Persistent Betti numbers for a noise tolerant shape-based approach to image retrieval

Patrizio Frosini[a,c], Claudia Landi[b,c]

[a] *Dipartimento di Matematica, Università di Bologna*
[b] *DiSMI, Università di Modena e Reggio Emilia*
[c] *ARCES, Università di Bologna*

Abstract

In content-based image retrieval a major problem is the presence of noisy shapes. Noise can present itself not only in the form of continuous deformations, but also as topological changes. It is well known that persistent Betti numbers are a shape descriptor that admits dissimilarity distances stable under continuous shape deformations. In this paper we focus on the problem of dealing with noise that alters the topology of the studied objects. We present a general method to turn persistent Betti numbers into stable descriptors also in the presence of topological changes. Retrieval tests on the Kimia-99 database show the effectiveness of the method.

Keywords: Multidimensional persistent homology, Hausdorff distance, symmetric difference distance
2010 MSC: Primary: 55N35, Secondary: 68T10, 68U05, 65D18, 55-04

1. Introduction

Persistence is a theory for studying objects related to computer vision and computer graphics, by adopting different functions (e.g., distance from the center of mass, distance from the medial axis, height, geodesic distance, color mapping) to measure the shape properties of the object under study (e.g., roundness, elongation, bumpiness, color). The object, considered as a topological space, is explored through the sequence of nested sub-level sets of the considered measuring function. A shape descriptor, called a persistent homology group, can be constructed by encoding at which scale a topological feature (e.g., a connected component, a tunnel, a void) is created, and when it is annihilated along this filtration. For application purposes, these groups are further encoded by considering only their dimension, yielding a parametrized version of Betti numbers, known in the literature as *persistent Betti numbers* [1], a *rank invariant* [2], and, for the 0th homology, a *size function* [3].

In the literature, a large number of methods for shape matching has been proposed, such has the shape-context [4], the shock

Email addresses:
patrizio.frosini@unibo.it (Patrizio Frosini), claudia.landi@unimore.it (Claudia Landi)

graph [5], and the inner distance [6], to name a few. Persistent Betti numbers are shape descriptors belonging to the class of shape-from-functions methods which are widely reviewed in [7].

The stability of persistent Betti numbers (hereafter, PBNs, for brevity) is quite an important issue, every data measurement being affected by noise. The stability problem involves both stability under perturbations of the topological space that represents the object, and stability under perturbations of the function that measures the shape properties of the object.

The problem of stability with respect to perturbations of the measuring function has been studied in [8] for scalar-valued measuring functions. For vector-valued measuring functions, the multidimensional matching distance between PBNs is introduced in [9], and is shown to provide stability in [10]. For the case of 0th homology, this problem is treated in [11] and [12] for scalar- and vector-valued functions, respectively.

In this paper we consider the problem of stability of PBNs with respect to changes of the topological space. This topic has been studied in [13] for sub-level sets of smooth functions satisfying certain conditions on the norm of the gradient. Unfortunately these conditions seem not to be satisfied in a wide variety of situations common in object recognition, such as point cloud data, curves in the plane, domains affected by salt & pepper noise.

We propose a general approach to the problem of stability of PBNs with respect to domain perturbations that applies to more general domains, i.e. compact subsets of $\mathbb{R}^n$. Moreover, according to the type of noise affecting the data, we propose to choose an appropriate set distance to measure the domain perturbation (for example, the Hausdorff distance in case of small position errors, the symmetric difference pseudo-distance in the presence of outliers). The core of our approach is to choose an appropriate continuous function to represent the domain, so that the problem of stability for noisy domains with respect to a given set distance can be reduced to that of stability with respect to changes of the functions. This is achieved by substituting the domain K with an appropriate function f_K defined on a fixed set X containing K. Assuming we were interested in the shape of K, as seen through the restriction to K of a measuring function $\vec{\varphi} : X \to \mathbb{R}^k$, we actually study the function $\vec{\Phi} : X \to \mathbb{R}^{k+1}$, with $\vec{\Phi} = (f_K, \vec{\varphi})$. Persistent Betti numbers of $\vec{\Phi}$ can be compared using the multidimensional matching distance, thus obtaining robustness of PBNs under domain perturbations.

In particular, we use this strategy when sets are compared by the Hausdorff distance and by the symmetric difference pseudo-distance. In both these cases we show stability results (Theorems 4.1 and 4.3). Moreover we show the relation existing between the shape of K as described by $\vec{\varphi}_{|K}$ and the shape described by $\vec{\Phi} = (f_K, \vec{\varphi})$ (Theorem 4.2).

We also consider the situation where sets are described in a fuzzy sense, by means of probability density functions, easily obtaining a stability result also in this case (Proposition 4.4).

Finally, we conclude our paper presenting some experiments in which our method is tested on the Kimia-99 database [14], using as query shapes noisy versions of the original shapes.

2. Preliminaries

2.1. Multidimensional persistent Betti numbers

Persistence may be used to construct shape descriptors that capture both geometrical and topological properties of objects $K \subset \mathbb{R}^n$. Geometrical properties of K are studied through the choice of a function $\vec{\varphi} = (\varphi_i) : K \to \mathbb{R}^k$, each component φ_i describing a shape property. The function $\vec{\varphi}$ is called a k-dimensional *measuring* (or *filtering*) *function*. Topological properties of K as seen through $\vec{\varphi}$ are studied by considering sub-level sets $K\langle\vec{\varphi} \preceq \vec{u}\,\rangle = \{x \in K : \varphi_i(x) \leq u_i,\ i = 1, \dots, k\}$. For $\vec{u} = (u_i), \vec{v} = (v_i) \in \mathbb{R}^k$ with $u_i \leq v_i$ for every index i (briefly, $\vec{u} \preceq \vec{v}$), the sub-level set $K\langle\vec{\varphi} \preceq \vec{u}\,\rangle$ is contained in the sub-level set $K\langle\vec{\varphi} \preceq \vec{v}\,\rangle$. A classical transform of algebraic topology, called homology, provides topological invariants. Working with homology coefficients in a field, it transforms topological spaces into vector spaces, and continuous maps (e.g., inclusions) into linear maps. This leads to the following definition, where the symbol $\vec{u} \prec \vec{v}$ means $u_i < v_i$ for $i = 1, \dots, k$.

Definition 2.1. Let $q \in \mathbb{Z}$. Let $\pi_q^{(\vec{u},\vec{v})} : \check{H}_q(K\langle\vec{\varphi} \preceq \vec{u}\rangle) \to \check{H}_q(K\langle\vec{\varphi} \preceq \vec{v}\rangle)$ be the homomorphism induced by the inclusion map $\pi^{(\vec{u},\vec{v})} : K\langle\vec{\varphi} \preceq \vec{u}\rangle \hookrightarrow K\langle\vec{\varphi} \preceq \vec{v}\rangle$ with $\vec{u} \preceq \vec{v}$, where $\check{H}_q$ denotes the qth Čech homology group. The *qth persistent Betti number function* of $\vec{\varphi}$ is the function $\beta_{\vec{\varphi}} : \{(\vec{u},\vec{v}) \in \mathbb{R}^k \times \mathbb{R}^k : \vec{u} \prec \vec{v}\} \to \mathbb{N} \cup \{\infty\}$ defined as $\beta_{\vec{\varphi}}(\vec{u},\vec{v}) = \dim \operatorname{im} \pi_q^{(\vec{u},\vec{v})}$.

The motivation for using Čech homology is that, unlike ordinary homology theories, it has the continuity axiom (cf. [15, Ch. X]). This will be important when we want to obtain information on homology groups by passing to the limit as in Theorem 4.2.

If K is a triangulable space embedded in some $\mathbb{R}^n$, then $\beta_{\vec{\varphi}}(\vec{u},\vec{v}) < +\infty$, for every $\vec{u} \prec \vec{v}$ and every $q \in \mathbb{Z}$ [16]. We point out that, in our setting, the finiteness of PBNs would not be guaranteed if they were defined on the set $\{(\vec{u},\vec{v}) \in \mathbb{R}^k \times \mathbb{R}^k : \vec{u} \preceq \vec{v}\}$ instead of $\Delta^+ = \{(\vec{u},\vec{v}) \in \mathbb{R}^k \times \mathbb{R}^k : \vec{u} \prec \vec{v}\}$. This motivates our choice of working only on Δ^+.

We refer to PBNs of functions taking values in $\mathbb{R}^k$ with $k > 1$ as to *multidimensional PBNs*, whereas PBNs of functions taking values in $\mathbb{R}$ are called *one-dimensional PBNs*. However, we simply use the term PBNs when it does not generate ambiguities.

The use of multidimensional PBNs instead of one-dimensional ones is crucial for the method presented here because adding the function f_K to the measuring functions makes the dimensionality increase from k to $k+1$, that is always greater than 1.

2.2. Comparison of sets

The problems of description and comparison of sets can be dealt with in a myriad of different ways, each one more or less suitable than another for a given application task.

In classical set theory, the membership of elements in a set is assessed in binary terms according to a bivalent condition – an element either belongs or does not belong to the set. By contrast, in fuzzy set theory [17, 18], a fuzzy set A in X is characterized by a membership function $f_A : X \to [0, 1]$, with the value $f_A(x)$ representing the grade of membership of x in A. Usually, the nearer the value of $f_A(x)$ to 1, the higher the grade of membership of x in A. The fuzzy set theory can be used in a wide range of domains in which information is incomplete or imprecise.

If classical set theory is adopted, then a number of different dissimilarity measures exist to compare two sets [19, 20]. A frequently used dissimilarity measure is the *Hausdorff distance*, which is defined for arbitrary non-empty compact subsets K_1, K_2 of $\mathbb{R}^n$. Let us assume that K_1, K_2 are contained in a compact subset X of $\mathbb{R}^n$, and, for a compact subset K of X, let us denote by d_K the distance to K, that is the function $d_K : X \to \mathbb{R}$ defined by $d_K(x) = \min_{y\in K} \|x - y\|$, $\|\cdot\|$ being any norm on $\mathbb{R}^n$ (e.g., the Euclidean norm). The Hausdorff distance can be defined by

$$\delta_H(K_1, K_2) = \max\{\max_{x\in K_2} d_{K_1}(x), \max_{y\in K_1} d_{K_2}(y)\}.$$

This can be reformulated as follows (cf. [21, Ch. 4, Sect. 2.2]):

$$\begin{aligned}\delta_H(K_1, K_2) &= \max_{x\in X} |d_{K_1}(x) - d_{K_2}(x)| \\ &= \|d_{K_1} - d_{K_2}\|_\infty. \qquad (1)\end{aligned}$$

The Hausdorff distance is robust against small deformations, but it is sensitive to outliers: a single far-away noise point drastically increases the Hausdorff distance. For example, with respect to the Hausdorff distance, the sets in Figure 1(a), (b), and (c) are similar to each other, whereas they are very dissimilar from the set in Figure 1 (d).

A dissimilarity measure that is based on the volume of the symmetric difference, such as the symmetric difference pseudo-metric, overcomes the problem of outliers. Denoting by μ the Lebesgue measure on $\mathbb{R}^n$, the *symmetric difference pseudo-metric* $d_\triangle$ is defined between two measurable sets A, B with finite measure by $d_\triangle(A, B) = \mu(A\triangle B)$ where $A\triangle B = (A\cup B)\setminus(A\cap B)$ is the symmetric difference of A and B. It holds that $d_\triangle(A, B) = 0$ if and only if A and B are equal almost everywhere. Identifying two sets A and B if $\mu(A\triangle B) = 0$, we obtain the symmetric difference metric.

Other dissimilarity measures are, for example, the bottleneck distance between finite point sets of the same cardinality, and the L_p-Hausdorff distance. However, since many other distances could be considered, we will limit our research to consider stability with respect to the Hausdorff and symmetric difference distances.

When fuzzy sets are used, their dissimilarity can be measured by any distance between functions. In this case we will confine ourselves to consider the max-norm distance between fuzzy sets.

3. Working assumptions

We will model objects under study as subsets K of some compact domain $X \subseteq \mathbb{R}^n$. Shape properties of the objects under study will be described by vector-valued functions $\vec{\varphi} : X \to \mathbb{R}^k$, with the measuring function $\vec{\varphi}$ defined on the entire ambient space X because the domain K will vary.

We think that this setting, which is used also in [13], is not very restrictive. Although apparently it prevents one from using filtering functions intrinsic of the domains K and K', and very common in applications (e.g., the distance from the center of mass, or the geodesic distance from a point), from a theoretical point of view this is not the case. Indeed, the well known Tietze's extension theorem states that if X is normal, K is a closed subset of X, and $f : K \to \mathbb{R}$ is a continuous function, then there is a continuous function $F : X \to \mathbb{R}$ such that $F_{|K} = f$ [22].

Noise on the measuring function $\vec{\varphi}$ will always be quantified using the max-norm,

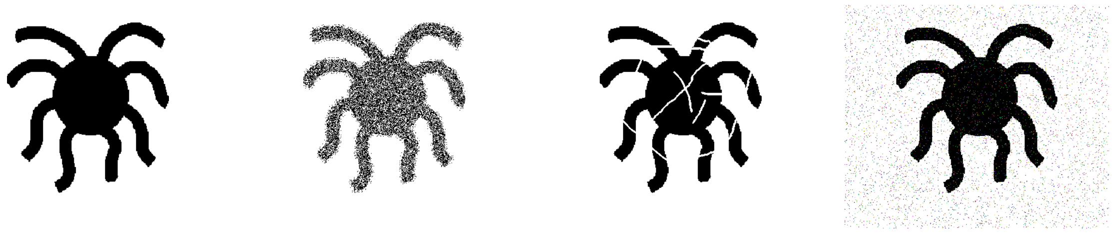

(a) (b) (c) (d)

Figure 1: Four binary images of an octopus. (b), (c), and (d) are noisy versions of (a).

as is standard (cf. [8, 12]): $\left\|\vec{\varphi} - \vec{\psi}\right\|_\infty = \max_{x \in X} \max_{1 \le i \le k} |\varphi_i(x) - \psi_i(x)|$.

On the other hand, since there is no standard way to measure the noise on the domain K, we confine our study to a selection of possible distances (the Hausdorff distance and the symmetric difference distance for classical sets, and the max-norm distance for fuzzy sets), indicating how our results can be adapted in other situations.

In order to measure perturbations on PBNs, we assume we are given a distance D on $\{\beta_{\vec{\varphi}} \,|\, \vec{\varphi} : X \to \mathbb{R}^k \text{ continuous}\}$ having the following property:

(S) $D(\beta_{\vec{\varphi}}, \beta_{\vec{\psi}}) \le \left\|\vec{\varphi} - \vec{\psi}\right\|_\infty$.

Property (S) will be called the Stability Property. Distances with this property exist: the multidimensional matching distance D_{match} introduced in [9] and proven to have Property (S) in [10] (see also [12] for the case $q = 0$), and the one induced by the interleaving distance presented in [23]. The reader can refer to Appendix B for details on D_{match}.

4. Stability of PBNs with respect to noisy domains

Our method to achieve stability of PBNs with respect to changes of the topological space K, even under perturbations that change its topology, is to consider K embedded in a larger space X in which K and its noisy version are similar with respect to some metric.

Next we substitute the set K with an appropriate function f_K defined on X, so that the perturbation of the set K becomes a perturbation of the function f_K. As a consequence, instead of studying the shape of K as seen through a measuring function $\vec{\varphi}_{|K} : K \to \mathbb{R}^k$, we study a new measuring function $\vec{\Phi} : X \to \mathbb{R}^{k+1}$, with $\vec{\Phi} = (f_K, \vec{\varphi})$. PBNs of $\vec{\Phi}$ can be compared using the distance D in a stable way, as a consequence of the Stability Property (S) for measuring function perturbations. The key issue here is that we can prove that the PBNs of $\vec{\Phi}$ are still descriptors of the shape of K.

4.1. Stability with respect to Hausdorff distance

In order to achieve stability under set perturbations that are measured by the Hausdorff distance, we can take the function f_K equal to the distance from K as the following result shows. In some sense this smooths spaces by a uniform thickening. A related method to do this in the applied algebraic topology literature is through the Rips filtration.

Theorem 4.1. *Let K_1, K_2 be non-empty closed subsets of a triangulable subspace X*

of $\mathbb{R}^n$. Let $d_{K_1}, d_{K_2} : X \to \mathbb{R}$ be their respective distance functions. Moreover, let $\vec{\varphi}_1, \vec{\varphi}_2 : X \to \mathbb{R}^k$ be vector-valued continuous functions. Then, defining $\vec{\Phi}_1, \vec{\Phi}_2 : X \to \mathbb{R}^{k+1}$ by $\vec{\Phi}_1 = (d_{K_1}, \vec{\varphi}_1)$ and $\vec{\Phi}_2 = (d_{K_2}, \vec{\varphi}_2)$, the following inequality holds:

$$D\left(\beta_{\vec{\Phi}_1}, \beta_{\vec{\Phi}_2}\right) \leq \max\{\delta_H(K_1, K_2), \|\vec{\varphi}_1 - \vec{\varphi}_2\|_\infty\}.$$

Proof. The Stability Property (S) implies that $D\left(\beta_{\vec{\Phi}_1}, \beta_{\vec{\Phi}_2}\right) \leq \|\vec{\Phi}_1 - \vec{\Phi}_2\|_\infty$. It follows that

$$D\left(\beta_{\vec{\Phi}_1}, \beta_{\vec{\Phi}_2}\right) \leq \max\left\{\|d_{K_1} - d_{K_2}\|_\infty, \|\vec{\varphi}_1 - \vec{\varphi}_2\|_\infty\right\}.$$

Hence, by equality (1), the claim is proved. □

In plain words, Theorem 4.1 states that small changes in the domain and in the measuring function imply small changes in the PBNs, i.e. in the shape descriptors. Clearly, the inequality of Theorem 4.1 tends to the classic bottleneck stability inequality [8] when K_1 tends to K_2 with respect to the Hausdorff distance.

An example illustrating Theorem 4.1 is shown in Figure 2. Figure 2(a), (b), and (c) show the 0th PBNs of the sets of black pixels K_1, K_2, K_3 of Figure 1(a), (b), and (c), respectively, with the measuring function equal to minus the distance from the center of mass of K_1. The PBNs $\beta_{\varphi_{|K_1}}$ displays eight relevant points in the persistence diagram, corresponding to the eight tentacles of the octopus. Only one of these points is at infinity (and therefore depicted by a vertical line rather than by a circle) since K_1 has only one connected component. As for $\beta_{\varphi_{|K_2}}$, due to the presence of a great quantity of connected components in the noisy octopus, its PBNs display a very large number of points at infinity, and a figure showing them all would be hardly readable. For this reason we show only a small subset of its persistence diagram. Finally, $\beta_{\varphi_{|K_3}}$, due to the presence of 11 connected components in Figure 1(c), shows 11 points at infinity. Figure 2 (a'), (b'), and (c') show a 2-dimensional slice of the 0th PBNs $\beta_{(d_{K_i}, \varphi)}$, $i = 1, 2, 3$ (more details on how multidimensional PBNs can be studied by slicing their domain can be found in Appendix A). It is easily perceivable how similar Figure 2(a'), (b'), and (c') are to each other, especially in contrast to the dissimilarity between (a), (b), and (c).

The key point of our approach is that the PBNs of $\vec{\Phi}$ still provide a shape descriptor for K as seen through $\vec{\varphi}_{|K}$. This fact is shown by the next result.

Theorem 4.2. *Let K be a non-empty triangulable subset of a triangulable subspace X of $\mathbb{R}^n$. Moreover, let $\vec{\varphi} : X \to \mathbb{R}^k$ be a continuous function. Setting $\vec{\Phi} : X \to \mathbb{R}^{k+1}$, $\vec{\Phi} = (d_K, \vec{\varphi})$, for every $\vec{u}, \vec{v} \in \mathbb{R}^k$ with $\vec{u} \prec \vec{v}$, there exists a real number $\hat{\eta} > 0$ such that, for any $\eta \in \mathbb{R}$ with $0 < \eta \leq \hat{\eta}$, there exists a real number $\hat{\varepsilon} = \hat{\varepsilon}(\eta)$, with $0 < \hat{\varepsilon} < \eta$, for which*

$$\beta_{\vec{\varphi}_{|K}}(\vec{u}, \vec{v}) = \beta_{\vec{\Phi}}\left((\varepsilon, \vec{u}), (\eta, \vec{v})\right),$$

for every $\varepsilon \in \mathbb{R}$ with $0 \leq \varepsilon \leq \hat{\varepsilon}$. In particular,

$$\beta_{\vec{\varphi}_{|K}}(\vec{u}, \vec{v}) = \lim_{\eta \to 0^+} \beta_{\vec{\Phi}}\left((0, \vec{u}), (\eta, \vec{v})\right).$$

Proof. For every $\vec{u} \in \mathbb{R}^k$, we have

$$\begin{aligned} K\langle \vec{\varphi}_{|K} \preceq \vec{u} \rangle &= \{x \in K : \vec{\varphi}(x) \preceq \vec{u}\} \\ &= \{x \in X : d_K(x) \leq 0\} \\ &\quad \cap \{x \in X : \vec{\varphi}(x) \preceq \vec{u}\} \\ &= \{x \in X : \vec{\Phi}(x) \preceq (0, \vec{u})\} \\ &= X\langle \vec{\Phi} \preceq (0, \vec{u}) \rangle. \end{aligned}$$

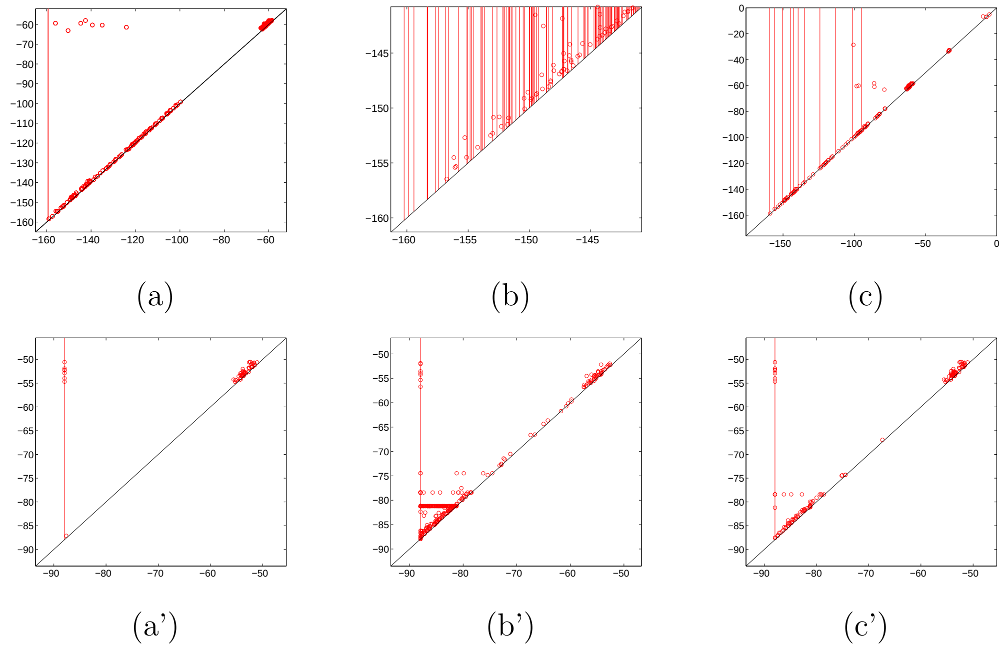

Figure 2: (a-c): The 0th PBNs $\beta_{\varphi_{|K_i}}$, $i = 1, 2, 3$, with K_1 the original octopus image, and K_2, K_3 two noisy versions, with respect to the same measuring function φ. (a'-c'): The 0th PBNs $\beta_{(d_{K_i},\varphi)}$, with K_i and φ as before (slice).

Hence, for every $q \in \mathbb{Z}$, denoting by $\pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})}$ the homology homomorphism induced by the inclusion $X\langle\vec{\Phi} \preceq (\varepsilon, \vec{u})\rangle \to X\langle\vec{\Phi} \preceq (\eta, \vec{v})\rangle$, with $(\varepsilon, \vec{u}) \preceq (\eta, \vec{v})$, it holds that

$$\beta_{\vec{\varphi}_{|K}}(\vec{u}, \vec{v}) = \dim \operatorname{im} \pi_q^{(0,\vec{u}),(0,\vec{v})}.$$

We claim that there exists a positive real number $\hat{\eta}$ such that

$$\operatorname{im} \pi_q^{(0,\vec{u}),(0,\vec{v})} \cong \operatorname{im} \pi_q^{(0,\vec{u}),(\eta,\vec{v})}$$

for every η with $0 < \eta \leq \hat{\eta}$ (the claim is trivial for $\eta = 0$). In particular, this fact proves that $\beta_{\vec{\varphi}_{|K}}(\vec{u}, \vec{v}) = \lim_{\eta\to 0^+} \beta_{\vec{\Phi}}\left((0, \vec{u}), (\eta, \vec{v})\right)$.

In order to prove this claim, we consider the inverse system of homomorphisms $\pi_q^{(0,\vec{u}),(\eta,\vec{v})} : \check{H}_q(X\langle\vec{\Phi} \preceq (0, \vec{u})\rangle) \to \check{H}_q(X\langle\vec{\Phi} \preceq (\eta, \vec{v})\rangle)$ over the directed set $\{\eta \in \mathbb{R} : \eta > 0\}$ decreasingly ordered. The following isomorphisms hold:

$$\begin{aligned}\operatorname{im} \pi_q^{(0,\vec{u}),(0,\vec{v})} &\cong \operatorname{im} \varprojlim \pi_q^{(0,\vec{u}),(\eta,\vec{v})} \\ &\cong \varprojlim \operatorname{im} \pi_q^{(0,\vec{u}),(\eta,\vec{v})}.\end{aligned}$$

Indeed, $\operatorname{im} \pi_q^{(0,\vec{u}),(0,\vec{v})} \cong \operatorname{im} \varprojlim \pi_q^{(0,\vec{u}),(\eta,\vec{v})}$ by the continuity of Čech homology, and $\operatorname{im} \varprojlim \pi_q^{(0,\vec{u}),(\eta,\vec{v})} \cong \varprojlim \operatorname{im} \pi_q^{(0,\vec{u}),(\eta,\vec{v})}$ because the inverse limit of vector spaces is an exact functor and therefore it preserves epimorphisms, and hence images.

It remains to prove that there exists a positive real number $\hat{\eta}$ such that, for every $0 < \eta \leq \hat{\eta}$, $\operatorname{im} \pi_q^{(0,\vec{u}),(\eta,\vec{v})}$ is isomorphic to $\varprojlim \operatorname{im} \pi_q^{(0,\vec{u}),(\eta,\vec{v})}$. To this end, let us consider the following commutative diagram, with $0 < \eta' \leq \eta''$:

$$\begin{tikzcd}
\check{H}_q(X\langle\vec{\Phi}\preceq(0,\vec{u})\rangle) \arrow[r, "id"] \arrow[d, "\pi_q^{(0,\vec{u}),(\eta',\vec{v})}"'] & \check{H}_q(X\langle\vec{\Phi}\preceq(0,\vec{u})\rangle) \arrow[d, "\pi_q^{(0,\vec{u}),(\eta'',\vec{v})}"] \\
\check{H}_q(X\langle\vec{\Phi}\preceq(\eta',\vec{v})\rangle) \arrow[r, "\pi_q^{(\eta',\vec{v}),(\eta'',\vec{v})}"] & \check{H}_q(X\langle\vec{\Phi}\preceq(\eta'',\vec{v})\rangle).
\end{tikzcd} \tag{2}$$

From the above diagram (2), we see that each $\pi_q^{(\eta',\vec{v}),(\eta'',\vec{v})}$ induces a map $\tau_q^{(\eta',\eta'')}$: $\operatorname{im}\pi_q^{(0,\vec{u}),(\eta',\vec{v})} \to \operatorname{im}\pi_q^{(0,\vec{u}),(\eta'',\vec{v})}$. From diagram (2) we see that these maps are surjective. On the other hand, by the finiteness of the dimension of $\operatorname{im}\pi_q^{(0,\vec{u}),(0,\vec{v})}$ and the monotonicity of PBNs, there exists $\hat{\eta} > 0$ such that the dimension of $\pi_q^{(0,\vec{u}),(\eta',\vec{v})}$ is finite and equal to the dimension of $\pi_q^{(0,\vec{u}),(\eta'',\vec{v})}$, whenever $0 < \eta' \leq \eta'' \leq \hat{\eta}$. Hence the maps $\tau_q^{(\eta',\eta'')}$ are surjections between vector spaces of the same finite dimension, i.e. isomorphisms for every $0 < \eta' \leq \eta'' \leq \hat{\eta}$. Thus, $\varprojlim \operatorname{im}\pi_q^{(0,\vec{u}),(\eta,\vec{v})}$ is the inverse limit of a system of finite dimensional vector spaces isomorphic to $\operatorname{im}\pi_q^{(0,\vec{u}),(\hat{\eta},\vec{v})}$, proving the claim.

We now claim that for every strictly positive real number η, there exists a strictly positive real number $\hat{\varepsilon} < \eta$ such that

$$\operatorname{im}\pi_q^{(0,\vec{u}),(\eta,\vec{v})} \cong \operatorname{im}\pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})}$$

for every ε with $0 \leq \varepsilon \leq \hat{\varepsilon}$.

This claim can be proved in much the same way as the previous one. We consider the inverse system of homomorphisms $\pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})} : \check{H}_q(X\langle\vec{\Phi} \preceq (\varepsilon,\vec{u})\rangle) \to \check{H}_q(X\langle\vec{\Phi} \preceq (\eta,\vec{v})\rangle)$ over the directed set $\{\varepsilon \in \mathbb{R} : 0 \leq \varepsilon < \eta\}$ decreasingly ordered. The following isomorphisms follow again from the continuity of Čech homology and the exactness of the inverse limit functor for vector spaces:

$$\begin{aligned}
\operatorname{im}\pi_q^{(0,\vec{u}),(\eta,\vec{v})} &\cong \operatorname{im}\varprojlim \pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})} \\
&\cong \varprojlim \operatorname{im}\pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})}.
\end{aligned}$$

To prove that there exists a strictly positive real number $\hat{\varepsilon}$ such that, for every $0 \leq \varepsilon \leq \hat{\varepsilon}$, $\operatorname{im}\pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})}$ is isomorphic to $\varprojlim \operatorname{im}\pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})}$, let us consider the following commutative diagram, with $0 \leq \varepsilon' \leq \varepsilon''$:

$$\begin{tikzcd}
\check{H}_q(X\langle\vec{\Phi}\preceq(\varepsilon',\vec{u})\rangle) \arrow[r, "\pi_q^{(\varepsilon',\vec{u}),(\varepsilon'',\vec{u})}"] \arrow[d, "\pi_q^{(\varepsilon',\vec{u}),(\eta,\vec{v})}"'] & \check{H}_q(X\langle\vec{\Phi}\preceq(\varepsilon'',\vec{u})\rangle) \arrow[d, "\pi_q^{(\varepsilon'',\vec{u}),(\eta,\vec{v})}"] \\
\check{H}_q(X\langle\vec{\Phi}\preceq(\eta,\vec{v})\rangle) \arrow[r, "id"] & \check{H}_q(X\langle\vec{\Phi}\preceq(\eta,\vec{v})\rangle).
\end{tikzcd} \tag{3}$$

From the above diagram (3), we see that each $\pi_q^{(\varepsilon',\vec{u}),(\varepsilon'',\vec{u})}$ induces a map $\sigma_q^{(\varepsilon',\varepsilon'')}$: $\operatorname{im}\pi_q^{(\varepsilon',\vec{u}),(\eta,\vec{v})} \to \operatorname{im}\pi_q^{(\varepsilon'',\vec{u}),(\eta,\vec{v})}$. From diagram (3) we see that these maps are injective. On the other hand, by the finiteness of the dimension of $\operatorname{im}\pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})}$, for any ε with $0 < \varepsilon < \eta$, and the monotonicity of PBNs, there exists $\hat{\varepsilon}$, with $0 < \hat{\varepsilon} < \eta$, such that the dimension of $\pi_q^{(\varepsilon',\vec{u}),(\eta,\vec{v})}$ is finite and equal to the dimension of $\pi_q^{(\varepsilon'',\vec{u}),(\eta,\vec{v})}$, whenever $0 \leq \varepsilon' \leq \varepsilon'' \leq \hat{\varepsilon}$. Hence the maps $\sigma_q^{(\varepsilon',\varepsilon'')}$ are injections between vector spaces of the same finite dimension, i.e. isomorphisms for every $0 \leq \varepsilon' \leq \varepsilon'' \leq \hat{\varepsilon}$. Thus, $\varprojlim \operatorname{im}\pi_q^{(\varepsilon,\vec{u}),(\eta,\vec{v})}$ is the inverse limit of a system of finite dimensional vector spaces isomorphic to $\operatorname{im}\pi_q^{(\hat{\varepsilon},\vec{u}),(\hat{\eta},\vec{v})}$, proving the claim. □

In other words, Theorem 4.2 ensures that we can recover the PBNs of $\vec{\varphi}_{|K}$, i.e. a description of the shape of K as seen by $\vec{\varphi}$, from the PBNs of $\vec{\Phi}$, simply by passing to the limit.

To illustrate this result we have considered again the octopus in Figure 1(a), together with the measuring function $\varphi : X \to \mathbb{R}$ equal to minus the distance from the center of mass of the set of black pixels K_1. Here X is the bounding box

Table 1: The values taken by the 0th PBNs $\beta_{(d_{K_1},\varphi)}$ of the octopus image of Figure 1(a) at $((\varepsilon,-100),(\eta,-80))$, as ε and η tend to 0, tend to the value $\beta_{\varphi_{|K_1}}(-100,-80)=8$.

ε	η	$\beta_{(d_{K_1},\varphi)}\left((\varepsilon,-100),(\eta,-80)\right)$
0.5	24	1
0.5	8	3
0.5	1	8
0.5	0.65	8
0.3	0.45	8
0.1	0.25	8

of the image, thus containing black and white pixels. We compared the value taken by the 0th PBNs $\beta_{\varphi_{|K_1}}$ at the point $(u,v)=(-100,-80)$, that is 8, with the value obtainable from $\beta_{(d_{K_1},\varphi)}$ by passing to the limit as in Theorem 4.2. Computations show that $\beta_{\varphi_{|K_1}}(-100,-80)=\beta_{(d_{K_1},\varphi)}\left((\varepsilon,-100),(\eta,-80)\right)=8$ for small but positive values of ε and η. The PBNs $\beta_{(d_{K_1},\varphi)}\left((\varepsilon,-100),(\eta,-80)\right)$ for the choices of ε and η considered in Table 1 are displayed in Figure 3 via a restriction to an appropriate slice of Δ^+.

4.2. Stability with respect to the symmetric difference pseudo-distance

In order to achieve stability under set perturbations that are measured by the symmetric difference pseudo-distance, we can take the function f_K to be the convolution of the characteristic function of K with that of a ball. More precisely, let $\lambda_K^\epsilon:\mathbb{R}^n\to\mathbb{R}$, with $\epsilon\in\mathbb{R}$, $\epsilon>0$, be defined as

$$\lambda_K^\epsilon(x)=\mu(B_\epsilon)^{-1}\cdot\int_{y\in B_\epsilon(x)}\chi_K(y)\ \mathrm{d}y$$

where $B_\epsilon(x)$ denotes the n-ball centered at x with radius ϵ, $B_\epsilon=B_\epsilon(0)$, and χ_K denotes the characteristic function of K. In other words, in this case, we smooth sets by convolving with the characteristic function of a ball.

In this case we have the following stability result.

Theorem 4.3. *Let K_1,K_2 be non-empty closed subsets of a triangulable subspace X of $\mathbb{R}^n$. Moreover, let $\vec{\varphi}_1,\vec{\varphi}_2:X\to\mathbb{R}^k$ be vector-valued continuous functions. Then, defining $\vec{\Psi}_1^\epsilon,\vec{\Psi}_2^\epsilon:X\to\mathbb{R}^{k+1}$ by $\vec{\Psi}_1^\epsilon=(-\lambda_{K_1}^\epsilon,\vec{\varphi}_1)$ and $\vec{\Psi}_2^\epsilon=(-\lambda_{K_2}^\epsilon,\vec{\varphi}_2)$, the following inequality holds:*

$$D\left(\beta_{\vec{\Psi}_1^\epsilon},\beta_{\vec{\Psi}_2^\epsilon}\right)\le\max\left\{\frac{d_\triangle(K_1,K_2)}{\mu(B_\epsilon)},\|\vec{\varphi}_1-\vec{\varphi}_2\|_\infty\right\}. \quad (4)$$

Proof. For every $x\in X$,

$$\begin{aligned}
&|\lambda_{K_1}^\varepsilon(x)-\lambda_{K_2}^\varepsilon(x)|\\
&\quad=\ \mu(B_\varepsilon)^{-1}\cdot\left|\int_{y\in B_\varepsilon(x)}\chi_{K_1}(y)-\chi_{K_2}(y)\ \mathrm{d}y\right|\\
&\quad\le\ \mu(B_\varepsilon)^{-1}\cdot\int_X|\chi_{K_1}(y)-\chi_{K_2}(y)|\ \mathrm{d}y\\
&\quad=\ \mu(B_\varepsilon)^{-1}\cdot\mu(K_1\triangle K_2).
\end{aligned}$$

Thus $\|\lambda_{K_1}^\varepsilon-\lambda_{K_2}^\varepsilon\|_\infty\le\mu(B_\varepsilon)^{-1}\cdot\mu(K_1\triangle K_2)$. The Stability Property (S) for measuring function perturbations guarantees that

$$D\left(\beta_{\vec{\Psi}_1^\varepsilon},\beta_{\vec{\Psi}_2^\varepsilon}\right)\le\|\vec{\Psi}_1^\varepsilon-\vec{\Psi}_2^\varepsilon\|_\infty.$$

It follows that

$$\begin{aligned}
&D\left(\beta_{\vec{\Psi}_1^\varepsilon},\beta_{\vec{\Psi}_2^\varepsilon}\right)\\
&\quad\le\ \max\left\{\|\lambda_{K_1}^\varepsilon-\lambda_{K_2}^\varepsilon\|_\infty,\|\vec{\varphi}_1-\vec{\varphi}_2\|_\infty\right\}\\
&\quad\le\ \max\left\{\mu(B_\varepsilon)^{-1}\cdot\mu(K_1\triangle K_2),\|\vec{\varphi}_1-\vec{\varphi}_2\|_\infty\right\}\\
&\quad=\ \max\left\{\mu(B_\varepsilon)^{-1}\cdot d_\triangle(K_1,K_2),\|\vec{\varphi}_1-\vec{\varphi}_2\|_\infty\right\}.
\end{aligned}$$

□

The previous theorem shows that, under our hypotheses, if two compact subsets K_1,K_2 of the real plane are close to each other in the sense that their symmetric difference has a small measure, and $\vec{\varphi}_1$ is close to $\vec{\varphi}_2$, then also the PBNs constructed using the functions $\vec{\Psi}_1^\varepsilon$, $\vec{\Psi}_2^\varepsilon$ are close to each other.

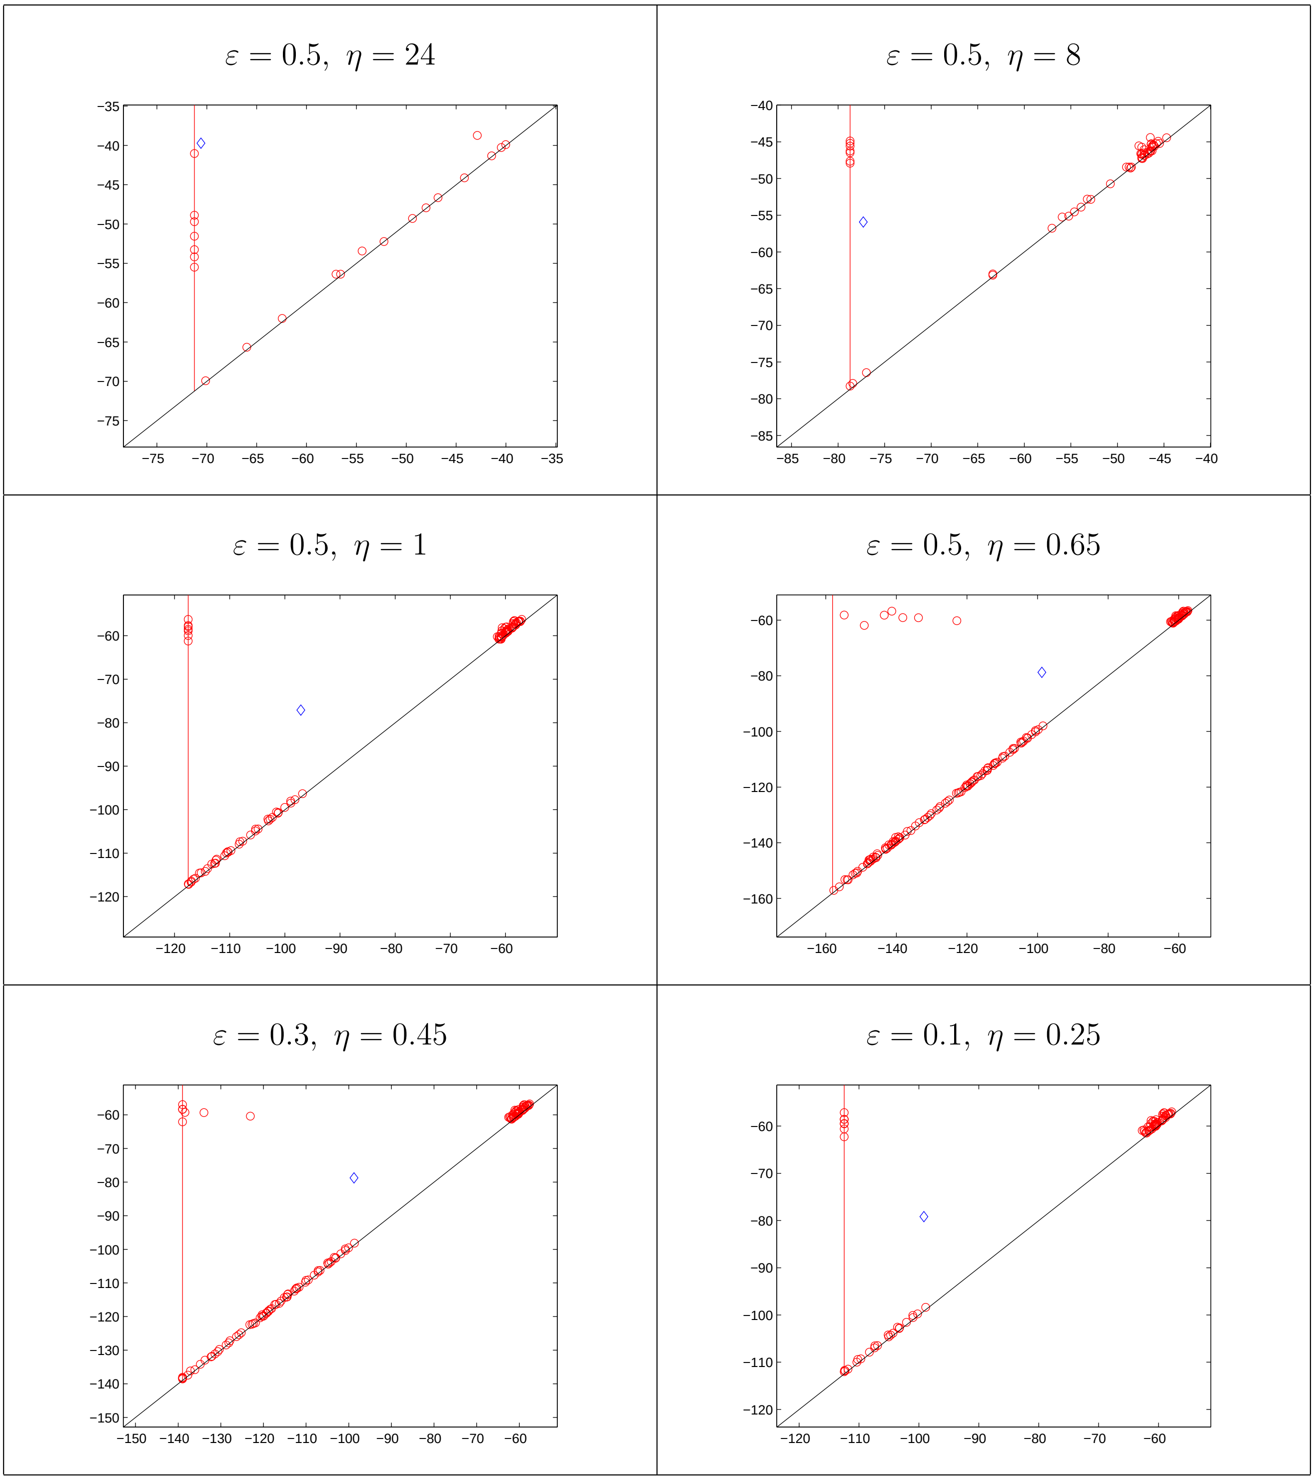

Figure 3: The 0th PBNs $\beta_{(d_{K_1},\varphi)}$ of the octopus image of Figure 1(a) as ε and η tend to 0, restricted to an appropriate 2-dimensional slice of Δ^+. Red circles and red lines denote the points (proper or at infinity) of the corresponding persistence diagram. The blue diamond denotes the point corresponding to $((\varepsilon, -100), (\eta, -80))$.

We observe that, for any $x \in X$,

$$\left|\int_{y\in B_\varepsilon(x)} \chi_{K_1}(y) - \chi_{K_2}(y)\ \mathrm{d}y\right|$$
$$\leq \int_{y\in B_\varepsilon(x)} |\chi_{K_1}(y) - \chi_{K_2}(y)|\ \mathrm{d}y$$
$$= \mu\left((K_1 \triangle K_2) \cap B_\epsilon(x)\right) \leq \mu(K_1 \triangle K_2)$$
$$= d_\triangle(K_1, K_2).$$

Moreover, since $\mu\left((K_1 \triangle K_2) \cap B_\epsilon(x)\right) \leq \mu(B_\epsilon(x))$, if $\mu(B_\epsilon) < \mu(K_1 \triangle K_2)$, then $\max_{x\in\mathbb{R}^n}\left|\int_{y\in B_\varepsilon(x)} \chi_{K_1}(y) - \chi_{K_2}(y)\ \mathrm{d}y\right| < d_\triangle(K_1, K_2)$. Therefore the estimate in inequality (4) can be improved by substituting $d_\triangle(K_1, K_2)$ with $\max_{x\in\mathbb{R}^n}\left|\int_{y\in B_\varepsilon(x)} \chi_{K_1}(y) - \chi_{K_2}(y)\ \mathrm{d}y\right|$.

4.3. Stability with respect to perturbations of fuzzy sets

Now we consider the case when sets are defined according to fuzzy theory, that is through functions representing the grade of membership of points to the considered set. One obtains a fuzzy set, for example, when a probability density $p(x)$ is given, $p(x)$ expressing the probability that a point of the considered set belongs to an infinitesimal neighborhood of x. We confine ourselves to considering only probability densities with compact support contained in a triangulable subspace X of $\mathbb{R}^n$. From the Stability Property (S) for measuring function perturbations we immediately deduce the following result, whose simple proof is omitted, concerning the stability with respect to perturbations of fuzzy sets defined by probability densities.

Proposition 4.4. *Let p_1, p_2 be two probability density functions having support contained in a compact and triangulable subspace X of $\mathbb{R}^n$. Defining $\vec{\Psi}_1, \vec{\Psi}_2 : X \to \mathbb{R}^{k+1}$ by $\vec{\Psi}_1 = (-p_1, \vec{\varphi}_1)$ and $\vec{\Psi}_2 = (-p_2, \vec{\varphi}_2)$, the following statement holds:*

$$D\left(\beta_{\vec{\Psi}_1}, \beta_{\vec{\Psi}_2}\right) \leq \max\{\|p_1 - p_2\|_\infty, \|\vec{\varphi}_1 - \vec{\varphi}_2\|_\infty\}.$$

5. Experimental results

In this section first we describe the more practical aspects of our method to compare PBNs in a way that is stable against domain perturbations, and next we present some numerical results.

5.1. Practical aspects

In view of the experiments that will be presented in the next subsection, here we illustrate how the method can work in practice on black and white images.

First of all the method requires some choices: the spaces K and X, the function $\vec{\varphi}$, and the distance D.

As for the set K, this is the set that contains the relevant information on the studied object. Generally, in a black and white image, this is the set of black pixels.

The space X must be a sort of ambient space for K, thus in the chosen setting it could be the set of all the pixels in the image. It is easy to make X a triangulable space by using the 8-adjacency relation between pixels. In this way pixels correspond to vertices of a triangulation.

The function $\vec{\varphi} : X \to \mathbb{R}^k$ is only required to be continuous. Therefore, we can assume it is defined on vertices of X, and extend it to other simplices by interpolation to achieve continuity. In practice, we can confine ourselves to compute $\vec{\varphi}$ at each pixel. This topic is widely treated in [24].

The distance D between PBNs can be taken to be the multidimensional matching distance D_{match}. Details on D_{match} can be found in Appendix B. The definition of

D_{match} is based on a foliation method reported in Appendix A. Intuitively, by the foliation method the set Δ^+ is sliced into infinitely many half-planes such that the restriction of multidimensional PBNs to each of these half-planes gives one-dimensional PBNs. Actually, D_{match} can be computed only up to some tolerance error [25] and the computation is very time consuming.

Another possibility is to compute a stable lower bound $\tilde{D}$ of D_{match}. This can be obtained by considering only some half-planes of the foliation, and taking the one-dimensional matching distance d_{match} to compare the one-dimensional PBNs we find on these half-planes, as described in Remark 5 of [12].

After these choices are made, the pipeline of our method to compare PBNs of black and white images in a way that is stable against domain perturbations that are small with respect to the Hausdorff distance consists of the following steps.

Given two black and white images I and I' with set of black pixels K and K', respectively, with bounding box X, and given a function $\vec{\varphi} : X \to \mathbb{R}^k$,

Step 1. Set $\vec{\Phi} = (d_K, \vec{\varphi})$, $\vec{\Phi}' = (d_{K'}, \vec{\varphi})$.

Step 2. Fix a finite set of half-planes π_h associated with the pairs $(\vec{l}^h, \vec{b}^h)$ from the foliation described in Appendix A.

Step 3. For each h, compute the values of $F^h = \min_{1\le i\le k+1} l_i^h \cdot \max_{1\le i\le k+1} \frac{\Phi_i - b_i^h}{l_i^h}$ at each pixel of I, and the values of $F'^h = \min_{1\le i\le k+1} l_i^h \cdot \max_{1\le i\le k+1} \frac{\Phi'_i - b_i^h}{l_i^h}$ at each pixel of I'.

Step 4. For each h, compute the PBNs of F^h and F'^h.

Step 5. Set $\tilde{D}$ equal to the maximum, varying h, of the one-dimensional matching distance d_{match} between the PBNs of F^h and F'^h.

In order to obtain stability against domain perturbations that are small with respect to the symmetric difference distance, it is sufficient to substitute d_K and $d_{K'}$ with $-\lambda_K^\varepsilon$ and $-\lambda_{K'}^\varepsilon$, respectively, in Step 1.

The complexity of computing the PBNs for the image I in Step 4 is $O(h \cdot (n \log n + m \cdot \alpha(2m + n, n)))$ operations, where n is the number of pixels of I, m behaves as $4n$ or $8n$ according to the chosen $4-$ or $8-$neighborhood relation among pixels, and α is the inverse of the Ackermann function. The complexity of computing the one-dimensional matching distance d_{match} between the PBNs of F^h and F'^h, for all h, in Step 5 is $O(h \cdot (p^{2.5} + k))$, with p the number of points in the persistence diagrams (see Section B).

5.2. Experiments

In order to demonstrate the effectiveness of the approach presented here, we performed some tests on the Kimia data set of 99 shapes [14], a selection of which is shown in Table 2. The dataset is classified in nine categories with 11 shapes in each category.

Each of the shapes has been corrupted by adding *salt & pepper* noise to a neighborhood of the set of its black pixels, as shown for some instances in Figure 3(Top). Salt & pepper noise is a form of noise typically seen on images, usually caused by errors in the data transmission. It appears as randomly occurring white and black pixels, the percentage of pixels which are corrupted quantifying the noise. For each image, the set of black pixels of the image obtained by adding salt & pepper noise as in Figure 3(Top) is

Table 2: Some instances from the database of 99 shapes with 9 categories and 11 shapes in each category used in our experiments. The complete database can be found in [14].

close to the set of black pixels of the original image with respect to the symmetric difference distance.

Salt & pepper noise can be partially removed by applying a morphological opening, thus obtaining shapes such as those in Figure 3(Bottom). The set of black pixels in the images so obtained is close to the set of black pixels of the original image with respect to the Hausdorff distance.

In both cases the topology of the set of black pixels in the noisy images is very different from that of the original images.

Three retrieval tests from the Kimia dataset were performed.

In order to provide a point of reference, the first retrieval test was performed without noise by matching each shape in the Kimia-99 dataset against every other shape in the database.

In the second retrieval test we used as models to be compared with all the shapes of the Kimia-99 database, the 99 images obtained by adding the salt & pepper noise and performing the morphological opening (examples of query shapes are given in Table 3(bottom)).

Finally, in the third experiment, we compared the images corrupted by the salt & pepper noise with all the original images (examples of query shapes are given in Table 3(top)).

In all cases, ideal result would be that the 11 closest matches (including the queried model itself) all be of the same category as the query shape. The actual results we obtained are reported in Table 4. For each experiment, a string of 11 numbers describes the performance rate, the nth number corresponding to the rate at which the nth nearest match was in the same category as the model. This performance test has been applied to retrieval experiments from the Kimia-99 database by several authors testing their methods (see, e.g., [26, 27, 14, 28, 29]). However, our results are not directly comparable with theirs since we aim at a method tolerant under noise that modifies the shape topology.

The results of Table 4 were obtained following the steps described in Subsection 5.1. As for the necessary choices, we proceeded as follows.

In each case we have used only the PBNs of zeroth homology (a.k.a. size functions). As for the choice of $\vec{\varphi} : X \to \mathbb{R}^k$, we have considered three different functions, with $k = 1$: $\varphi_0, \varphi_1, \varphi_2 : X \to \mathbb{R}$, where φ_0 is equal to minus the distance from the centroid of K, and φ_1, φ_2 are equal to minus the distance from the first and second axis of inertia of K, respectively.

In the first experiment, without noise, for each shape we computed three one-dimensional PBNs, corresponding to the functions $\varphi_0, \varphi_1, \varphi_2$ restricted to the set of black pixels K.

In the second experiment, the query shapes were corrupted by noise and partially cleaned by the morphological opening,

Table 3: Top: Shapes with salt & pepper noise. Bottom: The same shapes after morphological opening.

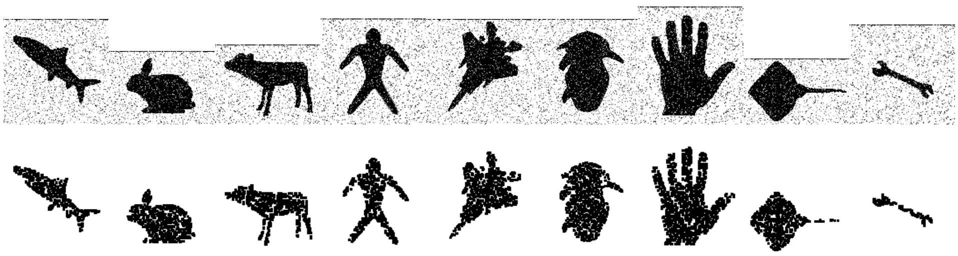

Table 4: The retrieval rates of our method for the Kimia-99 database.

Experiment	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th	11th
without noise	99	95	91	88	85	82	80	76	63	53	40
with noise after opening	99	95	88	82	81	75	71	69	60	42	39
with noise without opening	99	91	87	78	76	71	69	62	57	45	38

and the remaining noise was taken care of using the function d_K in Step 1. As for the choice of the half-planes in Step 2, we took the half-planes corresponding to the parameters $\vec{b} = (b, -b)$ with $b = 10, 13, 16$ and $\vec{l} = (\cos\theta, \sin\theta)$ with $\theta = 10°, 20°, \ldots, 80°$, with reference to Appendix A. Indeed, for these choices of b and θ the function d_K really interacts with the functions φ_0, φ_1, and φ_2.

In the third experiment, the query shapes were corrupted by noise and no preprocessing was performed. All the noise is smoothed out using in Step 1 the function $-\lambda^{\varepsilon}_K$, with $\epsilon = 10$, instead of d_K. As for the choice of the half-planes in Step 2, in this case we took those corresponding to the parameters $\vec{b} = (b, -b)$ with $b = 3, 5, 7, 9$ and $\vec{l} = (\cos\theta, \sin\theta)$ with $\theta = 10°, 20°, \ldots, 80°$. The motivation for these choices for b and θ is the same as before.

In all three experiments, the obtained one-dimensional PBNs were compared using the Hausdorff distance, as a lower bound of the matching distance to speed up computations. Next, these distances were normalized with mean equal to 0 and standard deviation equal to 1 so to obtain comparable values for different functions. Finally, as a dissimilarity measure between two shapes, we took the sum of the normalized Hausdorff distances.

The results proposed in Table 5 describe, for the octopus shape Figure 1(c) (269x256 pixels), the average time taken to extract the 1-dimensional PBNs for 0th homology on a half-plane of the foliation out of 40 half-planes, the total time required to compute the size function on the 40 half-planes considered, and the average and the total number of points of the persistence diagrams on the 40 half-planes.

[1] H. Edelsbrunner, D. Letscher, A. Zomorodian, Topological persistence and simplification, Discrete & Computational Geometry 28 (4) (2002) 511–533.

[2] G. Carlsson, A. Zomorodian, The theory

Table 5: Time requirements for the computation of the PBNs of the octopus image of Figure 1(c). Avg. time is the average time required to compute the PBNs on a single half-plane of the foliation, while Total time refers to the computation of the PBNs on 40 half-planes. Analogously, Avg. $|C|$ is the average number of points of a persistence diagram on a single half-plane of the foliation, and Total $|C|$ is the sum of the number of points of the persistence diagrams on 40 half-planes. These results are obtained using a processor T2400 at 1.83 GHz with 1 GB RAM.

Avg. time	Total time	Avg. $\lvert C\rvert$	Total $\lvert C\rvert$
0.421 sec	16.86 sec	10.325	413

of multidimensional persistence, Discrete & Computational Geometry 42 (1) (2009) 71–93.

[3] A. Verri, C. Uras, P. Frosini, M. Ferri, On the use of size functions for shape analysis, Biol. Cybern. 70 (1993) 99–107.

[4] S. Belongie, J. Malik, J. Puzicha, Shape matching and object recognition using shape contexts, IEEE Trans. Pattern Anal. Mach. Intell. 24(4) (2002) 509–522.

[5] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, S. W. Zucker, Shock graphs and shape matching, Int. J. Comput. Vis. 35(1) (1999) 13–32.

[6] S. Biswas, G. Aggarwal, R. Chellappa, An efficient and robust algorithm for shape indexing and retrieval, IEEE Transactions on Multimedia 12(5) (2010) 372–385.

[7] S. Biasotti, L. De Floriani, B. Falcidieno, P. Frosini, D. Giorgi, C. Landi, L. Papaleo, M. Spagnuolo, Describing shapes by geometrical-topological properties of real functions, ACM Comput. Surv. 40 (4) (2008) 1–87.

[8] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, Stability of persistence diagrams, Discrete Comput. Geom. 37 (1) (2007) 103–120.

[9] F. Cagliari, B. Di Fabio, M. Ferri, One-dimensional reduction of multidimensional persistent homology, Proc. Amer. Math. Soc. 138(8) (2010) 3003–3017.

[10] A. Cerri, B. Di Fabio, M. Ferri, P. Frosini, C. Landi, Betti numbers in multidimensional persistent homology are stable functions, Technical Report, Università di Bologna, http://amsacta.cib.unibo.it/2923/ (2010).

[11] M. d'Amico, P. Frosini, C. Landi, Natural pseudo-distance and optimal matching between reduced size functions, Acta Applicandae Mathematicae 109 (2010) 527–554.

[12] S. Biasotti, A. Cerri, P. Frosini, D. Giorgi, C. Landi, Multidimensional size functions for shape comparison, J. Math. Imaging Vision 32 (2) (2008) 161–179.

[13] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, D. Morozov, Persistent homology for kernels, images, and cokernels, in: Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '09, SIAM, 2009, pp. 1011–1020.

[14] T. Sebastian, P. Klein, B. Kimia, Recognition of shapes by editing shock graphs, in: ICCV'01, Vol. 1, 2001, pp. 755–762.

[15] S. Eilenberg, N. Steenrod, Foundations of algebraic topology, Princeton University Press, 1952.

[16] F. Cagliari, C. Landi, Finiteness of rank invariants of multidimensional persistent homology groups, Applied Mathematics Letters 24 (2011) 516–518.

[17] L. Zadeh, Fuzzy sets, Information Control 8 (1965) 338–353.

[18] D. Dubois, H. Prade, Fuzzy sets and systems - Theory and applications, Academic Press, New York, 1980.

[19] R. C. Veltkamp, M. Hagedoorn, State of the art in shape matching, Principles of visual information retrieval (2001) 87–119.

[20] M.-M. Deza, E. Deza, Dictionary of Distances, Elsevier, Burlington, 2006.

[21] M. C. Delfour, J.-P. Zolésio, Shapes and geometries: analysis, differential calculus, and optimization, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001.

[22] S. Willard, General Topology, Addison-Wesley Publishing Company, 1970.

[23] M. Lesnick, The optimality of the interleaving distance on multidimensional persistence modules, arXiv, arXiv:1106.5305v2 (2011).
[24] N. Cavazza, M. Ethier, P. Frosini, T. Kaczynski, C. Landi, Comparison of persistent homologies for vector functions: from continuous to discrete and back, Tech. Rep. 3205, Università di Bologna, http://amsacta.cib.unibo.it/3205/ (Gennaio 2012).
[25] S. Biasotti, A. Cerri, P. Frosini, D. Giorgi, A new algorithm for computing the 2-dimensional matching distance between size functions, Pattern Recognition Letters 32 (14) (2011) 1735 – 1746.
[26] T. Bernier, J. Landry, New method for representing and matching shapes of natural objects, Pattern Recognition 36(8) (2003) 1711–1723.
[27] Y. Ebrahim, M. Ahmed, S.-C. Chau, W. Abdelsalam, A template-based shape representation technique, in: Image Analysis and Recognition, Vol. 5112, LNCS, 2008, pp. 497–506.
[28] Z. Tu, A. Yuille, Shape matching and recognition using generative models and informative features, in: ECCV 2004, Vol. 3023, 2004, pp. 195–209.
[29] R. Zhou, L. Zhang, Shape retrieval using pyramid matching with orientation features, in: IEEE International Conference on Intelligent Computing and Intelligent Systems, Vol. 4, 2009, pp. 431–434.

Appendices

A. Foliation method

An effective way of studying multidimensional PBNs, whose domain Δ^+ is a subset of $\mathbb{R}^k \times \mathbb{R}^k$, is via a reduction to the one-dimensional case. This amounts to choose, for each $(\vec{u}, \vec{v}) \in \Delta^+$, a strictly increasing path through $\vec{u}$ and $\vec{v}$, and to consider the one-dimensional filtration defined by this path.

An appropriate choice of these paths allows us to obtain a foliation in half-planes of Δ^+ such that the restriction of the multidimensional PBNs to these half-planes turns out to give one-dimensional PBNs with respect to a filtration corresponding to the lower level sets of a certain (computable) scalar-valued function.

We start by recalling that the following parameterized family of half-planes in $\mathbb{R}^k \times \mathbb{R}^k$ is a foliation of Δ^+.

Definition A.1 ([12])**.** For every unit vector $\vec{l} = (l_1, \ldots, l_k)$ of $\mathbb{R}^k$ such that $l_i > 0$ for $i = 1, \ldots, k$, and for every vector $\vec{b} = (b_1, \ldots, b_k)$ of $\mathbb{R}^k$ such that $\sum_{i=1}^k b_i = 0$, we shall say that the pair $\left(\vec{l}, \vec{b}\right)$ is *admissible*. We shall denote the set of all admissible pairs in $\mathbb{R}^k \times \mathbb{R}^k$ by Adm_k. Given an admissible pair $\left(\vec{l}, \vec{b}\right)$, we define the half-plane $\pi_{\left(\vec{l},\vec{b}\right)}$ of $\mathbb{R}^k \times \mathbb{R}^k$ by the following parametric equations:

$$\begin{cases} \vec{u} = s\vec{l} + \vec{b} \\ \vec{v} = t\vec{l} + \vec{b} \end{cases}$$

for $s, t \in \mathbb{R}$, with $s < t$.

Since these half-planes $\pi_{\left(\vec{l},\vec{b}\right)}$ constitute a foliation of Δ^+, for each $(\vec{u}, \vec{v}) \in \Delta^+$ there exists one and only one $\left(\vec{l}, \vec{b}\right) \in Adm_n$ such that $(\vec{u}, \vec{v}) \in \pi_{\left(\vec{l},\vec{b}\right)}$. Observe that $\vec{l}$ and $\vec{b}$ only depend on $(\vec{u}, \vec{v})$. Intuitively, on each half plane $\pi_{\left(\vec{l},\vec{b}\right)}$ one can find the PBNs corresponding to the filtration obtained by sweeping the line through $\vec{u}$ and $\vec{v}$ parameterized by $\gamma_{\left(\vec{l},\vec{b}\right)} : \mathbb{R} \to \mathbb{R}^k$, with $\gamma_{\left(\vec{l},\vec{b}\right)}(\tau) = \tau\vec{l} + \vec{b}$.

We now recall that this filtration corresponds to the one given by the lower level sets of a certain scalar-valued continuous function.

Theorem A.1 ([9, 10]). *For every $(\vec{u},\vec{v}) \in \Delta^+$, let $\left(\vec{l},\vec{b}\right)$ be the only admissible pair such that $(\vec{u},\vec{v}) \in \pi_{\left(\vec{l},\vec{b}\right)}$. Let moreover $\varphi_{(\vec{u},\vec{v})} : X \to \mathbb{R}$ be the continuous filtering function defined by setting*

$$\varphi_{(\vec{u},\vec{v})}(x) = \min_i l_i \cdot \max_i \frac{\varphi_i(x) - b_i}{l_i}.$$

Then $X\langle \vec{\varphi} \preceq \vec{u}\rangle = X\langle (\min_i l_i)^{-1}\varphi_{(\vec{u},\vec{v})} \le s\rangle$. Therefore

$$\beta_{\vec{\varphi}}(\vec{u},\vec{v}) = \beta_{(\min_i l_i)^{-1}\varphi_{(\vec{u},\vec{v})}}(s,t)\,.$$

Finally, the most important property of this foliation method is that it allows us to obtain a distance for multidimensional PBNs, denoted by D_{match} and described in Appendix B, having a particularly simple form, yet yielding the Stability Property (S).

B. Multidimensional matching distance

We now recall the construction of the distance D_{match} to compare multidimensional PBNs. The key property of D_{match} is that it has the Stability Property (S). The construction is based on the foliation method described in Appendix A.

D_{match} was presented, and proved to yield stability of (multidimensional) PBNs, in [12] for 0th homology, and in [9] under a restrictive *max-tameness* assumption on the filtering functions. In [10], it was proved to yield stability of PBNs also in the wider setting of just continuous functions.

Definition B.1. Let X be a triangulable space endowed with continuous functions $\vec{\varphi} : X \to \mathbb{R}^k$, $\vec{\psi} : X \to \mathbb{R}^k$. The *multidimensional matching distance D_{match}* between $\beta_{\vec{\varphi}}$ and $\beta_{\vec{\psi}}$ is defined as

$$D_{match}\left(\beta_{\vec{\varphi}},\beta_{\vec{\psi}}\right) = \sup_{(\vec{u},\vec{v})\in\Delta^+} d_{match}\left(\beta_{\varphi_{(\vec{u},\vec{v})}},\beta_{\psi_{(\vec{u},\vec{v})}}\right). \quad \text{(B.1)}$$

We recall that d_{match} is a distance between one-dimensional PBNs that measures multi-bijections between persistence diagrams ([11, 8]). When $k = 1$, D_{match} coincides with the usual distance d_{match} between one-dimensional PBNs.

The following theorem states the Stability Property of multidimensional PBNs with respect to this distance: Small changes in a vector-valued filtering function induce small changes in the associated multidimensional PBNs, with respect to the distance D_{match}.

Theorem B.1 ([10]). *If X is a triangulable space, then D_{match} is a distance on the set $\{\beta_{\vec{\varphi}} \,|\, \vec{\varphi} : X \to \mathbb{R}^k \text{ continuous}\}$. Moreover,*

$$D_{match}\left(\beta_{\vec{\varphi}},\beta_{\vec{\psi}}\right) \le \left\|\vec{\varphi} - \vec{\psi}\right\|_\infty.$$