

© Springer Verlag. The copyright for this contribution is held by Springer Verlag. The original publication is available at www.springerlink.com.

Constructing Shape Spaces from a Topological Perspective

Christoph Hofer^{1,3}, Roland Kwitt¹, Marc Niethammer², Yvonne Höller^{4,5},
Eugen Trinka^{3,4,5} and Andreas Uhl¹ for the ADNI

¹Department of Computer Science, University of Salzburg, Austria

²UNC Chapel Hill, NC, USA

³Spinal Cord Injury & Tissue Regeneration Centre (SCI-TReCS) Salzburg,
Paracelsus Medical University, Salzburg, Austria,

⁴Department of Neurology, Christian Doppler Medical Centre, Paracelsus Medical
University, Salzburg, Austria,

⁵Centre for Cognitive Neuroscience, Paracelsus Medical University, Salzburg, Austria

Abstract. We consider the task of constructing (metric) shape space(s) from a topological perspective. In particular, we present a generic construction scheme and demonstrate how to apply this scheme when *shape* is interpreted as the *differences that remain after factoring out translation, scaling and rotation*. This is achieved by leveraging a recently proposed injective functional transform of 2D/3D (binary) objects, based on persistent homology. The resulting shape space is then equipped with a similarity measure that is (1) by design robust to noise and (2) fulfills all metric axioms. From a practical point of view, analyses of object shape can then be carried out *directly* on segmented objects obtained from some imaging modality *without* any preprocessing, such as alignment, smoothing, or landmark selection. We demonstrate the utility of the approach on the problem of distinguishing segmented hippocampi from normal controls *vs.* patients with Alzheimer’s disease in a challenging setup where volume changes are no longer discriminative.

1 Introduction

Characterization and representation of shapes, cornerstones of human perception of objects, are fundamental and well-studied problems in computer vision and medical image analysis [14,9,5]. Applications range from recognition tasks to the analysis of longitudinal changes. While computer vision mostly focuses on 2D objects, medical image analysis typically studies shape in the context of (binary) 2D/3D segmentations of an anatomical structure (e.g., the hippocampus) that has been extracted from some imaging modality.

Conceptually, to study shapes in a mathematical framework, we need a formalism, a *shape space*, which captures the semantics of the colloquial term ‘shape’ or ‘shape space’. In this work, we present a versatile construction scheme for shape spaces and then follow Kendall [14] who informally defines shape as “*what is left when the differences which can be attributed to translations, rotations, and dilations have been quotiented out*”, to implement a concrete realization of that scheme. Since, a mandatory property of a shape space is not only

to separate shapes, but also to assign some sort of similarity measure [21], we equip the shape space with an appropriate metric.

In many applications, where the primary objective is to characterize 2D/3D objects, a *shape space* is often not directly defined, but implicitly generated through the invariances of features extracted from the objects of interest (cf. [28,27,30]). In such approaches, the intermediate feature extraction step essentially condenses some discriminative properties (e.g., local curvature) into a compact representation that can be used in a subsequent learning step (e.g., using SVMs). While these approaches work remarkably well for various tasks, it is challenging to study the shape space due to the missing explicit definition.

In contrast, here we are specifically interested in a formal definition of shape. Approaches along this line of research predominantly study landmark-based representations [14,9] of objects, manifolds (e.g., planar curves and surfaces) [19,15,25,3], and even point clouds [18]. In the seminal work of Kendall [14] for instance, objects are represented via k landmarks in \mathbb{R}^n and shapes are identified by quotienting out rotations, translations and scalings. This strategy has led to various elegant approaches to study shapes, most prominently by introducing a Riemannian structure on the space of landmarks [22,13,12]. While this allows to lift many concepts from statistics (e.g., regression) to the shape space, the data needs to be carefully hand-labeled, or landmarks need to be found via optimization. This is a challenging task on its own and typically requires careful preprocessing (e.g., smoothing). In a conceptually different line of research, shapes are considered as equivalence classes of manifolds [19,25] and then studied by quotienting out reparametrizations. We refer the reader to [3] for a comprehensive review. While these approaches eliminate the need for landmarking, they are mostly theoretical, require considerable preprocessing and are non-trivial to implement. Alternatively, shapes can be represented via point clouds in \mathbb{R}^n , and distances can be established via a metric on the isometry classes of compact metric spaces (e.g., Gromov-Hausdorff), cf. [18], where the shape representation is invariant under rigid motions (excluding scalings). While this can be beneficial in certain situations, the Gromov-Hausdorff metric is rather coarse and might not allow to tease-out fine-grained differences. Recently, ideas from topological data analysis [6] have emerged to study 2D/3D objects, either in the form of computing topological invariants of features [16,24,23], or by directly analyzing the object of interest [26]. The latter is particularly relevant to our work, as the proposed *persistent homology transform* (1) is injective and (2) it allows to directly analyze the raw data. This idea will enable us to formally define shape space(s).

Contribution. In detail, we develop a construction scheme in which Kendall’s notion of shape can be easily translated into the persistent homology transform framework of Turner et al. [26]. A suitable similarity measure with desirable properties such as robustness (to noise) can then be defined and shown to be an actual metric. This has important practical implications, as no preprocessing steps are required to analyze shapes after data has been collected.

2 Theoretical background

While persistent homology (PH) is fundamental for our proposed shape space construction, one does not need a deep understanding of many parts of this framework to grasp the key ideas. In fact, we refer the interested reader to [10] (and references therein) for a detailed introduction to persistent homology. We only introduce a few necessary key concepts next.

Definition 1 (Filtration of a simplicial complex by a function). Let K be a finite simplicial complex of dimension n and $f : K^{(0)} \rightarrow \mathbb{R}$, where $K^{(0)}$ denotes the 0-skeleton of K . If $\sigma = [v_0, \dots, v_k]$ is a k -simplex of K , then we naturally expand f to σ by setting $f(\sigma) = \max(\{v_i, 0 \leq i \leq k\})$. If $f(K) = \{a_1, a_2, \dots, a_N\}$ with $a_i \leq a_{i+1}$, we call the sequence

$$K_0 \subset K_1 \subset \dots \subset K_N$$

a filtration of K by f , where $K_0 = \emptyset$, $K_i = f^{-1}((-\infty, a_i])$ for $1 \leq i \leq N$.

Definition 2 (Persistence diagram). Let $\Delta = \{(x, x) \in \mathbb{R}^2 : \text{mult}(x) = \infty\}$ be the multiset of the diagonal in \mathbb{R}^2 , where mult denotes the multiplicity function. A persistence diagram is a multiset of the following form:

$$\mathcal{D} = \{(b, d) \in \mathbb{R}^2 : d - b > 0 \text{ and } 1 \leq \text{mult}((b, d)) < \infty\} \cup \Delta .$$

For a (finite) n -dimensional simplicial complex K and a function f defined on its 0-skeleton, we can interpret persistent homology as a mapping that associates to the filtration of K by f , (K, f) , a vector of persistence diagrams, i.e.,

$$(K, f) \xrightarrow{\text{PH}} (\mathcal{D}_0(K, f), \dots, \mathcal{D}_{n-1}(K, f)) . \quad (1)$$

Each \mathcal{D}_i essentially encodes information about the homology of dimension i of the complexes as the filtration parameter grows, i.e., as we progress from \emptyset to K . For instance, in dimension 0, $\mathcal{D}_0(K, f)$ captures the evolution of *connected components* of (K, f) . Hence, $(b, d) \in \mathcal{D}_0(K, f)$ identifies a connected component which first occurs in $f^{-1}((-\infty, b])$ (i.e., it is *born*) and merges into an older one in $f^{-1}((-\infty, d])$ (i.e., it *dies*). To understand why this construction is useful, we briefly review the Bottleneck and Wasserstein distances.

Definition 3 (Bottleneck/Wasserstein distance). For two persistence diagrams \mathcal{D} and \mathcal{E} , we define their Bottleneck and Wasserstein distances by

$$w_\infty(\mathcal{D}, \mathcal{E}) = \inf_{\eta} \sup_{x \in \mathcal{D}} \|x - \eta(x)\|_\infty \quad \text{and} \quad w_p^q(\mathcal{D}, \mathcal{E}) = \inf_{\eta} \left(\sum_{x \in \mathcal{D}} \|x - \eta(x)\|_q^p \right)^{\frac{1}{p}} ,$$

where $p, q \in \mathbb{N}$ and the infimum is taken over all bijections $\eta : \mathcal{D} \rightarrow \mathcal{E}$.

Note that $w_p^q \rightarrow w_\infty$ for $p \rightarrow \infty$. Using the Bottleneck distance, we can equip the space of persistence diagrams with a metric structure [20], and hence with an

induced topology. In fact, Chazal et al. showed [7, Theorem 3.2] that (under reasonable constraints) for two filtrated complexes $(K, f_K), (L, f_L)$, the mapping in Eq. (1) is *stable*; for finite complexes this means that $w_\infty(\mathcal{D}_i(K, f_K), \mathcal{D}_i(L, f_L))$ is bounded by the maximum of (1) the Gromov-Hausdorff distance, $d_{\text{GH}}(K, L)$, and (2) the $\sup_{(x,y) \in K \times L} |f_K(x) - f_L(y)|$. Further, the Wasserstein distance, w_p^q , inherits those stability properties [8]. Informally, stability reflects the property that small changes in the domain *do not* result in large changes in the co-domain.

Persistent Homology Transform. Our approach is based on the persistent homology transform (PHT), introduced by Turner et al. [26]. It uses the fact that for each direction $v \in \mathbb{S}^{n-1}$ the mapping $x \mapsto \langle x|v \rangle$ induces a filtration on complexes (in \mathbb{R}^n). Let \mathcal{M}_n be the set of all subsets of \mathbb{R}^n which can be written as finite simplicial complexes, where we assume two elements of \mathcal{M}_n are equivalent if their linear embeddings¹ are equal. Then, the PHT assigns to each $K \in \mathcal{M}_n$ a mapping, $\text{PHT}(K)$, given by

$$\text{PHT}(K) : \mathbb{S}^{n-1} \rightarrow \mathbb{D}^n, \quad v \mapsto (\mathcal{X}_0(K, v), \dots, \mathcal{X}_{n-1}(K, v)) ,$$

where $\mathcal{X}_k(K, v) = \mathcal{D}_k(K, \langle \cdot | v \rangle)$ and \mathbb{D} denotes the space of persistence diagrams. [26] showed that this mapping is *injective* for $\mathcal{M}_2/\mathcal{M}_3$. With this and the aforementioned stability results, they introduce a *family of metrics* on $\mathcal{M}_2, \mathcal{M}_3$, i.e.,

$$m(K, L) = m_{\text{dist}, n}(K, L) = \sum_{k=0}^{n-1} \int_{\mathbb{S}^{n-1}} \text{dist}(\mathcal{X}_k(K, v), \mathcal{X}_k(L, v)) \, dv , \quad (2)$$

where $n = 2, 3$ and dist is either w_∞ or w_p^q . When not necessary, we will omit the subscripts (dist, n) for readability. Fig. 1 illustrates the PHT for two filtration directions $v_1, v_2 \in \mathbb{S}^1$ and shows the corresponding persistence diagram(s). The metric space (\mathcal{M}_3, m) will be the basis for our shape space construction.

3 Constructing shape spaces

To build a *shape space* on top of the metric space (\mathcal{M}_3, m) , we follow a group-theoretic construction. In particular, let (\mathbf{G}, \odot) be a group, S be a set and $\triangleright : \mathbf{G} \times S \rightarrow S, (g, s) \mapsto g \triangleright s$ a (left) action of \mathbf{G} on S . For $s \in S$ we denote by

$$\mathbf{G} \triangleright s = \{g \triangleright s : g \in \mathbf{G}\} \quad \text{and} \quad \mathbf{G} \triangleright S = \{\mathbf{G} \triangleright s : s \in S\}$$

the *orbit* of s and *the set of orbits* of S , resp., under the action \triangleright of \mathbf{G} on S . The following two lemmas will enable us to (1) obtain a metric on partitions of a metric space (Lemma 1) and (2) to establish a metric on the set of orbits induced by a group action on a metric space (Lemma 2).

Lemma 1 (Selection of representative). *Let (X, d) be a metric space with a partition $\mathcal{P} = \{X_i : i \in I\}$, i.e., $X_i \subset X, X_i \cap X_j = \emptyset$ for $i \neq j$ and $\bigcup X_i = X$. If $\iota : \mathcal{P} \rightarrow X$ is an injective mapping then*

$$\tilde{d}(X_i, X_j) = d(\iota(X_i), \iota(X_j)) \quad \text{is a metric on } \mathcal{P}.$$

¹ For dimension $k > 0$, we can take the interior of the convex hull of the defining vertices and, for dimension 0, a simplex is mapped to its defining vertex.




Fig. 1: Illustration of the PHT [26] for a hippocampus slice, represented as a (cubical) simplicial complex K , along two filtration directions v_1 and v_2 . The corresponding persistence diagram(s), $\mathcal{X}_0(K, v_1)$ and $\mathcal{X}_0(K, v_2)$, are shown on the *right*. For direction v_1 , we artificially added **noise** to the hippocampus; the corresponding points in the plot are highlighted in **red** (best-viewed in color).

Proof (Sketch). The identity of indiscernibles follows from the injectivity of ι ; the remaining metric axioms follow immediately from d . \square

Lemma 2. Let (X, d) be a metric space, (\mathbf{G}, \odot) a compact topological group and let $\triangleright : \mathbf{G} \times X \rightarrow X$ be a (left) action of \mathbf{G} on X , such that

- (i) $d(x, g \triangleright y)$ is continuous in g , and
- (ii) $d(x, y) = d(g \triangleright x, g \triangleright y)$, $\forall x, y \in X, \forall g \in \mathbf{G}$.

Then

$$\tilde{d}(\mathbf{G} \triangleright x, \mathbf{G} \triangleright y) = \min_{g \in \mathbf{G}} d(x, g \triangleright y)$$

is a metric on $\mathbf{G} \triangleright X$.

Proof. The proof consists of two parts.

Part (1). First, we have to show that $\min_g d(x, g \triangleright y)$ is well-defined, i.e., the minimum exists, and second, it does not depend on the representative of $\mathbf{G} \triangleright x$ and $\mathbf{G} \triangleright y$, respectively. The first property follows directly from condition (i) and the extreme value theorem, as d is continuous in g and \mathbf{G} is compact. To show the second property, we choose two representatives $u \triangleright x \in \mathbf{G} \triangleright x$, $v \triangleright y \in \mathbf{G} \triangleright y$, with $u, v \in \mathbf{G}$ and consider

$$\begin{aligned} \min_{g \in \mathbf{G}} d(u \triangleright x, g \triangleright (v \triangleright y)) &\stackrel{(ii)}{=} \min_{g \in \mathbf{G}} d(u^{-1} \triangleright (u \triangleright x), u^{-1} \triangleright (g \triangleright (v \triangleright y))) \\ &= \min_{g \in \mathbf{G}} d(u^{-1} \odot u \triangleright x, u^{-1} \odot g \odot v \triangleright y) \\ &= \min_{g \in \mathbf{G}} d(x, u^{-1} \odot g \odot v \triangleright y) \\ &\stackrel{(*)}{=} \min_{g \in \mathbf{G}} d(x, g \triangleright y) . \end{aligned}$$

The last identity (*) follows from the fact that if g runs through all of \mathbf{G} , so does $u^{-1} \odot g \odot v$.

Part (2). We will now verify the metric axioms.

Non-negativity. Follows from the non-negativity of d .

Symmetry. This is a consequence of condition (ii) as $d(x, g \triangleright y) = d(g \triangleright y, x) = d(g^{-1} \triangleright (g \triangleright y), g^{-1} \triangleright x) = d(y, g^{-1} \triangleright x)$ which does not affect the minimum by the same arguments as above, see (*).

Identity of indiscernibles. First let $\mathbf{G} \triangleright y = \mathbf{G} \triangleright x$. Then $\exists g' \in \mathbf{G}$ such that $y = g' \triangleright x$. Hence,

$$\tilde{d}(\mathbf{G} \triangleright x, \mathbf{G} \triangleright y) = \min_{g \in \mathbf{G}} d(x, g \triangleright (g' \triangleright x)) = 0 ,$$

as $(g')^{-1} \in \mathbf{G}$. Second, let $\tilde{d}(\mathbf{G} \triangleright x, \mathbf{G} \triangleright y) = 0$, then there exists some g_0 such that $d(x, g_0 \triangleright y) = 0$. Hence, $x = g_0 \triangleright y$ which means that $\mathbf{G} \triangleright x = \mathbf{G} \triangleright y$.

Triangle Inequality. Consider $\mathbf{G} \triangleright x, \mathbf{G} \triangleright y, \mathbf{G} \triangleright z \in \mathbf{G} \triangleright X$. For $g' \in \mathbf{G}$ we get

$$\begin{aligned} \tilde{d}(\mathbf{G} \triangleright x, \mathbf{G} \triangleright y) &= \min_{g \in \mathbf{G}} d(x, g \triangleright y) \\ &\stackrel{(**)}{\leq} \min_{g \in \mathbf{G}} \left(d(x, g' \triangleright z) + d(g' \triangleright z, g \triangleright y) \right) \\ &= d(x, g' \triangleright z) + \min_{g \in \mathbf{G}} d(g' \triangleright z, g \triangleright y) \\ &\stackrel{(ii)}{=} d(x, g' \triangleright z) + \min_{g \in \mathbf{G}} d(z, (g')^{-1} \triangleright (g \triangleright y)) \\ &\stackrel{(*)}{=} d(x, g' \triangleright z) + \tilde{d}(\mathbf{G} \triangleright z, \mathbf{G} \triangleright y). \end{aligned}$$

As g' is arbitrary, this inequality holds also for $g'_0 \in \mathbf{G}$ with $d(x, g'_0 \triangleright z) = \min_{g'} d(x, g' \triangleright z) = \tilde{d}(\mathbf{G} \triangleright x, \mathbf{G} \triangleright z)$; further, (**) holds due to the triangle inequality of d and (*) holds with the same argumentation as in Part (1). \square

4 Practical implementation

In this section, we show a practical realization of the concepts introduced in Sec. 3. In particular, we (1) construct a shape space in the spirit of Kendall, where translations, scalings and rotations are factored out and (2) show how the metric m can be pushed forward in this construction. At the end of the section, we demonstrate how these ideas translate into an implementation for binary segmentations defined on a pixel/voxel grid.

Remark 1. We follow the convention that if we interpret $K \in \mathcal{M}$ as a simplicial complex, applying $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ to K means to apply φ to each defining vertex, e.g., if $K = (\{[v_1, v_2], [v_1], [v_2]\})$ then $\varphi(K) = \{[\varphi(v_1), \varphi(v_2)], [\varphi(v_1)], [\varphi(v_2)]\}$.

Definition 4. Let \mathbf{S} be a subgroup of the homeomorphism group \mathbf{H} of \mathbb{R}^n and $\mathcal{K} \subset \mathcal{M}_n$ such that $\mathbf{S}(\mathcal{K}) = \{s(K) : s \in \mathbf{S}, K \in \mathcal{K}\} \subseteq \mathcal{K}$ for all $s \in \mathbf{S}$. Then

$$\triangleright_{\mathbf{S}} : \mathbf{S} \times \mathcal{K} \rightarrow \mathcal{K}, \quad (s, K) \mapsto s \triangleright_{\mathbf{S}} K = s(K)$$

is well-defined and we call \mathbf{S} a group of similarity transformations on \mathcal{K} . We further define by

$$\mathfrak{S}(\mathbf{S}, \mathcal{K}) = \mathbf{S} \triangleright_{\mathbf{S}} \mathcal{K}$$

the (simplicial) shape space of \mathcal{K} with respect to \mathbf{S} . If $\mathfrak{S}(\mathbf{S}, \mathcal{K})$ is equipped with a metric d then we call it a metric (simplicial) shape space.

If it is clear from the context, we will omit the subscript \mathbf{S} in $\triangleright_{\mathbf{S}}$ and simply write \triangleright . Also, we write \mathfrak{S} instead of $\mathfrak{S}(\mathbf{S}, \mathcal{K})$ if the particular selection of \mathbf{S} and \mathcal{K} is not of specific importance. We now fix $\mathcal{K} = \mathcal{M}_3$ (using metric m) and build a shape space that is invariant under translations $\mathbf{Tr}(\cong \mathbb{R}^n)$, barycentric scalings $\mathbf{Sc}(\cong \mathbb{R}^+)$ and barycentric rotations $\mathbf{Ro}(\cong \mathbf{SO}(3))$.

Remark 2. If (\mathfrak{S}, d) is a metric shape space and ι is a (bijective) isometry to some metric space (X, \tilde{d}) then we will identify $\iota(\mathfrak{S})$ with X and d with \tilde{d} . This simplifies notation, as we can refer to (X, \tilde{d}) as a metric simplicial shape space.

Translations. We start by factoring out *translations*.

Definition 5. Let, for $K \in \mathcal{K}$, b_K be the barycenter of K . Then, we set

$$\iota_{\mathbf{Tr}} : \mathbf{Tr} \triangleright \mathcal{K} \rightarrow \mathcal{K}, \quad \mathbf{Tr} \triangleright K \mapsto K - b_K .$$

Remark 3 ($\iota_{\mathbf{Tr}}$ is well defined and injective). Let $t \in \mathbf{Tr}$, then $b_{t \triangleright K} = b_{K+t} = b_K + t$. Now let $t \triangleright K$ be a representative of $\mathbf{Tr} \triangleright K$; then we get $t \triangleright K - b_{t \triangleright K} = K + t - b_{K+t} = K + t - b_K - t = K - b_K$ which implies that $\iota_{\mathbf{Tr}}$ does not depend on the representative of $\mathbf{Tr} \triangleright K$. If $\iota_{\mathbf{Tr}}(\mathbf{Tr} \triangleright K) = K - b_K = K - b_L = \iota_{\mathbf{Tr}}(\mathbf{Tr} \triangleright L)$, then $L = K - b_K + b_L \in \mathbf{Tr} \triangleright K$ which implies $\mathbf{Tr} \triangleright K = \mathbf{Tr} \triangleright L$, and injectivity follows.

Since $\iota_{\mathbf{Tr}}$ is injective and $\mathbf{Tr} \triangleright \mathcal{K}$ is a partition of \mathcal{K} we can invoke Lemma 1 and therefore define a metric $d_{\mathbf{Tr}}$ on $\mathfrak{S}(\mathcal{K}, \mathbf{Tr})$ by setting $d_{\mathbf{Tr}}(\mathbf{Tr} \triangleright K, \mathbf{Tr} \triangleright L) = m(\iota_{\mathbf{Tr}}(K), \iota_{\mathbf{Tr}}(L))$. This metric simplicial shape space is isometric to $(\mathcal{K}_{\mathbf{Tr}}, m)$ where $\mathcal{K}_{\mathbf{Tr}} = \iota_{\mathbf{Tr}}(\mathfrak{S}(\mathcal{K}, \mathbf{Tr}))$. After this first step, we have a subset of \mathcal{M}_3 equipped with m .

Scalings. In order to factor-out *barycentric scalings*, we construct the next shape space $\mathfrak{S}(\mathcal{K}_{\mathbf{Tr}}, \mathbf{Sc})$ from $(\mathcal{K}_{\mathbf{Tr}}, m)$ and equip it with a metric based on m . This is possible, as the barycenter remains unchanged by scalings.

Definition 6. Let, for $K \in \mathcal{K}_{\mathbf{Tr}}$, rad_K be the radius of K , i.e., $\text{rad}_K = \max_{x \in K} \|x\|_2$. Then, we set

$$\iota_{\mathbf{Sc}} : \mathbf{Sc} \triangleright \mathcal{K}_{\mathbf{Tr}} \rightarrow \mathcal{K}, \quad \mathbf{Sc} \triangleright K \mapsto \frac{1}{\text{rad}_K} \cdot K .$$

Remark 4 ($\iota_{\mathbf{Sc}}$ is well defined and injective). For $s \in \mathbf{Sc}$ it holds that $\text{rad}_{s \triangleright K} = \text{rad}_{s \cdot K} = s \cdot \text{rad}_K$. Now let $s \triangleright K$ be a representative of $\mathbf{Sc} \triangleright K$, then $\text{rad}_{s \triangleright K}^{-1} \cdot (s \triangleright K) = s^{-1} \cdot \text{rad}_K^{-1} \cdot s \cdot K = \text{rad}_K^{-1} \cdot K$. This implies that $\iota_{\mathbf{Sc}}$ does not depend on the representative in $\mathbf{Sc} \triangleright K$. If $\iota_{\mathbf{Sc}}(\mathbf{Sc} \triangleright K) = \text{rad}_K^{-1} \cdot K = \text{rad}_L^{-1} \cdot L = \iota_{\mathbf{Sc}}(\mathbf{Sc} \triangleright L)$, then $L = \text{rad}_L \cdot \text{rad}_K^{-1} \cdot K \in \mathbf{Sc} \triangleright K$ which implies $\mathbf{Sc} \triangleright K = \mathbf{Sc} \triangleright L$, and injectivity follows.

In the same manner as above, we can invoke Lemma 1. In detail, we set $\mathcal{K}_{\mathbf{Sc}} = \iota_{\mathbf{Sc}}(\mathfrak{S}(\mathcal{K}_{\mathbf{Tr}}, \mathbf{Sc}))$ and get the translation- and barycentric scaling-invariant metric simplicial shape space $(\mathcal{K}_{\mathbf{Sc}}, \mathfrak{m})$.

Rotations. In our final step, we factor-out *barycentric rotations*.

Lemma 3. *For $r \in \mathbf{Ro}$ and $K, L \in \mathcal{K}_{\mathbf{Sc}}$, it holds that*

- (i) $\mathfrak{m}(K, r \triangleright L)$ is continuous in r , and
- (ii) $\mathfrak{m}(K, L) = \mathfrak{m}(r \triangleright K, r \triangleright L)$.

Remark 5. In the following proof, we use two identities: (1) for a metric space (X, d) and $x, y, y' \in X$, we have $|d(x, y) - d(x, y')| \leq d(y, y')$; (2) $\mathcal{X}_k(K, d) = \mathcal{X}_k(r \triangleright K, r \triangleright d)$ for $r \in \mathbf{Ro}$ (as our filtration uses inner products, see Sec. 2).

Proof (of (i)). Consider Eq. (2); it suffices to show that $\text{dist}(\mathcal{X}_k(K, d), \mathcal{X}_k(r \triangleright L, d))$ is continuous in r . Let $r, r' \in \mathbf{Ro}$. For $\varepsilon > 0$, we obtain

$$\begin{aligned} |\text{dist}(\mathcal{X}_k(K, d), \mathcal{X}_k(r \triangleright L, d)) - \text{dist}(\mathcal{X}_k(K, d), \mathcal{X}_k(r' \triangleright L, d))| &\stackrel{\text{Rem. 5 (1)}}{\leq} \\ &\text{dist}(\mathcal{X}_k(r \triangleright L, d), \mathcal{X}_k(r' \triangleright L, d)) \stackrel{\text{Rem. 5 (2)}}{=} \\ &\text{dist}(\mathcal{X}_k(L, r^{-1} \triangleright d), \mathcal{X}_k(L, r'^{-1} \triangleright d)) \stackrel{[26][\text{Lem. 2.1}]}{\leq} \\ &C_L \| (r^{-1} \triangleright d) - (r'^{-1} \triangleright d) \|_2 \leq C_L \| r^{-1} - r'^{-1} \|_F \leq \varepsilon , \end{aligned} \quad (3)$$

where $C_L > 0$ is a constant dependent on L . Since, \mathbf{Ro} is a topological group, the last term in this transformation depends continuously on r and r' and hence, for $\|r - r'\|_F < \delta$ sufficiently small, Eq. (3) is always fulfilled. \square

Proof (of (ii)). It was shown in [2][Proposition 2 (iv)] that for $f : \mathbb{S}^{n-1} \rightarrow \mathbb{R}$, the integral over the spherical surface is rotation invariant, i.e., $\int_{\mathbb{S}^{n-1}} f = \int_{\mathbb{S}^{n-1}} f \circ r$ for $r \in \mathbf{Ro}$. In combination with Rem. 5 (2), this is sufficient for (ii) to hold. \square

Most importantly, Lemma 3 now enables us to invoke Lemma 2 with $(X, d) = (\mathcal{K}_{\mathbf{Sc}}, \mathfrak{m})$ and $\mathbf{G} = \mathbf{Ro}$. Consequently, we can equip $\mathfrak{S}(\mathcal{K}_{\mathbf{Sr}}, \mathbf{Ro})$ with the metric

$$\tilde{\mathfrak{m}}(\mathbf{Ro} \triangleright K, \mathbf{Ro} \triangleright L) = \min_{r \in \mathbf{Ro}} \mathfrak{m}(K, r \triangleright L) .$$

In summary, we obtain the following *sequence of metric spaces*

$$\boxed{(\mathcal{K}, \mathfrak{m}) \xrightarrow{\iota} (\mathcal{K}_{\mathbf{Tr}}, \mathfrak{m}) \xrightarrow{\iota} (\mathcal{K}_{\mathbf{Sr}}, \mathfrak{m}) \xrightarrow{\pi} (\mathfrak{S}(\mathcal{K}_{\mathbf{Sr}}, \mathbf{Ro}), \tilde{\mathfrak{m}})} \quad (4)$$

where π is the canonical mapping $K \mapsto \mathbf{Ro} \triangleright K$. For $K, L \in \mathcal{K}$ we can thus calculate the distance between K and L as

$$\tilde{\mathfrak{m}}(K, L) = \min_{r \in \mathbf{Ro}} \mathfrak{m} \left(\frac{1}{\text{rad}_{K-b_K}} (K - b_K), r \triangleright \left(\frac{1}{\text{rad}_{L-b_L}} (L - b_L) \right) \right) . \quad (5)$$

Implementation. To use the constructed shape space in the context of binary voxel data (e.g., segmentations of some anatomical structure), we need a transition of the binary voxel structure, V , to $K_V \in \mathcal{K}$. In detail, $V \mapsto K_V \in \mathcal{K}$ is achieved by interpreting each voxel as a cube in \mathbb{R}^3 and setting the 0-skeleton of K_V , $K_V^{(0)}$, as the union of the centers of those cubes. Then we choose a construction procedure, denoted as **Build**, which iteratively constructs the higher dimensional skeletons from $K_V^{(0)}$, i.e., $K = \text{Build}(K_V^{(0)})$. However, the choice of **Build** is not canonical. In our experiments, we used the algorithm in [29] to compute persistence diagrams for cubical data (implemented in **DIPHA**²) from which we obtain the PHTs. Fig. 2 (*right*) illustrates this particular choice of **Build**. To calculate the integral in Eq. (2), we decided to use a Lebedev integration scheme with 26 points, which is a quadrature rule optimized for functions defined on a sphere. The key idea is to integrate $f : \mathbb{S}^2 \rightarrow \mathbb{R}$ by sampling on a point set which is invariant under the octahedral symmetry group \mathbf{O} (including reflections), i.e., a Lebedev grid. The resulting quadrature formula (see [4] for theoretical details) is invariant under \mathbf{O} . To improve efficiency, we calculated the PHT for each object with respect to the 26 directions contained in the Lebedev grid. The minimization in Eq. (5) to calculate \tilde{m} can then be approximated by reducing the original minimization domain from $\mathbf{SO}(3)$ to \mathbf{R} , where $\mathbf{R} \subset \mathbf{O}$ is the octahedral rotation group. The stability of the Lebedev grid under \mathbf{O} guarantees that the application of each $r \in \mathbf{R}$ just causes a permutation of the 26 directions. Since \mathbf{R} has rank 24, this means each evaluation of \tilde{m} results in (less than) $24 \cdot 26$ evaluations of dist . In our experiments, we set $\text{dist} = w_2^2$, but remark that other choices might be better, depending on the specific application. In addition, we map points of the form $(b, \infty) \in \mathcal{X}_k(K, v)$, see Fig. 1, to $(b, \max_{x \in K^{(0)}} \langle x|v \rangle)$.

5 Experiments

We demonstrate the utility of our approach on (1) a simple toy example and (2) on the problem of separating binary 3D segmentations of the *hippocampus* obtained from healthy controls and patients with Alzheimer’s disease (AD). As it is well-known (cf. [1]) that volume changes of the hippocampus are indicative of AD, we construct a particularly challenging setup where we only consider hippocampi from patients that are *close* in volume; in other words, volumetric changes are no longer discriminative in this setup.

Toy example. Fig. 2 presents a toy example of one hippocampus slice, represented as K . The distance plot shows the behavior of the metric \tilde{m} when computing $\tilde{m}(K, K')$ where K' is obtained by rotating the hippocampus in steps of $(360^\circ/64) \cdot k$, for $0 \leq k \leq 64$ (with nearest-neighbor interpolation). For comparison, we list the distances to two artificially modified versions of the hippocampus (changes in **orange**). As we can see, the distances $\tilde{m}(K, K')$ are *lower* than the distances to the modified hippocampi (i.e., $\tilde{m}(K, M), \tilde{m}(K, L)$). As the hippocampus is rather small (< 500 pixel) this behavior is important, as the interpolation creates artifacts that are similar in size to salient object features.

² available online at <https://github.com/DIPHA/dipha>




Fig. 2: 2D version of \tilde{m} , approx. with 8 directions: $(0, 1), \frac{1}{\sqrt{2}}(1, 1), \dots, \frac{1}{\sqrt{2}}(1, -1)$ (left); Illustration of the construction scheme, **Build**, of a simplicial complex on a (pixel/voxel) grid in 2D (right). Simplices are shown in blue, pixel as squares.

Real data. We use segmentations of the *left hippocampus*, obtained from the ADNI dataset via MAPER [11]. All segmentations are publicly available. To focus on actual shape changes, we restrict the data to hippocampi with voxel volumes 1000 ± 50 . Under this constraint, we obtain 18 hippocampi associated with AD and 24 hippocampi associated with healthy controls. By design, a Wilcoxon rank-sum test for equality of the median volume does not allow to reject the null-hypothesis at any reasonable significance level (e.g., 0.01, 0.05).

Setup. We consider two setups: *First*, we assess a setup where rigid pre-alignment (implemented via NiftiReg; one healthy control randomly selected as target) is used and minimization of rotations is omitted in the metric. In other words, we work in the shape space $(\mathcal{K}_{\text{sr}}, m)$, see Eq. (4). In the *second* setup, we invoke the full metric \tilde{m} (w/o pre-alignment). For classification, we use a simple k -NN classifier and report average classification performance over 1000 balanced cross-validation splits, i.e., 15/15 AD and controls are randomly chosen to configure the k -NN classifier, 3/3 AD and controls are chosen for testing in each run.

Results. Table 1 lists the classification results over a range of k for the k -NN classifier. Several points are worth pointing out. *First*, the results indicate that rigid pre-alignment performs worse than letting the metric take care of rotations. This can be attributed to the fact that the metric does not rely on any registration procedure which, in case of small objects such as the hippocampus, can easily lead to slight misalignment and consequently confound the similarity measure used for classification. *Second*, although volumetric changes cannot separate the groups (as previously shown by the hypothesis test), our metric still allows to distinguish subjects from the AD and control group with performance comparable to what has been previously reported in the literature (although our approach is not specifically-tailored to the problem). *Third*, we highlight that *raw segmentations*, as produced by MAPER, were used in the experiment. While segmentation artifacts might still be present, we argue that it is undesirable to eliminate them during preprocessing, as we cannot reliably distinguish between noise and potentially discriminative/salient features *a-priori*.

$k(-\text{NN})$	Pre-Align	Ours (full metric \tilde{m})	Ours (full metric \tilde{m} , LOO)
3	0.78 ± 0.15	0.81 ± 0.15	0.81 ± 0.39
7	0.79 ± 0.16	0.83 ± 0.14	0.88 ± 0.32
11	0.75 ± 0.15	0.84 ± 0.14	0.83 ± 0.37
15	0.78 ± 0.14	0.85 ± 0.13	0.83 ± 0.37
19	0.75 ± 0.14	0.84 ± 0.15	0.83 ± 0.37

Table 1: Classification accuracies for distinguishing hippocampi of subjects with AD *vs.* healthy controls: (1) via (rigid) pre-alignment of segmentations where the metric *does not* take care of rotations (*first* column) and (2) by using the *full metric* \tilde{m} (*second/third* column). LOO denotes leave-one-out cross-validation.

6 Discussion

We introduced a versatile construction scheme for shape spaces (equipped with appropriate metrics) on top of a functional transform for objects, based on persistent homology. In a sense this framework (1) is *exact*, as no preprocessing is required and (2) allows to easily refine/coarsen the metric by choosing a higher or lower number of filtration directions, respectively. We argue that avoiding preprocessing is a desirable property, as the data processing theorem [17] specifically advises against reducing information in early stages of a processing chain. While we focused on the construction of a shape space in the sense of Kendall, the principles outlined in Sec. 3 constitute a generic *construction kit* to build all sorts of different shape spaces, with one particular realization shown in Eq. (4). At this stage, it is however unclear if our particular choice of base metric, i.e., the w_2^2 distance, is the most appropriate choice. In fact, a data-driven approach to select the parameters of the metric might be better. Another, conceptually interesting, extension could be to lift the metric to reproducing Kernel Hilbert spaces (cf. [24]), which would readily enable statistical computations.

References

1. L. G. Apostolova, A. E. Green, S. Babakchianian, K. S. Hwang, Y. Y. Chou, A. W. Toga, and P. M. Thompson. Hippocampal atrophy and ventricular enlargement in normal aging, mild cognitive impairment, and Alzheimer Disease. *Alzheimer Dis. Assoc. Disord.*, 26(1):17–27, 2012.
2. J. A. Baker. Integration over spheres and the divergence theorem for balls. *Am. Math. Monthly*, 104(1):36–47, 1997.
3. M. Bauer, M. Bruveris, and P. W. Michor. Overview of the geometries of shape spaces and diffeomorphism groups. *Math. Imaging Vis.*, 50(1):60–97, 2014.
4. C. H. L. Beentjes. Quadrature on a Spherical Surface, 2015. Online, <http://people.maths.ox.ac.uk/beentjes/Essays/QuadratureSphere.pdf>.
5. F. Bookstein. Size and shape spaces for landmark data in two dimensions. *Statist. Sci.*, 1(2):181–242, 1986.
6. G. Carlsson. Topology and data. *Bull. Amer. Math. Soc.*, 46:255–308, 2009.
7. F. Chazal, D. Cohen-Steiner, L. J. Guibas, F. Mémoli, and S. Y. Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum*, 28(5), 2009.

8. D. Cohen-Steiner, H. Edelsbrunner, J. Harer, and Y. Mileyko. Lipschitz functions have L_p -stable persistence. *Found. Comput. Math.*, 10(2):127–139, 2010.
9. I. L. Dryden and K. V. Mardia. *Statistical Shape Analysis*. Wiley series in probability and statistics. Wiley & Sons, New York, 1998.
10. H. Edelsbrunner and J. L. Harer. *Computational Topology : An Introduction*. American Mathematical Society, 2010.
11. R. A. Heckemann, S. Keihaninejad, P. Aljabar, C. Nielsen, K. R. Gray, D. Rueckert, J. V. Hajnal, and A. Hammers. Automatic morphometry in Alzheimer’s disease and mild cognitive impairment. *Neuroimage*, 56(4):2024–2037, 2011.
12. J. Hinkle, P. T. Fletcher, and S. Joshi. Intrinsic polynomials for regression on Riemannian manifolds. *Imaging Vis.*, 50:32–52, 2014.
13. J. Hinkle, P. Muralidharan, P. T. Fletcher, and S. Joshi. Polynomial regression on Riemannian manifolds. *ECCV*, 2012.
14. D. G. Kendall. Shape manifolds, Procrustean metrics, and complex projective spaces. *Bull. London Math. Soc.*, 16(2):81–121, 1984.
15. S. Kurttek, E. Klassen, Z. Ding, and A. Srivastava. A novel Riemannian framework for shape analysis of 3D objects. In *CVPR*, 2010.
16. C. Li, M. Ovsjanikov, and F. Chazal. Persistence-based structural recognition. *CVPR*, 2014.
17. D. MacKay. *Information theory, inference, and learning algorithms*. Cambridge University Press, 2003.
18. F. Mémoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.*, 5(3):313–347, 2005.
19. P. W. Michor and D. B. Mumford. Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc.*, 8(1):1–48, 2006.
20. Y. Mileyko, S. Mukherjee, and J. Harer. Probability measures on the space of persistence diagrams. *Inverse Probl.*, 27(12), 2011.
21. D. B. Mumford. Mathematical theories of shape: do they model perception? *Proc. of SPIE*, 1570:2–10, 1991.
22. P. T. Fletcher. Geodesic regression and the theory of least squares on Riemannian manifolds. *Int. J. Comput. Vision*, 105(2):171–185, 2012.
23. D. Pachauri, C. Hinrichs, M. K. Chung, S. C. Johnson, and V. Singh. Topology-based kernels with application to inference problems in Alzheimers disease. *IEEE Trans. Med. Imaging*, 30(10):1760–1770, 2011.
24. R. Reininghaus, U. Bauer, S. Huber, and R. Kwitt. A stable multi-scale kernel for topological machine learning. *CVPR*, 2015.
25. J. Shah. An H^2 Riemannian metric on the space of planar curves modulo similitudes. *Adv. Appl. Math.*, 51(4):483–506, 2013.
26. K. Turner, S. Mukherjee, and D. M. Boyer. Persistent homology transform for modeling shapes and surfaces. *Inf. Inference*, 3(4):310–344, 2014.
27. O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011.
28. R. C. Veltkamp and M. Hagedoorn. State of the art in shape matching. *Principles of Visual Information Retrieval*, pages 87–119, 2001.
29. H. Wagner, C. Chen, and E. Vučini. Efficient computation of persistent homology for cubical data. *Topological Methods in Data Analysis and Visualization II: Theory, Algorithms, and Applications*, pages 91–106, 2012.
30. M. Yang, K. Kpalma, and J. Ronsin. A survey of shape feature extraction techniques. *Pattern Recognit.*, 15(7):43–90, 2008.