, 2014, 355–364. ISBN: 9781450325943.

[BL17] Andrew J. Blumberg and Michael Lesnick. "Universality of the Homotopy Interleaving Distance". In: *CoRR* abs/1705.01690 (2017). arXiv: 1705.01690.

[BL20] Ulrich Bauer and Michael Lesnick. "Persistence Diagrams as Diagrams: A Categorification of the Stability Theorem". In: *Topological Data Analysis*. Cham: Springer International Publishing, 2020, pp. 67–96. ISBN: 978-3-030-43408-3.

[BP19] Nicolas Berkouk and Francois Petit. *Ephemeral persistence modules and distance comparison*. 2019. arXiv: 1902.09933 [math.AT].

[BS14] P. Bubenik and J. Scott. "Categorification of Persistent Homology". In: *Discrete & Computational Geometry* 51.3 (2014), pp. 600–627.

[BSN17] P. Bubenik, V. Silva, and V. Nanda. “Higher Interpolation and Extension for Persistence Modules”. In: *SIAM J. Appl. Algebra Geom.* 1.1 (2017), pp. 272–284.

[BSS13] P. Bubenik, V. Silva, and J. Scott. “Metrics for Generalized Persistence Modules”. In: *Foundations of Computational Mathematics* 15 (Dec. 2013).

[BV18] Peter Bubenik and Tane Vergili. “Topological spaces of persistence modules and their properties”. In: *Journal of Applied and Computational Topology* 2.3-4 (2018), pp. 233–269.

[CCBS14] Frédéric Chazal, William Crawley-Boevey, and Vin Silva. “The observable structure of persistence modules”. In: *Homology, Homotopy and Applications* 18 (Apr. 2014).

[CCMSW17] Samir Chowdhury, Nathaniel Clause, Facundo Memoli, Jose Angel Sanchez, and Zoe Wellner. *New families of simplicial filtration functors*. 2017. arXiv: 1712.00696 [math.AT].

[CCSGGO09] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. “Proximity of Persistence Modules and Their Diagrams”. In: *Proceedings of the Twenty-Fifth Annual Symposium on Computational Geometry*. SCG ’09. Aarhus, Denmark: Association for Computing Machinery, 2009, 237–246. ISBN: 9781605585017.

[CCSGMO09] Frédéric Chazal, David Cohen-Steiner, Leonidas J. Guibas, Facundo Mémoli, and Steve Y. Oudot. “Gromov-Hausdorff Stable Signatures for Shapes Using Persistence”. In: *Proceedings of the Symposium on Geometry Processing*. SGP ’09. Berlin, Germany: Eurographics Association, 2009, 1393–1403.

[CFP01] Francesca Cagliari, Massimo Ferri, and Paola Pozzi. “Size Functions from a Categorical Viewpoint”. In: *Acta Applicandae Mathematicae* 67 (July 2001), pp. 225–235.

[CGOS13] Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. “Persistence-Based Clustering in Riemannian Manifolds”. In: *J. ACM* 60.6 (Nov. 2013). ISSN: 0004-5411.

[Cho19a] Simon Cho. *Quantales, persistence, and magnitude homology*. 2019. arXiv: 1910.02905 [math.AT].

[Cho19b] Samir Chowdhury. *Geodesics in persistence diagram space*. 2019. arXiv: 1905.10820 [math.MG].

[CL05] Frédéric Chazal and André Lieutier. "Weak Feature Size and Persistent Homology: Computing Homology of Solids in Rn from Noisy Data Samples". In: *Proceedings of the Twenty-First Annual Symposium on Computational Geometry*. SCG '05. Pisa, Italy: Association for Computing Machinery, 2005, 255–262. ISBN: 1581139918.

[CM10a] Gunnar Carlsson and Facundo Mémoli. "Characterization, stability and convergence of hierarchical clustering methods". In: *J. Mach. Learn. Res.* 11 (2010), pp. 1425–1470. ISSN: 1532-4435.

[CM10b] Gunnar Carlsson and Facundo Mémoli. "Characterization, Stability and Convergence of Hierarchical Clustering methods". In: *Journal of Machine Learning Research* 11 (2010), pp. 1425–1470.

[CM10c] Gunnar Carlsson and Facundo Mémoli. "Multiparameter Hierarchical Clustering Methods". In: *Classification as a Tool for Research*. Ed. by Hermann Locarek-Junge and Claus Weihs. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 63–70. ISBN: 978-3-642-10745-0.

[CM18a] Samir Chowdhury and Facundo Mémoli. "A functorial Dowker theorem and persistent homology of asymmetric networks". In: *Journal of Applied and Computational Topology* (Sept. 2018).

[CM18b] Samir Chowdhury and Facundo Mémoli. "Explicit geodesics in Gromov-Hausdorff space". In: *Electronic Research Announcements* 25 (2018), p. 48. ISSN: 1935-9179.

[CO17] Mathieu Carrière and Steve Oudot. "Local Equivalence and Intrinsic Metrics between Reeb Graphs". In: *33rd International Symposium on Computational Geometry, SoCG 2017, July 4-7, 2017, Brisbane, Australia*. Ed. by Boris Aronov and Matthew J. Katz. Vol. 77. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 25:1–25:15.

[Cru19] J. Cruz. *Metric Limits in Categories with a Flow*. 2019. arXiv: 1901.04828.

[CSEH05] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. "Stability of Persistence Diagrams". In: *Proceedings of the Twenty-First Annual Symposium on Computational Geometry*. SCG '05. Pisa, Italy: Association for Computing Machinery, 2005, 263–271. ISBN: 1581139918.

[CSGO16] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. *The Structure and Stability of Persistence Modules*. SpringerBriefs in Mathematics. Springer International Publishing, 2016. ISBN: 9783319425450.

[CSO14] Frédéric Chazal, de Vin Silva, and Steve Oudot. "Persistence stability for geometric complexes". In: *Geometriae Dedicata* (2014).

[Day70] Brian Day. "On closed categories of functors". In: *Reports of the Midwest Category Seminar IV*. Ed. by S. MacLane, H. Applegate, M. Barr, B. Day, E. Dubuc, Phreilambud, A. Pultr, R. Street, M. Tierney, and S. Swierczkowski. Berlin, Heidelberg: Springer Berlin Heidelberg, 1970, pp. 1–38. ISBN: 978-3-540-36292-0.

[Day72] Brian Day. "A reflection theorem for closed categories". In: *Journal of Pure and Applied Algebra* 2.1 (1972), pp. 1 –11. ISSN: 0022-4049.

[DK69] B. J. Day and G. M. Kelly. "Enriched functor categories". In: *Reports of the Midwest Category Seminar III*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1969, pp. 178–191. ISBN: 978-3-540-36150-3.

[Dub70] E.J. Dubuc. *Kan Extensions in Enriched Category Theory*. Lecture notes in mathematics. Springer-Verlag, 1970.

[EBW15] Justin Eldridge, Mikhail Belkin, and Yusu Wang. "Beyond Hartigan Consistency: Merge Distortion Metric for Hierarchical Clustering". In: *Proceedings of The 28th Conference on Learning Theory*. Ed. by Peter Grünwald, Elad Hazan, and Satyen Kale. Vol. 40. Proceedings of Machine Learning Research. Paris, France: PMLR, 2015, pp. 588–606.

[EJM15] Herbert Edelsbrunner, Grzegorz Jabłoński, and Marian Mrozek. "The Persistent Homology of a Self-Map". In: *Found. Comput. Math.* 15.5 (Oct. 2015), 1213–1244. ISSN: 1615-3375. DOI: 10.1007/s10208-014-9223-y. URL: https://doi.org/10.1007/s10208-014-9223-y.

[FLM17] Patrizio Frosini, Claudia Landi, and Facundo Mémoli. "The Persistent Homotopy Type Distance". In: *Homology, Homotopy and Applications* 21 (Feb. 2017).

[FLRWBS14] Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, Larry Wasserman, Sivaraman Balakrishnan, and Aarti Singh. "Confidence sets for persistence diagrams". In: *Ann. Statist.* 42.6 (Dec. 2014), pp. 2301–2339.

[GMOTWW19] Ellen Gasparovic, Elizabeth Munch, Steve Oudot, Katharine Turner, Bei Wang, and Yusu Wang. *Intrinsic Interleaving Distance for Merge Trees.* 2019. arXiv: 1908.00063 [cs.CG].

[Hau95] Jean-Claude Hausmann. "On the Vietoris–Rips complexes and a cohomology theory for metric spaces". In: *Annals of mathematics studies* 138 (1995), pp. 175–188.

[Hir09] P.S. Hirschhorn. *Model Categories and Their Localizations.* Mathematical surveys and monographs. American Mathematical Society, 2009.

[Hov07] M. Hovey. *Model Categories.* Mathematical surveys and monographs. American Mathematical Society, 2007. ISBN: 9780821843611.

[INT16] A. O. Ivanov, N. K Nikolaeva, and A. A. Tuzhilin. "The Gromov–Hausdorff metric on the space of compact metric spaces is strictly intrinsic". In: *Mathematical Notes* 100 (2016), pp. 883–885.

[KCBRW16] Jisu KIM, Yen-Chi Chen, Sivaraman Balakrishnan, Alessandro Rinaldo, and Larry Wasserman. "Statistical Inference for Cluster Trees". In: *Advances in Neural Information Processing Systems 29.* Ed. by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett. Curran Associates, Inc., 2016, pp. 1839–1847.

[Kel74] G. M. Kelly. "Doctrinal adjunction". In: *Category Seminar.* Ed. by Gregory M. Kelly. Berlin, Heidelberg: Springer Berlin Heidelberg, 1974, pp. 257–280. ISBN: 978-3-540-37270-7.

[Kel82] G.M. Kelly. *Basic Concepts of Enriched Category Theory.* Lecture note series / London mathematical society. Cambridge University Press, 1982. ISBN: 9780521287029.

[KM20] Woojin Kim and Facundo Mémoli. "Spatiotemporal Persistent Homology for Dynamic Metric Spaces". In: *Discrete & Computational Geometry* (Feb. 2020), pp. 1–45.

[KSCRW19] Jisu Kim, Jaehyeok Shin, Frédéric Chazal, Alessandro Rinaldo, and Larry Wasserman. *Homotopy Reconstruction via the Cech Complex and the Rips Complex*. 2019. arXiv: 1903.06955 [math.AT].

[Lan98] S.M. Lane. *Categories for the Working Mathematician*. Graduate Texts in Mathematics. Springer New York, 1998. ISBN: 9780387984032.

[Lat01] J. Latschev. "Vietoris-Rips complexes of metric spaces near a closed Riemannian manifold". In: *Archiv der Mathematik* 77 (Dec. 2001), pp. 522–528.

[Law73] F. William Lawvere. "Metric spaces, generalized logic, and closed categories". In: *Rendiconti del Seminario Matematico e Fisico di Milano* 43.1 (1973), pp. 135–166. ISSN: 1424-9294.

[Les12] Michael Lesnick. "Multidimensional Interleavings and Applications to Topological Inference". PhD thesis. June 2012.

[Les15] Michael Lesnick. "The Theory of the Interleaving Distance on Multidimensional Persistence Modules". In: *Foundations of Computational Mathematics* 15 (2015).

[LW15] Michael Lesnick and Matthew Wright. "Interactive visualization of 2-D Persistence Modules". In: *eprint arXiv:1512.00180v1* (2015).

[Mel] Paul-André Mellies. "Parametric monads and enriched adjunctions". In progress. URL: https://www.irif.fr/~mellies/tensorial-logic/8-parametric-monads-and-enriched-adjunctions.pdf.

[MH17] L. McInnes and J. Healy. "Accelerated Hierarchical Density Based Clustering". In: *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*. 2017, pp. 33–42.

[Mie09] Grégory Miermont. "Tessellations of random maps of arbitrary genus". In: *Ann. Sci. Éc. Norm. Supér. (4)* 42.5 (2009), pp. 725–781. ISSN: 0012-9593.

[MMH11] Yuriy Mileyko, Sayan Mukherjee, and John Harer. “Probability measures on the space of persistence diagrams”. In: *Inverse Problems* 27.12 (2011), p. 124007.

[Mé17] Facundo Mémoli. *A Distance Between Filtered Spaces Via Tripods.* 2017. arXiv: 1704.03965 [math.AT].

[NSW08] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. “Finding the Homology of Submanifolds with High Confidence from Random Samples”. In: *Discrete Comput. Geom.* 39.1 (Mar. 2008), 419–441. ISSN: 0179-5376.

[PCR96] S. Priess-Crampe and Paulo Ribenboim. “Generalized ultrametric spaces I”. In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 66 (Dec. 1996), pp. 55–73.

[PCR97] S. Priess-Crampe and Paulo Ribenboim. “Generalized Ultrametric Spaces II”. In: *Abhandlungen Aus Dem Mathematischen Seminar Der Universitat Hamburg - ABH MATH SEM UNIV HAMBURG* 67 (Dec. 1997), pp. 19–31.

[Qui67] Daniel G. Quillen. *Homotopical algebra.* Springer, 1967.

[Qui73] Daniel Quillen. “Higher algebraic K-theory: I”. In: *Higher K-Theories.* Ed. by H. Bass. Berlin, Heidelberg: Springer Berlin Heidelberg, 1973, pp. 85–147. ISBN: 978-3-540-37767-2.

[Rie17] E. Riehl. *Category Theory in Context.* Aurora: Dover Modern Math Originals. Dover Publications, 2017. ISBN: 9780486820804.

[RS20] Alexander Rolle and Luis Scoccola. *Stable and consistent density-based clustering.* 2020. arXiv: 2005.09048 [math.ST].

[Rut98] J.J.M.M. Rutten. “Weighted colimits and formal balls in generalized metric spaces”. In: *Topology and its Applications* 89.1 (1998). Domain Theory, pp. 179 –202. ISSN: 0166-8641. DOI: https://doi.org/10.1016/S0166-8641(97)00224-1. URL: http://www.sciencedirect.com/science/article/pii/S0166864197002241.

[Seg16] Santiago Segarra. PhD thesis. University of Pennsylvania, 2016. URL: https://repository.upenn.edu/edissertations/2575.

[Shu06] Michael Shulman. *Homotopy limits and colimits and enriched homotopy theory*. 2006. arXiv: math/0610194 [math.AT].

[SMS18] V. Silva, E. Munch, and A. Stefanou. "Theory of interleavings on categories with a flow". In: *Theory and Applications of Categories* 33 (June 2018), pp. 583–607.

[SN10] Werner Stuetzle and Rebecca Nugent. "A Generalized Single Linkage Method for Estimating the Cluster Tree of a Density". In: *Journal of Computational and Graphical Statistics* 19.2 (2010), pp. 397–418. DOI: 10.1198/jcgs.2009.07049.

[Spa12] E.H. Spanier. *Algebraic Topology*. Springer New York, 2012. ISBN: 9781468493221.

[Sto79] A.H. Stone. "Inverse limits of compact spaces". In: *General Topology and its Applications* 10.2 (1979), pp. 203 –211. ISSN: 0016-660X.

[Stu12] Karl-Theodor Sturm. *The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces*. 2012. arXiv: 1208.0434 [math.MG].

[TMMH12] Katharine Turner, Yuriy Mileyko, Sayan Mukherjee, and John Harer. "Frechet Means for Distributions of Persistence Diagrams". In: *Discrete & Computational Geometry* 52 (June 2012).

[Wol74] Harvey Wolff. "V-cat and V-graph". In: *Journal of Pure and Applied Algebra* 4.2 (1974), pp. 123 –135. ISSN: 0022-4049.

[ZC04] Afra Zomorodian and Gunnar Carlsson. "Computing Persistent Homology". In: *Proceedings of the Twentieth Annual Symposium on Computational Geometry*. SCG '04. Brooklyn, New York, USA: Association for Computing Machinery, 2004, 347–356. ISBN: 1581138857.