Eurographics Symposium on Geometry Processing 2016
Maks Ovsjanikov and Daniele Panozzo
(Guest Editors)

Volume 35 (2016), Number 5

Stable Region Correspondences Between Non-Isometric Shapes

V. Ganapathi-Subramanian[1], B. Thibert[1,2], M. Ovsjanikov[3] & L. Guibas[1]

[1]Stanford University
[2]Université Grenoble Alpes, LJK, France
[3]LIX, Ecole Polytechnique, France

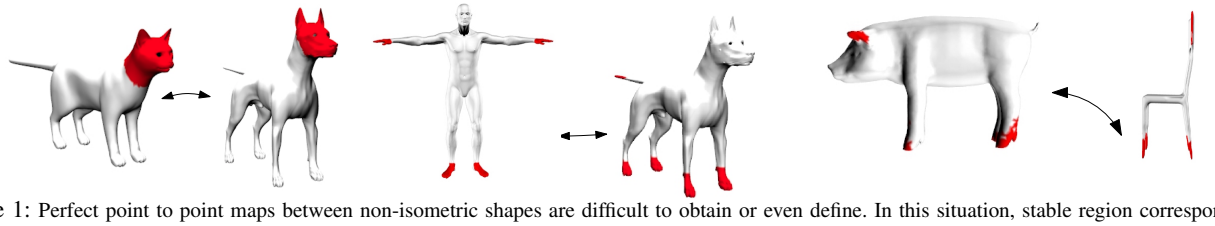

Figure 1: Perfect point to point maps between non-isometric shapes are difficult to obtain or even define. In this situation, stable region correspondences are a meaningful alternative for a set of matches. The three shape pairs above each illustrate a pair of stable matching regions obtained when we compute a map between cat and dog, human and dog, and pig and chair, respectively. Note that in all cases the obtained correspondences capture regions with similar semantics.

Abstract

We consider the problem of finding meaningful correspondences between 3D models that are related but not necessarily very similar. When the shapes are quite different, a point-to-point map is not always appropriate, so our focus in this paper is a method to build a set of correspondences between shape regions or parts. The proposed approach exploits a variety of feature functions on the shapes and makes use of the key observation that points in matching parts have similar ranks in the sorting of the corresponding feature values. Our algorithm proceeds in two steps. We first build an affinity matrix between points on the two shapes, based on feature rank similarity over many feature functions. We then define a notion of stability of a pair of regions, with respect to this affinity matrix, obtained as a fixed point of a nonlinear operator. Our method yields a family of corresponding maximally stable regions between the two shapes that can be used to define shape parts. We observe that this is an instance of the biclustering problem and that it is related to solving a constrained maximal eigenvalue problem. We provide an algorithm to solve this problem that mimics the power method. We show the robustness of its output to noisy input features as well its convergence properties. The obtained part correspondences are shown to be almost perfect matches in the isometric case, and also semantically appropriate even in non-isometric cases. We provide numerous examples and applications of this technique, for example to sharpening correspondences in traditional shape matching algorithms.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—3D Shape Matching, Geometric Modeling, Shape Matching and Retrieval

1. Introduction

Understanding shapes and the connections/relations between them is a classical problem in geometry processing. When shapes are very similar, it is appropriate to match them using point-to-point correspondences. However, when they are related but are significantly different, a point-to-point map might not be meaningful, as illustrated in Figure 1. In this paper, we consider the problem of extracting part or region correspondences between related but possibly dissimilar shapes.

A common approach for establishing correspondences between points on geometric shapes consists of associating a set of values (called signatures or feature descriptors) with every point on the shape, values that characterize the geometric properties of the point or of its neighborhood, often in a multi-scale way. Common examples of such features include various notions of curvature (Gaussian, mean), diffusion-based descriptors, such as the Heat or Wave Kernel Signatures [SOG09, ASC11], or more classical descriptors such as spin images or shape contexts [JH99, BMP02]. In most scenarios, after computing the features, correspondences are obtained by considering pairs of points or shape parts that have similar fea-

ture values, for example by doing nearest-neighbor search in the descriptor space.

A crucial observation which drives our work is that for many geometric features or descriptors, such as for example curvature, geometrically similar parts between different shapes obey a similar ranking with respect to the sorting of their feature values, even if the exact values of the features might differ. In other words, if each feature is considered as *a function* on the shape, and its values are sorted, then the rank of the value on corresponding points on two different shapes will be close, even if the descriptor value itself differs significantly. This is illustrated in Figure 2, where one can observe that the highest mean curvature values correspond to the tail of both the cat and the dog and that the lowest values correspond to the face. We also observe on this figure that different feature functions still be consistent in spite of being ranked differently on the same shape. The idea of our approach is to aggragate information obtained from many such features.

In this paper, we provide a method that, given a set of feature functions computed on a pair shapes, outputs a set of corresponding regions that partition each of the two shapes. Our approach consists of two steps: we first create an affinity matrix that captures the similarity between pairs of points on the two shapes. The affinity is high for a pair of points whose ranks in the sorted feature functions is similar across many features. Then, in a second step, we extract a set of pairwise corresponding regions that maximize their mutual affinity. For this we develop a novel iterative algorithm that is able to efficiently extract pairs of corresponding regions from the affinity matrix, by solving an optimization problem, which is closely related to biclustering and constrained eigenvalue maximization.

To summarize, our contributions are the following:
- We define a novel affinity matrix between two shapes that encodes the "similarity" of every pair of points, using the ranks of corresponding feature functions.
- We define the notion of stable pair of regions with respect to an affinity matrix, as the solution of an optimization problem, which is related to constrained eigenvalue maximization problem.
- We provide a practical iterative algorithm that converges to a stable pair of matching regions.
- We provide theoretical stability guarantees for the construction of the affinity matrix, and of the output of our algorithm, in the presence of noisy feature functions, as well as a proof of convergence of our iterative algorithm for extracting stable parts.
- We show that these stable parts provide new features that are both semantically meaningful for dissimilar shapes and can aid existing correspondence methods for nearly isometric shapes.

For every step of our algorithm, we put special emphasis on *robustness* of the approach in the presence of possibly extreme geometric variability. This is in contrast to most existing shape matching methods which are geared towards finding accurate point or part correspondences, but can easily fail even under moderate geometric variability. As a result, our approach can return plausible part correspondences even for very different shape pairs. We validate the output of our algorithm via quantitative accuracy measures for sufficiently similar shape pairs with known correspondences. We also observe on various examples that stable regions often correspond to geometrically meaningful parts between non-isometric shapes.

2. Related Work

Shape matching or correspondence finding is one of the oldest and best-studied problems in digital geometry processing involving a diverse set of approaches. Below we only mention the methods directly related to ours, putting emphasis on the non-rigid and especially non-isometric shape correspondence algorithms. We also invite the interested reader to consider existing surveys, including [VKZHCO11, TCL*13] and a tutorial [CLM*11].

The problem of finding correspondences between a pair of shapes undergoing a non-rigid deformation has received a lot of attention in the past decade. The most common theoretical framework for studying this problem is that of intrinsic isometries, which assumes that the underlying transformation approximately preserves geodesic distances on shape surfaces. Introduced by Bronstein et al. [BBK06] and Mémoli [Mém07] , this model has inspired a large number of efficient techniques for finding correspondences, including [TBW*09, SY11, OMMG10] to name a few, which are especially useful in the case of articulated shapes. Note that in [RPSS10], the features are sorted to find a one-to-one map between sample points that sends the feature values on the first shape as close as possible to the feature values on the second shape, and not to find a confidence value between pairs of points as in our approach.

Although appealing from the theoretical and practical points of view, the isometric shape deformation model is rather restrictive. A more general model of conformal deformations (i.e. ones that only preserve angles) has also led to a number of efficient shape matching methods, including [LF09, KLF11, APL14, APL15, ZSCO*08]. These methods are able to handle larger shape variability, but also fail in the presence of extreme geometric deformations. Moreover, even for shapes with moderate deformations, in part due to the presence of multiple solutions these methods can also strongly benefit from fixed feature correspondences.

A number recent techniques have also considered a more general problem of finding soft (or functional) correspondences between shapes, rather than point-to-point ones [SNB*12, OBCS*12, SRGB14]. These methods are based on the idea that in many cases, it is both more efficient and less error-prone to establish correspondences between probability distributions, or more generally real-valued functions on the shapes, rather than between pairs of points. This is especially true when the shapes undergo large deformations, so that even defining a point-to-point map outside of a small set of feature points might not be meaningful. Nevertheless, most of the existing techniques are still limited, and often incorporate the isometric deformation assumption as a regularizer in the optimization.

Non-rigid shape matching has also been tackled with machine learning-techniques, where one typically assumes some training data and learns a deformation model that can be applied on new shape pairs, which is possible both in the context of pointwise [ZB15, WHC*15] and functional [COC14] correspondences.

Finally, other methods have formulated the matching problem via classical quadratic assignment matching, for which several efficient relaxations recently been devised, e.g. [KKBL15, CK15]. These methods typically still rely on variants of isometric shape deformation model and aim to obtain pointwise correspondences.

Shape region decomposition has also been handled using topological tools such as topological persistence to provide a multiscale isometry invariant mesh segmentation [SOCG10], or to complete a shape or find correspondences between isometric shapes [DLL*10].

In this context, our main goal is to lift the restrictive assumptions on shape deformations and to provide a robust and efficient method for establishing part-based correspondences between a given pair of shapes, without any prior. As mentioned above, we target significant shape variability and emphasize robustness at every stage of the pipeline. However, we also show that our part-based correspondence algorithm can be used to improve point-to-point maps when these are required and appropriate.

Our method is also somewhat related to the 3D shape co-segmentation techniques, such as [SvKK*11, HKG11, HFL12, AXZ*15]. However, unlike these approaches, we specifically target the part-level correspondence of a *pair* shapes that undergo a potentially significant deformation. This way, our method for extracting stable parts from an affinity between pairs of points can potentially be used *to enhance* co-segmentation of a shape collection, as a future application.

Our approach bears similarity with spectral clustering, where shapes or points are embedded in a domain associated to the eigenvectors of some Laplacian matrix. Then a clustering or a map is calculated in the spectral domain and pushed back onto the shapes, as in [JZvK07]. The normalized graph cut problem is in particular close to our problem: it is formulated as a discrete minimization problem, which is in practice relaxed into a continuous problem that amounts to computing the Fiedler vector. The continuous solution is then discretized by sorting and truncating. In our method, we do not relax the problem and directly solve a discrete optimization problem by computing two binary vectors.

Our problem can also be considered as an instance of the biclustering problem similar to the minimization of normalized cut in a bipartite graph [ZHD*01]. There, the goal is also to find a simultaneous partition on two sets and the relaxed continuous problem is solved by calculating an eigenvector associated to the second largest left and right singular vectors of the affinity matrix, and a truncation is applied to get a partition. In our setting, the objective function is different, we directly solve the problem without relaxing the non-convex constraint and get a local maximum. Since we directly solve the discrete problem, we apply truncations at each iteration of a power method, while in the normalized cut problem, the eigenvector is calculated with the power method and the truncation is only done at the end.

3. Method Overview

In this paper we assume that we are given a pair of shapes and a set of corresponding feature functions computed on each shape. Our underlying assumption is that even if the values of the feature functions might be different at points we want to correspond, the ranks of the values of the sorted feature functions might still agree. For example, the regions corresponding to the maxima or minima of the feature functions are likely to correspond, even if the exact function values might be different. We also assume that

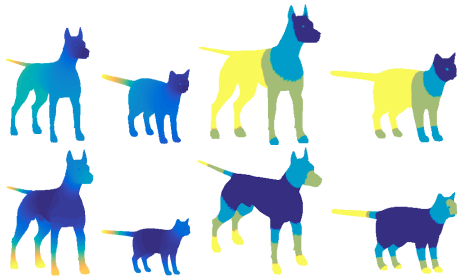

Figure 2: Feature functions and characteristic functions on two different shapes. The first row shows a multiscale mean curvature [MDSB03] and the second row shows a WKS signature [ASC11]. The two shapes are the TOSCA *cat0* and *dog0*. We visualize the feature values (Left two columns) as well as the partitions of the shapes into 4 characteristic feature value bands $\mathbb{1}_{C_j}$ and $\mathbb{1}_{C'_j}$ for $j = 1, \cdots, 4$ (Right two columns). The characteristic functions are consistent and provide independent information. In this paper we consistently combine information obtained from many such features.

the feature functions are sufficiently smooth and discriminative, so that different characteristic regions (pre-images of a fixed feature value range) correspond to contiguous regions on the shapes. One can consider diffusion-based descriptors such as the HKS or WKS that fulfill both of these assumptions [ASC11, SOG09]. This is illustrated in Figure 2, where one visualizes the sorting values of two corresponding features in two different shapes. Our algorithm proceeds in two stages: first we construct an *affinity matrix* that captures the correspondence likelihood of every pair of points on the shapes. This affinity matrix is constructed by using the correlations between the ranks of the pairs of points across different feature functions. We stress that the affinity matrix can be interpreted as a discrete rank transport plan and is stable in the presence of noisy features.

Given the affinity matrix, we then develop an algorithm that extracts corresponding parts of the two shapes in a robust, efficient and stable way. We provide a theoretical analysis of convergence and demonstrate its efficiency and accuracy in practice.

The final output of our method is a set of part/region correspondences on the pair of shapes. As our emphasis is on robustness in the presence of large geometric variability, our approach does not provide point-to-point correspondences, and in some cases, the corresponding regions might be relatively large. However, as we show below, such correspondences can still be useful in a variety of tasks even when more fine matchings are desired.

Remark. In our approach we do not directly enforce topological or geometric constraints. However, since our stable correspondences depend on input feature functions, and since the features we use contain geometric information, we observe geometric similarities on our output correspondences, even on pairs of topologically different shapes. Thus, our output correspondences are as robust to various perturbations as our input features. However, our method can break down when we have inconsistent features, such as extrinsic features on non-aligned shapes. Finally, although our method is

robust to random noise, our affinity matrix is not robust to noisy features that would provide bad matches on the shapes consistently.

The rest of the paper is organized as follows. Section 4 defines the notion of affinity matrix and describes its construction from input shape feature correspondences. Section 5 defines the notion of stability for regions on the pairs of shapes related by the affinity matrix. We present two algorithms to compute these stable regions in Section 6. Section 7.1 discusses the evolution of the stable regions on a pair of shapes. Section 7.2 provides examples of matching of non-isometric shapes and evaluation techniques using ground truth correspondences, and Section 7.3 provides an evaluation of how stable region correspondences could be used as features on shape pairs to improve existing maps between shapes.

Notations. Throughout the paper we let $S_1 = \{p_1, \cdots, p_{d_1}\}$ and $S_2 = \{q_1, \cdots, q_{d_2}\}$ denote the set of vertices of two shapes represented by triangle meshes. We denote by $\mathbb{1}_C$ the indicator function of any subset $C \subset S_i$ or of any interval $C \subset \mathbb{R}$. Remark that each feature function on S_i can be seen as a vector of $\mathbb{R}^{d_i}$ and is represented by a column matrix of size $d_i \times 1$. We denote by $\mathbf{1}_{k,l}$ the matrix of dimension $k \times l$ whose all entries are 1. For any matrix A, we denote by $\|A\|_1$ and $\|A\|_2$ respectively the induced matrix 1-norm and 2-norm, by $\|A\|_{1,1}$ the sum of the absolute values of the entries of A, by $\|A\|_{\mathcal{F}}$ its Frobenius norm and by $\rho(A)$ the highest eigenvalue in absolute value. For any vector x, we denote by $\|x\|_2$ its Euclidean norm, by $\|x\|_1$ its 1-norm and by x_i or $(x)_i$ its i^{th}-coordinate.

4. Affinity Matrix as Transport Plan

In this section we an *affinity matrix* between two shapes that measures the likelihood that two points on the shapes in question correspond. We define this affinity matrix in Section 4.1 and note that it can be seen as a discrete transport plan between two uniform probability measures. We show in Section 4.2 that, up to a scale factor and a translation factor, this affinity matrix is robust to random, independent and numerous noisy features.

4.1. Definition

From here on, we assume that we are given a pair of shapes and a set of N corresponding feature functions on each shape. The exact choice of the feature functions that we use in practice is detailed in Section 7.1. For simplicity we assume that vertex set cardinalities d_1 and d_2 are multiples of some integer $K > 0$.

Given a feature function f_k on S_1 and the corresponding feature function d_k on S_2 we partition the sets S_1 and S_2 into K subsets each, which we denote by $C_{k,j}$ and $C'_{k,j}$ (where $j = 1, \cdots, K$) respectively, by sorting the values f_k and d_k taking consecutive equally-sized blocks. This means that if $j \leq j'$ and $p \in C_{k,j}$ then for every $p' \in C_{k,j'}$, one has $f(p) \leq f(p')$. Given the construction above, we define the *affinity matrix* of size $d_2 \times d_1$ as:

$$W^{(N)} = \sum_{k=1}^{N} \sum_{j=1}^{K} \mathbb{1}_{C'_{k,j}} \cdot \mathbb{1}^T_{C_{k,j}},$$

where $\cdot$ denotes the matrix product. We remark that the value $W^{(N)}_{i,j}$ represents the number of features that classify vertex p_j of S_1 and vertex q_i of S_2 in the same characteristic class. Note also that the sum of every row of $W^{(N)}$ equals to Nd_1/K, and the sum of every column equals to Nd_2/K. The *normalized affinity matrix* is

$$W_N = K/(Nd_1d_2)W^{(N)}. \tag{1}$$

Again, we observe that W_N can be interpreted as a discrete transport plan between the uniform probability measures $\mu = (1/d_1)\mathbf{1}_{d_1,1}$ and $\nu = (1/d_2)\mathbf{1}_{d_2,1}$ [Vil03], *i.e.* the matrix W_N is positive and satisfies

$$W_N \cdot \mathbf{1}_{d_1,1} = \nu \quad \text{and} \quad W_N^T \cdot \mathbf{1}_{d_2,1} = \mu. \qquad \text{(Transport Plan)}$$

Each entry $(W_N)_{i,j}$ can be seen as a probability for vertices p_j and q_i to be matched. This implies in particular that $\mathbf{1}_{d_1,1}$ is an eigenvector of $W_N^T W_N$ associated to the eigenvalue $1/(d_1d_2)$. Furthermore, the induced 1-norm $\|W_N^T W_N\|_1$ is equal to $1/(d_1d_2)$ implying, using $\rho(A) \leq \|A\|_1$ that $1/(d_1d_2)$ is the largest eigenvalue of $W_N^T W_N$.

4.2. Robustness to noisy features

We show in this subsection that the affinity matrix is robust to the addition of noisy feature functions, as long as these feature functions are numerous and are the realization of independent, identically distributed (i.i.d) random variables. More precisely, let W_n^{init} be an affinity matrix of size $d_2 \times d_1$ built by a set of n "reliable" pair of feature functions (f_k, d_k), W_{N-n}^{noise} built using the $N-n$ pairs of real noisy feature functions (ϕ_k, ψ_k) and W_N^{all} built over all the N pairs. Then, one has

$$W_N^{all} = \kappa\, W_n^{init} + (1-\kappa) W_{N-n}^{noise},$$

where $\kappa = n/N$ is the fraction of reliable features.

The following proposition states that one can recover with high probability (up to a scale factor and a translation factor) an approximation of the non-noisy affinity matrix. Roughly speaking, the intuition is that a large number of noisy features will compensate for each other and simply modify all affinities in a uniform way.

Proposition 1 Suppose that $\phi_k : S_1 \to \mathbb{R}^{d_1}$ are $(N-n)$ i.i.d multivariate random variables and $\psi_k : S_2 \to \mathbb{R}^{d_2}$ are $(N-n)$ i.i.d multivariate random variables. Suppose in addition that for every k, l, ϕ_k and ψ_l are independent. Then for every $\delta > 0$

$$P(\|W_N^{all} - \kappa\, W_n^{init} - C\mathbf{1}_{d_2,d_1}\|_{\mathcal{F}} < \delta) \geq 1 - \frac{(1-\kappa)\sigma^2}{N\delta^2},$$

where $\|\cdot\|_{\mathcal{F}}$ is the Frobenius norm, $C = (1-\kappa)/(d_1d_2K)$ and σ is a constant independent on N, κ and δ.

Remark 1 We can define alternate affinity matrices. In our experiments we tried, for example, to weight the features using confidence values (by putting a high confidence value for features with a low entropy correspondence) but the results did not improve. Although this can, in part, be explained by the robustness of the affinity matrix, we leave the exploration of other possible constructions of the affinity matrix for future work.

5. Stable pairs of regions

Here we introduce the notion of stable pair of regions associated to an affinity matrix. As mentioned in the previous section, we interpret the normalized affinity matrix W_N as a discrete transport plan

between the two uniform probability measures supported on the two shapes S_1 and S_2. The matrix W_N sends the uniform measure on S_1 to the one on S_2 and the matrix W_N^T sends the uniform measure on S_2 back to the one on S_1. Intuitively, we would like to find a pair of subsets (Ω_1, Ω_2) such that W sends Ω_1 close to Ω_2 and W_N^T sends back Ω_2 close to Ω_1 — these will be our stable regions.

In Section 5.1 we express this problem as a constrained optimization problem and explain how it is related to eigenvalue maximization problems. We also show its equivalence to a maximization problem over a set of submatrices and define in Section 5.2 the stable regions as local maximizers, which can be expressed as fixed points of a nonlinear function (Section 5.3). We also show the stability of this notion with respect to the affinity matrix. Note that our approach works for any positive rectangular matrix.

5.1. Optimization problem

Let W be a positive matrix of size $d_2 \times d_1$, and let $q < d_1$ and $p < d_2$ be two integers. Our problem amounts to finding regions of S_1 and S_2 of respective sizes q and p that correspond to the highest values of the matrix W. This can be stated as follows:

$$\text{argmax}_{x,y} y^T W x \qquad \text{(Opt\_Pb)}$$

over the set of points $x \in \{0,1\}^{d_1}$ and $y \in \{0,1\}^{d_2}$ that satisfy $\|x\|_1 = q$ and $\|y\|_1 = p$.

We remark that the set of constraints is non-convex, finite and its size makes the optimization not tractable in practice. We show below that this problem can be reformulated as selecting a submatrix of W that maximizes the $\|\cdot\|_{1,1}$ norm. This characterization is a key point that will be used in Section 5.2 to introduce the notion of stable pairs of regions.

We say that a matrix A of size $q \times p$ is a *submatrix* of W if it is obtained by removing $(d_1 - q)$ columns and $(d_2 - p)$ rows from W. Given two regions $\Omega_1 \subset S_1$ and $\Omega_2 \subset S_2$, we denote by W_{Ω_1,Ω_2} the matrix of size $q \times p$ obtained by removing the columns of W whose index correspond to points in $S_1 \setminus \Omega_1$ and removing the rows of W whose index correspond to points in $S_2 \setminus \Omega_2$. Remark that W_{Ω_1,Ω_2} is a submatrix of size $q \times p$ of W.

Let x and y be two vectors that realize the maximum in Problem (Opt_Pb). Note that x is the indicator function of a set C_1 of S_1 and y is the indicator function of a set C_2 of S_1. Since the matrix W is positive, one has: $y^T W x = \mathbb{1}_{C_2}^T W \mathbb{1}_{C_1} = \mathbf{1}_{p,1}^T W_{C_1,C_2} \mathbf{1}_{q,1} = \|W_{C_1,C_2}\|_{1,1}$. Conversely, it is easy to see that for every submatrix of W of size $q \times p$ there correspond indicator functions x and y. Hence Problem (Opt_Pb) is equivalent to

$$\text{argmax}_A \|A\|_{1,1}, \qquad \text{(Opt\_Pb}')$$

where the maximum is over all the submatrices A of W of size $q \times p$.

Remark 2 If we do not enforce x and y to be indicator functions, and replace the assumption on their L_1 with a bound on the L_2 norm, then Problem (Opt_Pb) is equivalent to maximizing the highest eigenvalue of the positive semi-definite matrix $W^T W$. This is just a consequence of the classical equality

$$\max_{\|x\|_2 \le 1, \|y\|_2 \le 1} y^T W x = \max_{\|x\|_2 \le 1} \|Wx\|_2 = \sqrt{\rho(W^T W)}.$$

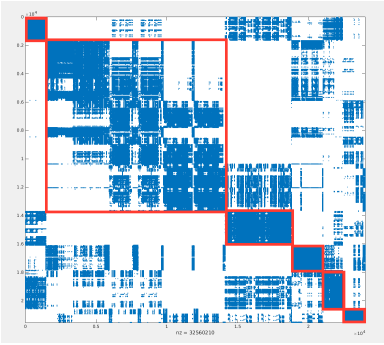

Figure 3: Structure of affinity matrix between a cat and dog. Only values $w_{i,j}$ greater than 30% of the maximum value are shown in blue. Algorithm 2 returns 6 stable pair of parts. The rows and the columns of W are reordered, so that each pair j correspond to one of the (diagonal) red rectangles. We observe in particular that the coefficients of W are higher close to the diagonal.

Note, however, that under generic conditions, the eigenvectors of matrix $W^T W$ do not provide, even approximately, regions of Shape S_1. This is a consequence of Perron-Frobenius Theorem. By generic condition, we mean the matrix $W^T W$ is irreducible (*e.g.* the associated graph is connected), which is the case in practice in our experiments. The Perron-Frobenius Theorem then implies that the eigenvector associated to the highest eigenvalue has strictly positive components (it is in fact $x = \mathbf{1}_{d_1,1}$ when W_N is the discrete transport plan between two uniform probability measures defined, as in Section 4), and that all other eigenvectors have at least one negative component.

Remark 3 Our optimization problem is related to sparse eigenvalues and eigenvectors problems, and in particular to sparse principal component analysis (see [JNRS10, YZ13] for example). Our approach here is in particular close to the one considered in [YZ13], where the following problem is considered

$$\max_{\|x\|_2 \le 1,\ \|x\|_0 \le q} x^T W x,$$

where $\|x\|_0$ is the number of non-zero elements. Note however that in their problem, the matrix W is supposed to be positive semi-definite, while it is only supposed to be positive in our framework. One may also remark that even though their approach could be applied to our matrix $W^T W$ (instead of W), the maximum is realized by a vector x which is not binary, and no vector y is provided.

5.2. Stable pairs as local maxima

The problem we want to solve amounts to finding global maximizers of the norm $\|\cdot\|_{1,1}$ over the set of submatrices of size $q \times p$ of W. The size of the search space makes the problem not tractable in practice. Thus, we define a stable pair of regions as being associated to a submatrix of W that locally maximizes the norm $\|\cdot\|_{1,1}$. To give a sense of locality, we first need to define a notion of neighborhood in the set of submatrices.

Up to a reordering of the rows and columns of W, we can write

$$W = \begin{pmatrix} W_{\Omega_1,\Omega_2} & W_{S_1\setminus\Omega_1,\Omega_2} \\ W_{\Omega_1,S_2\setminus\Omega_2} & W_{S_1\setminus\Omega_1,S_2\setminus\Omega_2} \end{pmatrix}, \tag{2}$$

Remark that to any submatrix A of W of size $q\times p$ correspond a pair of regions Ω_1 and Ω_2 such that $A = W_{\Omega_1,\Omega_2}$. We say that two submatrices $A = W_{\Omega_1,\Omega_2}$ and $A' = W_{\Omega_1',\Omega_2'}$ of size $q\times p$ are *neighbors* if there exists $i\in\{1,2\}$, a point $p\in S_i$ and a point $p'\in S_i'$ such that $\Omega_i\cup\{p'\}\setminus\{p\} = \Omega_i'$. Note that the set of all submatrices of W of size $q\times p$ is connected with this notion of neighborhood.

Definition. We say that the pair of regions $\Omega_1\subset S_1$ and $\Omega_2\subset S_2$ is *W-stable* if the submatrix W_{Ω_1,Ω_2} is a local maximizer of $\|\cdot\|_{1,1}$, namely if for every matrix neighbor $W_{\Omega_1',\Omega_2'}$, one has

$$\|W_{\Omega_1',\Omega_2'}\|_{1,1}\leq\|W_{\Omega_1,\Omega_2}\|_{1,1}.$$

We note that since the global maximum in Problem (Opt_Pb′) is always attained, there must always exist a stable pair of regions (Ω_1,Ω_2) of any size q and p.

5.3. Stable pairs as fixed points of a nonlinear function

In this section we show that stable pairs of regions can be expressed as fixed points of nonlinear functions. This observation is crucial and will be used in the algorithms presented in Section 6. We say that $\mathcal{P}:\mathbb{R}^d\to\{0,1\}^d$ is a *thresholding function* if for every $x\in\mathbb{R}^d$, if $x_i\leq x_j$, then $(\mathcal{P}(x))_i\leq(\mathcal{P}(x))_j$. Remark that for almost every $x\in\mathbb{R}^d$ there exists a cutoff $\tau(x)$, such that for every i, $(\mathcal{P}(x))_i = 0$ if $x_i<\tau(x)$ and $(\mathcal{P}(x))_i = 1$ otherwise.

Proposition 2 Let W be a discrete transport plan between the uniform probability measures on S_1 and S_2, and let (Ω_1,Ω_2) be a *W-stable* pair of regions. Then there exist two thresholding functions $\mathcal{P}:\mathbb{R}^{d_1}\to\{0,1\}^{d_1}$ and $\mathcal{Q}:\mathbb{R}^{d_2}\to\{0,1\}^{d_2}$ such that

$$\mathcal{P}(W\mathbb{1}_{\Omega_1}) = \mathbb{1}_{\Omega_2}\quad\text{and}\quad\mathcal{Q}(W^T\mathbb{1}_{\Omega_2}) = \mathbb{1}_{\Omega_1}$$

In particular, $\mathbb{1}_{\Omega_1}$ is a fixed point of $\mathcal{Q}\circ W^T\circ\mathcal{P}\circ W$ and $\mathbb{1}_{\Omega_2}$ is a fixed point of $\mathcal{P}\circ W\circ\mathcal{Q}\circ W^T$.

5.4. Robustness to noisy features

We show here the robustness of the stable pairs with respect to the affinity matrix. More precisely, if the noisy features used to build the affinity matrix are numerous and the realization of i.i.d random variables, then with a high probability, it does not change stable regions. Our result relies on the notion of δ-independency defined in the Appendix. Although this notion is quite restrictive, we can see from the proof that it is only used in rare cases.

Proposition 3 Suppose that W_N^{noise} is δ-independent and that conditions of Proposition 1 are satisfied. Then there is a probability at least $1-8(1-k)\sigma^2 d_\infty^2/(N\delta^2)$ that for every $\Omega_1\subset S_1$ and every $\Omega_2\subset S_2$

$$(\Omega_1,\Omega_2)\ W_N^{all}-\text{stable}\quad\Leftrightarrow\quad(\Omega_1,\Omega_2)\ W_N^{init}-\text{stable}.$$

where $d_\infty = \max(d_1,d_2)$.

Algorithm 1: stable_pair (given $\mathcal{P},\mathcal{Q}$)

input : W matrix of size $d_2\times d_1$
$f_0\in\{0,1\}^{d_1}$
output: Stable subspaces Ω_1,Ω_2
$f^{(0)} = f_0$
$f^{(1)} = \mathcal{Q}(W^T(\mathcal{P}(W(f^{(0)}))))$
$j = 1$
while $f^{(j)}\neq f^{(j-1)}$ **do**
  $g^{(j)} = \mathcal{P}(Wf^{(i)})$
  $f^{(i+1)} = \mathcal{Q}(W^Tg^{(j)})$
  $j = j+1$
end
Return $\Omega_1 := \{p_i|f_i = 1\}$ and $\Omega_2 = \{p_i|g_i = 1\}$

6. Algorithms to compute stable pair of regions

In this section, we provide algorithms to compute a set of stable pairs of regions between two shapes. As mentioned in the previous section, a stable pair is a nonlinear counterpart to an eigenvector associated to the highest eigenvalue in the linear setting. Our algorithms use the fixed point characterization mentioned before and are mimicking the power method in a nonlinear manner.

We propose in Section 6.1 a fixed point algorithm that outputs a pair of stable regions. We show in Proposition 4 that its output is robust to random noisy features, as long as these noisy features are independent and numerous. We also show in Proposition 5 that this algorithm terminates in a finite number of steps when the number of points in the stable regions remain fixed along the process. We propose in Section 6.2 Algorithm 2 that builds iteratively a set of stable pair of regions, given two thresholding functions. We also propose an alternate algorithm, Algorithm 3, that automatically chooses thresholding functions with a criterion that tends to create stable pairs of regions with similar and small areas.

6.1. A stable pair of regions

Algorithm 1 is a fixed point algorithm, that can be seen as a truncated power-method. It calculates a W-stable pair of regions, starting from an initial vector f_0 and a matrix W. Note here that the thresholding functions $\mathcal{P}$ and $\mathcal{Q}$ can be chosen arbitrarily.

In this algorithm, we transport a feature function from S_1 to S_2 and back to S_1, and iterate until convergence. This transport is done nonlinearly. Given a function f on S_1, we send it to S_2 by applying Wf. Now, we perform the non-linear step by applying the thresholding function. The intuition is that, since W is a confidence matrix, Wf is a confidence vector that gives us an idea of which vertices on Wf have been received well in the transportation using W. The thresholding retains the selected vertices which have been transported with high confidence by setting them to 1, and sets the remaining vertices to 0. This sort of thresholding on both shapes ensures that the vertices that stay as 1 have been transported with high confidence according to W. The algorithm converges to a fixed point, which is known to be W-stable, for $\mathcal{P}$ and $\mathcal{Q}$. Note that our algorithm bears some similarity with the truncating power-method developed for positive semi-definite matrices in [YZ13].

To avoid infinite loops, when several points have the same feature value which is exactly the cutoff value, we select appropriately

the points. In practice, we use the following thresholding functions (defined almost everywhere):

- The thresholding function $\mathcal{P}_p : \mathbb{R}^{d_2} \to \{0,1\}^{d_2}$ that associates to a vector x a vector $\mathcal{P}_p(x)$ with exactly p non-zero values. We define $\mathcal{Q}_q : \mathbb{R}^{d_1} \to \{0,1\}^{d_1}$ similarly.
- We can associate to each vertex p_i of S_1 a weight ω_i which is an area measure. Given a cutoff value $\alpha \in (0,1)$, we define the thresholding function $\mathcal{P}_\alpha$ by $\mathcal{P}_\alpha(x) := \mathcal{P}[\tau(x)](x)$, where the cutoff is given by $\tau(x) = \inf\{t, \sum_i \omega_i \mathbb{1}_{[t,+\infty)}(x_i) \geq \alpha \sum_i \omega_i\}$. We define $\mathcal{Q}_\alpha : \mathbb{R}^{d_1} \to \{0,1\}^{d_1}$ similarly.

The following proposition states that with a high probability the output of Algorithm 1 is robust to noisy features, as long as they are numerous and the realization of i.i.d random variables, under the assumption of δ-independency.

Proposition 4 (Stability of Algorithm 1) Suppose that W_N^{all} is δ-independent, for a given $\delta > 0$ and that assumptions of Proposition 1 are satisfied. Then there is a probability at least $1 - 8(1-k)\sigma^2 d_\infty^2/(N\delta^2)$ that for any $f_0 \in \mathbb{R}^{d_1} \setminus \{\mathbf{1}_{d_1,1}\}$, the output of Algorithm 1 with f_0 and W_N^{all} is the same than with f_0 and W_n^{init}.

We prove in the following proposition that our algorithm terminates in the case where the number of points in each shape remain fixed along the process. The key idea is that at each iteration, either we have a fixed point, or the norm $\|W_{\Omega_1,\Omega_2}\|_{1,1}$ is increasing by at least a strictly positive constant.

Proposition 5 (Termination of Algorithm 1) Let $q < d_1$ and $p < d_2$ be two positive integers. Algorithm 1 with ($\mathcal{P} = \mathcal{P}_p$ and $\mathcal{Q} = \mathcal{Q}_q$) converges in a finite number of steps to a stable pair (Ω_1, Ω_2) of sizes respectively q and p.

Remark 4 We prove the convergence when the sizes are fixed along the process. However if the number of rows or columns only increases, then the norm $\|W_{\Omega_1,\Omega_2}\|_{1,1}$ increases and the algorithm terminates in a finite number of steps. In practice, the algorithm converges in our experiments in less than 10 steps, see Fig. 5.

6.2. Iterative construction of several stable pairs of regions

We propose in this subsection two iterative algorithms to compute a set of stable pairs of regions. The idea is to first compute a stable pair of regions (Ω_1, Ω_2) of S_1 and S_2, and to apply again the algorithm on the remaining regions $S_1 \setminus \Omega_1$ and $\mathcal{S}_2 \setminus \Omega_2$. The process is then iterated. Alg. 2 correspond to the case where the thresholding functions are fixed. In Alg. 3, the thresholding functions are automatically calculated.

We first introduce notations. Let $s_i \subset S_i$ for $i = 1,2$. Given a matrix W, $W[s_2, s_1]$ denotes the matrix of the same size of W, where $W[s_2,s_1]_{ij} = w_{ij}$, if i is the index of a point of s_1 and if j is the index of a point of s_2, and $W[s_2,s_1]_{ij} = 0$ otherwise. Similarly, for a vector f of $\mathbb{R}^{d_k}$, $f[s_k]$ denotes the vector of $\mathbb{R}^{d_k}$ that satisfies $f[s_k]_i = f_i$ if i is the index of a point of s_k and $f(s_k)_i = 0$ otherwise.

In this algorithm, at every step of the iteration, we discard the previously obtained stable regions from the affinity matrix, and proceed to find stable parts on the remaining shape as in Algorithm 1. This way, we enforce that obtained stable regions are disjoint. In

Algorithm 2: Several stable pairs of regions

input : S_1 and S_2: sets of (weighted) points of size d_1 and d_2
W: matrix of size $d_2 \times d_1$
$\mathcal{P}$ and $\mathcal{Q}$: thresholding functions
K: integer
output: K stable pairs of parts $(\Omega_{1,j}, \Omega_{2,j})$ for $j = 1, \cdots, K$
$W^{(0)} = W$,
$s_1 = S_1, s_2 = S_2$.
for i = 1:K-1 **do**
  $f_{in}[s_1] =$ **random**$(\{0,1\}^{|s_1|})$
  $(\Omega_{1,i}, \Omega_{2,i}) =$ **stable_pair**$(W^{(i-1)}, f_{in})$ (Algorithm 1)
  $s_1 = s_1 \setminus \Omega_{1,i}, s_2 = s_2 \setminus \Omega_{2,i}$
  $W^{(i)}[s_2, s_1] = W[s_2, s_1]$
end
$\Omega_{1,K} = S_1 \setminus s_1$ and $\Omega_{2,i} = S_2 \setminus s_2$

Figure 3, we observe a sparsified version of the output of Algorithm 2. Here we observe that the obtained (rearranged) affinity matrix is not symmetric, and also the fraction of vertices of each shape that contribute to the stable correspondence are also very different.

In practice, one wants to compute stable pair of regions of small sizes as these are more informative. Furthermore, one may want to enforce for every i the sizes of $\Omega_{1,i}$ and $\Omega_{2,i}$ to be comparable. In order to get a notion of "size", we take an area measure $\omega_{1,j}$ associated to each vertex p_j of S_1 and an area measure $\omega_{2,j}$ associated to each vertex q_j of S_2. We also introduce the cutoff function $\tau_t(x) := \mu_x + t\sigma_x$, where $t \in \mathbb{R}$, μ_x and σ_x^2 are respectively the mean value and the variance of the vector x (x is a vector of $\mathbb{R}^{d_1}$ or $\mathbb{R}^{d_2}$).

At each iteration, Alg. 3 first computes the highest cutoff values τ_t (when t is only allowed to be a multiple of a given $\delta_t > 0$), such that the output of Alg. 1 (with the thresholding functions $\mathcal{P}(x) := \mathcal{P}[\tau_t(x)](x)$ and $\mathcal{Q}(x) := \mathcal{Q}[\tau_t(x)](x)$) is a pair of nonempty sets, denoted by (Ω_1^t, Ω_2^t). Then, Algorithm 1 is applied using this cutoff and while enforcing the same measure. More precisely, we apply Alg. 1 using the thresholding functions $\mathcal{P}_\alpha$ and $\mathcal{Q}_\alpha$, where $\alpha = \max(\alpha_1, \alpha_2)$ and $\alpha_i = \sum_{p_j \in \Omega_i^t} \omega_{i,j} / \sum_{p_j \in \Omega_i} \omega_{i,j}$ (for $i = 1,2$). In practice, if S_i is the set of vertices of a triangulation, $\omega_{i,j}$ is the sum of the area of all the triangles adjacent to the vertex of index j, divided by 3.

7. Experiments and Results

7.1. Algorithm Evolution

In this section, we experimentally investigate the stable regions obtained by the algorithms specified in Section 6. For all the experiments in this section, we use shapes from the TOSCA [BBK08] and SHREC watertight [GBP07] datasets, apart from certain handcrafted shapes from the Non-Rigid shape datasets. In all experiments we used a fixed set of 100 WKS, 100 HKS and 100 multiscale mean curvature features [ASC11, SOG09, MDSB03] unless specified otherwise. For the SHREC watertight dataset alone, we used shape features, as described in the Appendix A of [KHS10]. We used $K = 20$ characteristic bands on each feature function for all our experiments, unless specified otherwise.

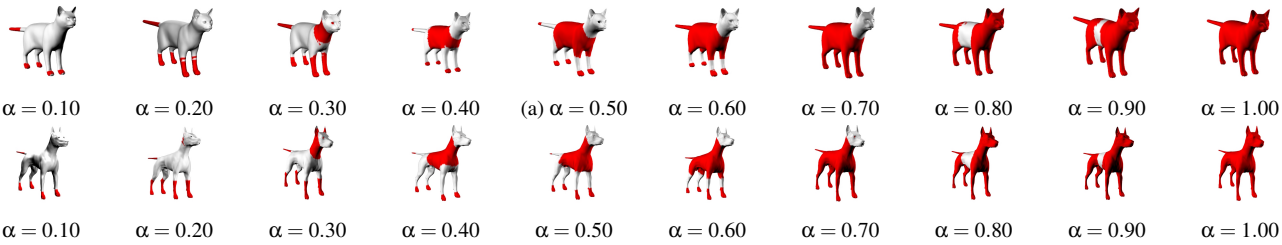

Figure 4: Evolution of stable pair of parts correspondences on TOSCA shapes *cat0* and *dog0*, with respect to the area proportion. Here α is the ratio of the surface area occupied by the stable parts divided by the surface area of the entire shape.

Algorithm 3: Stable pairs of regions with automatic threshold

input : S_1 and S_2: sets of (weighted) points of size d_1 and d_2
W: matrix of size $d_2 \times d_1$
δ_t: threshold step-size
K: integer
output: K stable pairs of regions $(\Omega_{1,j}, \Omega_{2,j})$ for $j = 1, \cdots, K$

$W^{(0)} = W$,
$s_1 = S_1, s_2 = S_2$.
for $i = 1 : K-1$ **do**
  $f_{in}[s_1] =$ **random**$(\{0,1\}^{|s_1|})$
  // Computation of the largest parameter t
  $t = 0$; $\Omega_1^t = s_1$
  while $\Omega_1^t \neq \mathbf{0}$ **do**
    $(\Omega_1^t, \Omega_2^t) =$ **stable_pair**$(W^{(i-1)}, f_{in})$
    (with $\mathcal{P}(x) := \mathcal{P}[\tau_t(x)](x)$ *and* $\mathcal{Q}(x) := \mathcal{Q}[\tau_t(x)](x)$*)*
    $t_{old} = t$
    $t = t + \delta_t$
  end
  $t = t_{old}$
  // Computation of a stable pair of regions
  $\alpha = \max(\alpha_1, \alpha_2)$
  $(\Omega_{1,i}, \Omega_{2,i}) =$ **stable_pair**$(W^{(i-1)}, \mathbb{1}_{\Omega_{1,i}})$(with $\mathcal{P}\alpha, \mathcal{Q}\alpha$)
  $s_1 = s_1 \backslash \Omega_{1,i}$ and $s_2 = s_2 \backslash \Omega_{2,i}$
  $W^{(i)}[s_2, s_1] = W[s_2, s_1]$
end
$\Omega_{1,K} = S_1 \setminus s_1$ and $\Omega_{2,i} = S_2 \setminus s_2$

Evolution of stable pair of regions We demonstrate the performance of Alg.1, with change in the threshold area value, for a fixed pair of cat and dog shapes. The results can be seen in Fig.4 where the two rows depict the evolution of the stable parts obtained on the two shapes for various choices of the area threshold. Observe that whenever a new part appears on one shape, it simultaneously appears on the other shape as well. To validate this, in Fig.(5-left) we plot relative change of the part area, $\text{Area}(\Omega_{\alpha+\delta_\alpha} \cap \Omega_\alpha)/\text{Area}(\Omega_\alpha)$, with respect to the evolution of the threshold α. If the new stable part is just an extension of the old stable part, then the value is 1, while if the new stable part has a new area, with a reduction in the old stable part, then the value drops below 1.

Algorithm termination Next, we illustrate the termination of the algorithm, by plotting the change in the error at each iteration, given by $||f^{(i)} - f^{(i-1)}||_2$ for a choice of thresholds (see Fig. (5-right)).

Algorithm Timing		
Shape dimensions	Algorithm 2(s)	Algorithm 3(s)
5400, 5619	6.12	54.13
6448, 8679	10.99	95.69
25290, 27894	116.96	1128. 74
25290, 45659	191.50	2001.29
45659, 52565	645.10	3720.43

Table 1: Algorithm Timing: The time taken to compute stable regions on both algorithms, for various dimensions of affinity matrix is provided. The second column provides time in seconds to run Algorithm 2, with a 10% cutoff for all stable regions, and the third column provides time in seconds to run Algorithm 3.

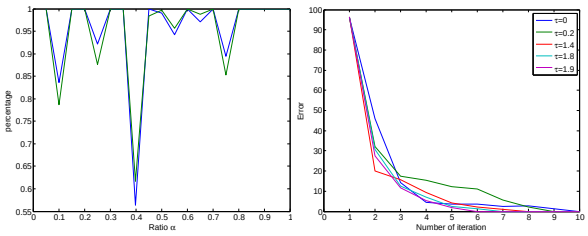

Figure 5: ***Left***: Evolution of the stable parts. We plot $\text{Area}(\Omega_{\alpha+\delta_\alpha} \cap \Omega_\alpha)/\text{Area}(\Omega_\alpha)$ for each ratio area parameter α. The green curve is the for the TOSCA *dog0* shape and the blue curve is for the TOSCA *cat0* shape. ***Right***: Fixed point error. We plot the error $\|f^{(i+1)} - f^{(i)}\|_2$ at each iteration step i of Algorithm 1. Each curve corresponds to a different threshold value: $\tau_0(x), \tau_{0.2}(x), \tau_{1.4}(x), \tau_{1.8}(x)$ and $\tau_{1.9}(x)$. The shapes used here are the TOSCA *cat*0 and *dog*0.

In particular, we observe that for almost all threshold values, the algorithm terminates very quickly, almost always within 10 iterations. With increasing threshold, the number of iterations often is reduced since the degrees of freedom of the Wf vector decreases.

7.2. Non-isometric shape correspondences

Next, we provide examples of stable region correspondences between non-rigid shapes, which is the main focus of the approach presented in this paper. As mentioned earlier, our method does not provide a point-to-point correspondences between two shapes, but rather correspondence between automatically-extracted shape regions. Since most of the existing shape matching techniques produce a point-to-point map, we use a heuristic to evaluate our ap-

Correspondence ratios on Rigid Shapes Dataset		
Shape pairs	Algorithm 2	Algorithm 3
	Ours/BIM$_{max}$/BIM$_{\mu}$	Ours/BIM$_{max}$/BIM$_{\mu}$
Elephant, Horse	62.37 / 26.46 / 22.16	73.88 / 43.82 / 36.44
Homer, Alien	48.09 / 30.86 / 23.59	59.56 / 36.96 / 26.23
Alien, Robot	40.76 / 50 / 42.35	54.32 / 54.81 / 50.03
Baby, Boy	24.26 / 58.10 / 42.50	47.46 / 66.37 / 47.52
Elephants	96.54 / 93.82 / 92.3	95.04 / 95.19 / 94.15
Average of correspondence ratios on Watertight Dataset		
Airplane	41.93 / 41.36 / 35.59	63.64 / 60 / 48.77
Ant	40.29 / 50.9 / 42.06	63.75 / 69.96 / 56.69
Armadillo	27.32 / 19.90 /17.27	40.85 / 41.98 / 34.30
Bearing	21.53 / 20.83 / 10.42	42.26 / 35.94 / 25.75
Bird	28.30 / 40.93 / 29.42	44.17 / 54.20 / 39.33
Bust	20.57 / 22.77 / 18.80	29.48 / 36.73 / 28.13
Chair	51.10 / 38.96 / 31.79	61.53 / 47.99 / 24.69
Cup	48.30 / 43.64 / 32.91	46.52 / 36.82 / 25.93
Fish	34.95 / 21.56 / 21.42	59.99 / 43.31/ 38.72
Fourleg	26.80 / 25.82 / 17.80	36.19 / 36.52 / 31.51
Glasses	39.35 / 34.88 / 25.96	60.09 / 55.71 / 35.26
Hand	16.76 / 33.12 / 26.56	24.26 / 32.77 / 31.745
Human	33.75 / 39.39 / 35.38	60.04 / 57.59 / 47.25
Mech	12.63 / 17.21 / 15.87	14.23 / 18.26 / 17.46
Octopus	39.28 / 38.84 / 30.04	70.37 / 59.26 / 51.20
Plier	79.36 / 90.47 / 71.67	94.44 / 98.45 / 82.14
Table	50 / 25.4 / 20.66	85.18 / 46.30 / 42.84
Teddy	51.64 / 59.88 / 48.21	61.77 / 62.4 / 53.64

Table 2: Comparison with BIM on the Rigid Shapes Dataset and the Watertight Dataset. The correspondence ratio is given in percentage. On the left column, the evaluation is done on a set of stable regions given by Algorithm 2. On the right, the evaluation is done on a set of stable regions given by Algorithm 3.

proach and compare it with existing methods against a ground truth correspondence. More precisely, given an existing ground-truth point-to-point map $L : C_1 \subset S_1 \to S_2$ between a a subset C_1 of a shape S_1 and shape S_2, and a set of stable pairs of parts $\mathcal{S} = \{(\Omega_{1,i}, \Omega_{2,i})\}_{i=1,\cdots,K}$, we define the *correspondence ratio* as the ratio of correspondences $(p, L(p))$ for which $L(p) \in \Omega_{2,i}$, if $p \in \Omega_{1,i}$, namely:

$$\mathrm{Cor}_L(\mathcal{S}) = \frac{|\{p \in C_1 |\ \exists i,\ \text{s.t}\ p \in \Omega_{1,i} \text{ and } L(p) \in \Omega_{2,i}\}|}{|C_1|}.$$

Comparison with BIM. We compare our method to the Blended Intrinsic Maps (BIM - [KLF11]) technique on both the SHREC watertight and on a non-rigid shape correspondences datasets.

We provide in Table 2 the results for the Rigid Shape dataset correspondences. Given two shapes shape S_1 and S_2, BIM outputs a map $T : S_1 \to S_2$. We also consider a set of stable parts $\{\Omega_{1,i}, \Omega_{2,i}\}_{i=1,\cdots,K}$ given by one of our algorithms. We calculate the correspondence ratio $\mathrm{Cor}_L(S_1, S_2)$ of the sets of stable parts defined by $\{\Omega_{1,i}, T(\Omega_{1,i})\}_{i=1,\cdots,K}$, where $T(\Omega_{1,i}) := \{T(x) |\ x \in \Omega_{1,i}\}$. We also calculate the correspondence ratio $\mathrm{Cor}_L(S_2, S_1)$ of the sets of stable parts defined by $\{T^{-1}(\Omega_{2,i}), \Omega_{2,i})\}_{i=1,\cdots,K}$, where $T^{-1}(\Omega_{2,i}) := \{x |\ T(x) \in \Omega_{2,i}\}$. The average of these two numbers is recorded in BIM$_{\mu}$, and the greatest value is recorded in BIM$_{max}$. The set $\{\Omega_{1,i}, \Omega_{2,i}\}_{i=1,\cdots,K}$ is either obtained by Algorithm 2 with an area-cutoff of 10 % or by Algorithm 3.

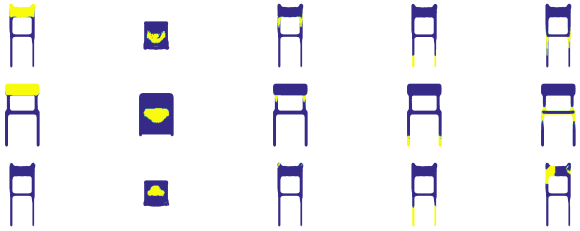

Figure 6: Comparison with BIM: The first and second rows show the corresponding stable regions on chair 1 and chair 2, while the third row shows the BIM pre-image of the second row. Though BIM works very well for keypoint correspondences, we can often show an improvement over BIM on region correspondences.

We use Algorithm 2 with an area thresholding function, so that the stable parts are all equal in size, and the accumulation of the number of correspondence vertices across stable parts is straightforward. In the case where Algorithm 3 is used, the results are better, since the stable parts are free from any area constraint. Thus, the advantage of using Algorithm 2 over Algorithm 3 is that the stable regions are more uniformly sized, while the negative is that the evolution of the final stable regions is not natural, as the area thresholding is forced.

We also compare our technique to the BIM on the Watertight benchmark. We consider 18 object categories of the Watertight dataset with 360 shapes overall. (20 shapes per category), and the Surface Correspondence Benchmark provides ground-truth correspondences within these object categories. We observe for how many of these well-defined correspondences do the corresponding points belong to the same stable region pair. Note that in spite of averaging over many pairs of shapes, the values in this table are erratic since the number of ground truth correspondences for each of the shape pairs is not very high/dense as in the Rigid Shapes Dataset. If $x \in S_1$ and $y \in S_2$ correspond to each other, it is expected that if $x \in \Omega_{i,1}$ and $y \in \Omega_{j,2}$, then $i = j$. The evaluation is identical to that in the previous table. Figure 6 shows an example of how the stable region correspondences compare with BIM correspondences.

Remark. The features we use are predominantly geometric features, and do not capture topological information. This is shown in the first two rows of Figure 9, where cups with and without a handle are used. Here, we observe that in spite of significantly different topological differences between the shapes, it is possible to draw meaningful correspondences between them.

7.3. Region correspondences as new feature functions

Our set of pairs of stable regions can be viewed as pair of feature functions that can be themselves used in different applications, extending the original feature set that was used to obtain them. In this section, we illustrate the use of these new features in sharpening existing point-to-point shape mapping algorithms. To construct these maps, feature correspondences are used in a least squares formulation. Since on isometric shapes our technique provides stable pair of parts that correspond almost perfectly (see Figure 7), we can

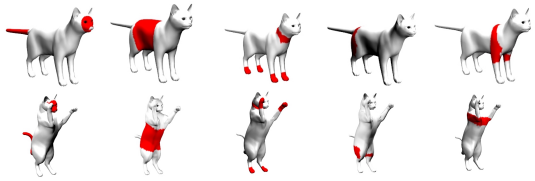

Figure 7: Stable pairs of parts between near-isometric shapes. Every column corresponds to a stable pair of parts, obtained using Algorithm 2 with an area-cutoff of 10%.

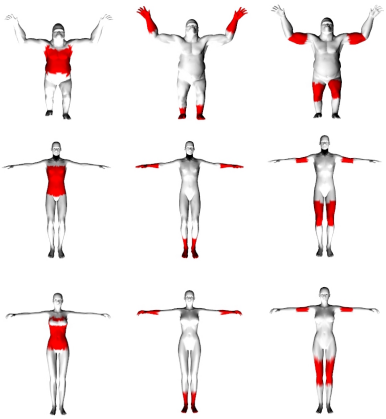

Figure 8: Stable regions generated by extending Algorithm 2 to three shapes using an area threshold.

use them to define new pairs of corresponding feature functions for functional maps. Given a stable pair of parts Ω_1 and Ω_2 on S_1 and S_2, we construct, as in [OBCS*12], a set of functional constraints using the Wave Kernel Map [OMMG10] based on segment correspondences. These new feature functions can be used alongside existing feature correspondences to further sharpen the map.

We sharpen the functional maps obtained in [OBCS*12] (100 base correspondences using WKS features). The sharpening is done using 4 stable parts, each of them contributing to 50 Wave Kernel Map based features. In these experiments, we observe the reduction in geodesic error when we use the stable correspondence based feature functions, as compared to just using the WKS feature functions. Here, we compute the functional maps (using our features obtained from the stable correspondences, alongside the existing WKS features), and compute point to point maps from them, using the technique employed in [OBCS*12]. We also use the same evaluation of average geodesic error with respect to a ground truth correspondence, as was used in the original functional maps paper. The results of these experiments on various near-isometric shapes in the TOSCA dataset is shown in Table 3. Clearly, on average, the stable parts extracted using our approach provide additional information compared to using the WKS feature functions directly. Thus, when we use the stable part correspondences together with WKS feature functions, the results improve significantly.

We can also extend Algorithms 2 and 3 to obtain stable regions on more than two shapes, so as to obtain a cleaner and more stable set of shape regions. As shown in Proposition 2, our algorithms compute a fixed point of $\mathcal{Q} \circ W^T \circ \mathcal{P} \circ W$. For the case of three

TOSCA Dataset			
Shapes	Error: WKS Mean/Median	Error: stable parts Mean / Median	Error: Both Mean / Median
cat	15.61 / 17.54	13.18 / 10.77	12.79 / 9.42
dog	6.50 / 6.89	6.68 / 6.75	6.16 / 6.43
david	12.13 / 13.21	8.14 / 7.28	6.44 / 5.48
victoria	24.44 / 21.31	19.56 / 15.63	14.70 / 17.36
horse	9.34 / 8.55	10.42 / 11.39	9.50 / 11.29
wolf	3.16 / 4.29	7.48 / 3.32	4.92 / 4.81
michael	24.11 / 17.78	17.74 / 14.71	12.21 / 11.76

Table 3: Comparing Functional Maps approach on just WKS feature function correspondences and the approach on Stable part correspondences, that have been converted to feature function correspondences by Wave Kernel Maps, for shapes TOSCA categories.

shapes, this is just done by finding the fixed point of a function $\mathcal{P}_1 \circ W_{13} \circ \mathcal{P}_3 \circ W_{32} \circ \mathcal{P}_2 \circ W_{21}$, where the functions $\mathcal{P}_i$ are thresholding functions on shape S_i and W_{ij} is the affinity matrix between shapes S_j and S_i. We have illustrated this in Figure 8. This idea can be extended to more than three shapes as well.

8. Conclusion and Future Work

We have presented in this paper a tool to extract a set of stable pairs of corresponding regions between two shapes, using as input a set of corresponding feature functions on each of the two shapes. We use the observation that points in matching parts often have similar ranks in the sorting of the feature functions values and observe that on many examples that stable pairs of regions correspond to geometrically meaningful parts.

Since all our knowledge of a shape comes from its feature functions, our method does not distinguish symmetric regions from each other. We plan to investigate local methods of our technique, where stable regions are required to contain a given seed point and be connected. We can then exploit various quadratic assignment assignment techniques, such as the one in [LH05] or flips in a coarse to fine approach such as in [SY13] to handle symmetry ambiguities.

One interesting future direction is to better relate these stable regions to semantic meaningful parts. In order to do that, one idea is to enforce that one of the two parts is a semantically meaningful segment. A potential application would be the fast annotation of shape parts in a network of shapes, if we have annotated information on a small number of shapes. Since our iterative algorithm converges to a fixed point quite quickly, we could propagate the annotation information to other shapes, in the same spirit as in Fig. 8.

The central algorithm relies on two thresholding parameters, whose selection is automated. However, determining the optimal threshold is still an open question and is application-dependent. In this article, we enforce the area ratios of the regions to be similar in the two shapes. It would be interesting to consider other constraints that handle correspondences between shapes under general non-linear deformations, such as partial scalings of shape parts.

Acknowledgements: This work was supported by NSF grants DMS-1228304, DMS-1521608, and DMS-1546206, French ANR grants TopData ANR-13-BS01-0008 and Comedic ANR-15-CE40-0006, Marie-Curie CIG-334283-HRGP, a CNRS chaire

d'excellence, chaire Jean Marjoulet from Ecole Polytechnique, FUI project "TANDEM 2" as well as a Google Focused Research Award, a Google Faculty Research Award, the Max Planck Center for Visual Computing and Communications and the National Center for Scientific Research (CNRS). The authors would also like to thank Panos Achlioptas for making the FMapLib framework available and Sören Pirk for helping render images.

References

[APL14] AIGERMAN N., PORANNE R., LIPMAN Y.: Lifted bijections for low distortion surface mappings. ACM Transactions on Graphics (TOG) 33, 4 (2014), 69. 2

[APL15] AIGERMAN N., PORANNE R., LIPMAN Y.: Seamless surface mappings. ACM Transactions on Graphics (TOG) 34, 4 (2015), 72. 2

[ASC11] AUBRY M., SCHLICKEWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In Computer Vision Workshops (ICCV Workshops) (2011), IEEE, pp. 1626–1633. 1, 3, 7

[AXZ*15] ALHASHIM I., XU K., ZHUANG Y., CAO J., SIMARI P., ZHANG H.: Deformation-driven topology-varying 3d shape correspondence. ACM Transactions on Graphics (TOG) 34, 6 (2015), 236. 3

[BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. PNAS 103, 5 (2006). 2

[BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Numerical geometry of non-rigid shapes. Springer, 2008. 7

[BMP02] BELONGIE S., MALIK J., PUZICHA J.: Shape matching and object recognition using shape contexts. PAMI 24, 4 (2002), 509–522. 1

[CK15] CHEN Q., KOLTUN V.: Robust nonrigid registration by convex optimization. In The IEEE International Conference on Computer Vision (ICCV) (December 2015). 2

[CLM*11] CHANG W., LI H., MITRA N., PAULY M., RUSINKIEWICZ S., WAND M.: Computing correspondences in geometric data sets. Eurographics Tutorials (2011). 2

[COC14] CORMAN É., OVSJANIKOV M., CHAMBOLLE A.: Supervised descriptor learning for non-rigid shape matching. In Computer Vision-ECCV 2014 Workshops (2014), Springer, pp. 283–298. 2

[DLL*10] DEY T. K., LI K., LUO C., RANJAN P., SAFA I., WANG Y.: Persistent heat signature for pose-oblivious matching of incomplete models. In Computer Graphics Forum (2010), vol. 29, Wiley Online Library, pp. 1545–1554. 3

[GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: Shape retrieval contest 2007: Watertight models track. SHREC competition 8 (2007). 7

[HFL12] HU R., FAN L., LIU L.: Co-segmentation of 3d shapes via subspace clustering. In Computer Graphics Forum (2012), vol. 31, Wiley Online Library, pp. 1703–1713. 3

[HKG11] HUANG Q., KOLTUN V., GUIBAS L.: Joint shape segmentation with linear programming. In ACM Transactions on Graphics (TOG) (2011), vol. 30, ACM, p. 125. 3

[JH99] JOHNSON A. E., HEBERT M.: Using spin images for efficient object recognition in cluttered 3d scenes. PAMI 21, 5 (1999), 433–449. 1

[JNRS10] JOURNÉE M., NESTEROV Y., RICHTÁRIK P., SEPULCHRE R.: Generalized power method for sparse principal component analysis. The Journal of Machine Learning Research 11 (2010), 517–553. 5

[JZvK07] JAIN V., ZHANG H., VAN KAICK O.: Non-rigid spectral correspondence of triangle meshes. International Journal of Shape Modeling 13, 01 (2007), 101–124. 3

[KHS10] KALOGERAKIS E., HERTZMANN A., SINGH K.: Learning 3d mesh segmentation and labeling. ACM Transactions on Graphics (TOG) 29, 4 (2010), 102. 7, 13

[KKBL15] KEZURER I., KOVALSKY S. Z., BASRI R., LIPMAN Y.: Tight relaxation of quadratic matching. In Computer Graphics Forum (2015), vol. 34, Wiley Online Library, pp. 115–128. 2

[KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. ACM TOG (Proc. SIGGRAPH) 30, 4 (2011). 2, 9

[LF09] LIPMAN Y., FUNKHOUSER T.: Mobius voting for surface correspondence. ACM Transactions on Graphics (Proc. SIGGRAPH) 28, 3 (Aug. 2009). 2

[LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on (2005), vol. 2, IEEE, pp. 1482–1489. 10

[MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. In Visualization and mathematics III. Springer, 2003, pp. 35–57. 3, 7

[Mém07] MÉMOLI F.: On the use of Gromov-Hausdorff Distances for Shape Comparison. In Symposium on Point Based Graphics (2007), pp. 81–90. 2

[OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. ACM Transactions on Graphics (TOG) 31, 4 (2012), 30. 2, 10

[OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. CGF 29, 5 (2010), 1555–1564. 2, 10

[RPSS10] RUGGERI M. R., PATANÈ G., SPAGNUOLO M., SAUPE D.: Spectral-driven isometry-invariant matching of 3d shapes. International Journal of Computer Vision 89, 2-3 (2010), 248–265. 2

[SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft maps between surfaces. In Computer Graphics Forum (2012), vol. 31, Wiley Online Library, pp. 1617–1626. 2

[SOCG10] SKRABA P., OVSJANIKOV M., CHAZAL F., GUIBAS L.: Persistence-based segmentation of deformable shapes. In Computer Vision and Pattern Recognition Workshops (CVPRW) (2010), IEEE, pp. 45–52. 3

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In Computer graphics forum (2009), vol. 28, Wiley Online Library, pp. 1383–1392. 1, 3, 7

[SRGB14] SOLOMON J., RUSTAMOV R., GUIBAS L., BUTSCHER A.: Earth mover's distances on discrete surfaces. ACM Transactions on Graphics (TOG) 33, 4 (2014), 67. 2

[SvKK*11] SIDI O., VAN KAICK O., KLEIMAN Y., ZHANG H., COHEN-OR D.: Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering, vol. 30. ACM, 2011. 3

[SY11] SAHILLIOĞLU Y., YEMEZ Y.: Coarse-to-fine combinatorial matching for dense isometric shape correspondence. Computer Graphics Forum 30, 5 (2011), 1461–1470. 2

[SY13] SAHILLIOĞLU Y., YEMEZ Y.: Coarse-to-fine isometric shape correspondence by tracking symmetric flips. In Computer Graphics Forum (2013), vol. 32, Wiley Online Library, pp. 177–189. 10

[TBW*09] TEVS A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.-P.: Isometric registration of ambiguous and partial data. In Proc. CVPR (2009), pp. 1185–1192. 2

[TCL*13] TAM G. K., CHENG Z.-Q., LAI Y.-K., LANGBEIN F. C., LIU Y., MARSHALL D., MARTIN R. R., SUN X.-F., ROSIN P. L.: Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. Visualization and Computer Graphics, IEEE Transactions on 19, 7 (2013), 1199–1217. 2

[Vil03] VILLANI C.: Topics in optimal transportation. No. 58. American Mathematical Soc., 2003. 4

[VKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. Computer Graphics Forum 30, 6 (2011), 1681–1707. 2

[WHC*15] WEI L., HUANG Q., CEYLAN D., VOUGA E., LI H.: Dense human body correspondences using convolutional networks. arXiv preprint arXiv:1511.05904 (2015). 2

[YZ13] YUAN X.-T., ZHANG T.: Truncated power method for sparse eigenvalue problems. The Journal of Machine Learning Research 14, 1 (2013), 899–925. 5, 6

[ZB15] ZUFFI S., BLACK M. J.: The stitched puppet: A graphical model of 3d human shape and pose. In Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on (2015), IEEE, pp. 3537–3546. 2

[ZHD*01] ZHA H., HE X., DING C., SIMON H., GU M.: Bipartite graph partitioning and data clustering. In Proceedings of the tenth international conference on Information and knowledge management (2001), ACM, pp. 25–32. 3

[ZSCO*08] ZHANG H., SHEFFER A., COHEN-OR D., ZHOU Q., VAN KAICK O., TAGLIASACCHI A.: Deformation-driven shape correspondence. In Computer Graphics Forum (2008), vol. 27, Wiley Online Library, pp. 1431–1439. 2

9. Appendix: proofs

PROOF OF PROPOSITION 1. One has $W_{N-n}^{noise} = 1/(N-n)\sum_{k=1}^{N-n} Y_k$, where Y_k is a sum over the characteristic functions derived from the two feature functions ϕ_k and ψ_k. We can thus write $Y_k = G(\phi_k, \psi_k)$, where G is a measurable function. By assumption, we thus have that the functions Y_j for $j = 1, \cdots, N-n$ are i.i.d. By the Central Limit Theorem, W_{N-n}^{noise} thus converges to the expected value $\mathrm{E}(Y_k)$. By a symmetry argument, the expected value for each entry are the same. A small calculation shows that $\mathrm{E}(Y_k) = (1-\kappa)/(d_1 d_2 K) \times \mathbf{1}_{d_2,d_1}$.

We now consider the error term $E := W_{N-n}^{noise} - (1-\kappa)/(d_1 d_2 K) \times \mathbf{1}_{d_2,d_1}$. Every entry of this matrix $E_{h,l}$ is identically distributed with mean 0. By a symmetry argument, the variance σ of each entry is the same, thus does not depend on h and l. Tchebyshev's inequality then implies that for every $\delta > 0$, one has

$$P\left(|E_{h,l}|^2 \geq \frac{\delta^2}{d_1 d_2}\right) \leq \frac{(1-\kappa)\sigma^2}{d_1 d_2 N \delta^2}.$$

From the fact that if for every h, l $|E_{h,l}|^2 < \delta^2/(d_1 d_2)$, then $\|E\|_{\mathcal{F}} < \delta$, one has that $P(\|E\|_{\mathcal{F}} < \delta)$ is bounded from below by

$$P\left(E, \forall h, l |E_{h,l}|^2 < \frac{\delta^2}{d_1 d_2}\right) \geq 1 - \sum_{h,l} P\left(|E_{h,l}|^2 \geq \frac{\delta^2}{d_1 d_2}\right),$$

which implies the result since the sum is over $d_1 d_2$ terms.

PROOF OF PROPOSITION 2. We reorder the rows and columns of W such that it is of the form (2). We denote by R_i the sum of the coefficients of the row i of the matrix W_{Ω_1,Ω_2} if $i \in \{1, \cdots p\}$ or row $i-p$ of the matrix $W_{\Omega_1,S_2\setminus\Omega_2}$ if $i \in \{p+1, \cdots d_2\}$. Then for every $i \leq p < i'$, one has $R_i \geq R_{i'}$. Otherwise, there exists $i \leq p < i'$, such that the exchange the row i of W_{Ω_1,Ω_2} and the row $i'-p$ of $W_{\Omega_1,S_2\setminus\Omega_2}$ would strictly increase $\|W_{\Omega_1,\Omega_2}\|_{1,1}$. The contradiction comes from the fact that this exchange amounts to replace W_{Ω_1,Ω_2} with a neighbor matrix whose norm cannot be larger (by assumption). Now, we define $\mathcal{P}$ is being the thresholding function that outputs exactly p non-zero values and which respect the canonical order (if some coordinates are equal). We clearly have $\mathcal{P} \cdot W \mathbb{1}_{\Omega_1} = \mathbb{1}_{\Omega_2}$. The second equality comes from the same analysis with the columns of W_{Ω_1,Ω_2} and $W_{S_1\setminus\Omega_1,\Omega_2}$.

We now need to introduce the following definition. A matrix $W : \mathbb{R}^{d_1} \to \mathbb{R}^{d_2}$ is said to be δ-independent, if for any $x \in \{0,1\}^{d_1} \setminus \{\mathbf{1}_{d_1,1}\}$ and $i \neq j$, one has $|(Wx)_i - (Wx)_j| > \delta$, and if for any $y \in \{0,1\}^{d_2} \setminus \{\mathbf{1}_{d_2,1}\}$ and $i \neq j$ one has $|(W^T y)_i - (W^T y)_j| > \delta$.

PROOF OF PROPOSITION 3. Let $p = 4(1-k)^2\sigma^2 d_\infty^2/[(N-n)\delta^2]$. By applying Proposition 1 with $\delta/(2d_\infty)$, we know there a probability greater than $1-p$ that all the coefficient of $E = W_N^{all} - \kappa W_n^{init} - C\mathbf{1}_{d_2,d_1}$ are bounded from above by $\delta/(2d_\infty)$. Let now $x \in \{0,1\}^{d_1}$ and denote by $\|x\|_1$ its number of non-zero coordinates. The i^{th}-row of Ex is given by

$$(Ex)_i = (W_N^{all}x)_i - k(W_n^{init}x)_i - C\|x\|_1$$

and satisfies $|(Ex)_i| \leq \delta/2$. Using the previous equation, one has

$$(W_N^{all}x)_j - (W_N^{all}x)_i = (Ex)_j - (Ex)_i + \kappa[(W_n^{init}x)_j - (W_n^{init}x)_i].$$

Since $|(Ex)_j - (Ex)_i| < \delta$ and $|(W_N^{all}x)_j - (W_N^{all}x)_i| \geq \delta$ by δ-minimal assumption, then $(W_n^{init}x)_j - (W_n^{init}x)_i$ and $(W_N^{all}x)_j - (W_N^{all}x)_i$ are the same sign. Hence the sorting of the vectors $W_N^{all}x$ and $W_N^{all}x$ is the same for every $x \in \{0,1\}^{d_1} \setminus \mathbf{1}_{d_1,1}$. Applied to $x = \mathbb{1}_{\Omega_1}$, one gets $\mathcal{P}(W_N^{all}\mathbb{1}_{\Omega_1}) = \mathcal{P}(W_n^{init}\mathbb{1}_{\Omega_1})$. The same arguments applied to W^T allows to show that there is a probability greater than $1-p$ that for $\mathcal{Q}((W_N^{all})^T\mathbb{1}_{\Omega_2}) = \mathcal{Q}((W_n^{init})^T\mathbb{1}_{\Omega_2})$. Hence there is a probability greater than $1-2p$, that the conclusion of the proposition hold.

PROOF OF PROPOSITION 4. From the proof of Proposition 3, for every $x \in \{0,1\}^{d_1} \setminus \{\mathbf{1}_{d_1,1}\}$ and every $y \in \{0,1\}^{d_2} \setminus \{\mathbf{1}_{d_2,1}\}$, one has $\mathcal{P}_p(W_N^{all}x) = \mathcal{P}_p(W_n^{init}x)$ and $\mathcal{Q}_q((W_N^{all})^T y) = \mathcal{Q}_q((W_n^{init})^T y)$, which implies the conclusion.

PROOF OF PROPOSITION 5. We reorder the rows and columns of W such that it is of the form (2). Starting from $f^{(i)}$ (with $i \geq 1$), we compute $g^{(i)} := \mathcal{P}(Wf^{(i)})$. If $g^{(i)} = g^{(i-1)}$, the algorithm is done. Otherwise, there exists at least a row A_i of W_{Ω_1,Ω_2} and a row $C_{i'}$ of $W_{\Omega_1,S_2\setminus\Omega_2}$, such that the sum of the coefficients of A_i is strictly less than the sum of the coefficients of $C_{i'}$. We reorder the lines of W so that $\mathcal{P}(Wf^{(i)})$ is a decreasing vector. This operation exchanges $l \geq 1$ lines of A with l lines of $W_{\Omega_1,S_2\setminus\Omega_2}$ (and in particular A_i and $C_{i'}$), which implies that $\|A\|_{1,1}$ increases by a strictly positive value. The same happens when we compute $f^{(i+1)}$ from $g^{(i)}$.

We denote by η_L the minimum of $|l_1 \cdot \mathbf{1}_{1,p} - l_2 \cdot \mathbf{1}_{1,p}|$ over all sub-matrices $L = (l_1, l_2)$ of W of size $2 \times p$ such that $l_1 \cdot \mathbf{1}_{1,p} \neq l_2 \cdot \mathbf{1}_{1,p}$. Similarly, let η_C be the minimum of $|c_1 \cdot \mathbf{1}_{q,1} - c_2 \cdot \mathbf{1}_{q,1}|$ over all sub-matrices $L = (c_1, c_2)^T$ of W of size $q \times 2$ such that $c_1 \cdot \mathbf{1}_{q,1} \neq c_2 \cdot \mathbf{1}_{q,1}$. Since the two sets under which we take the minimum is finite, the number $\eta = \min(\eta_L, \eta_C)$ is strictly positive. At each iteration, $\|A\|_{1,1}$ increases by at least this value η. Since $\|A\|_{1,1}$ is bounded by $\|W\|_{1,1}$, the algorithm terminates.

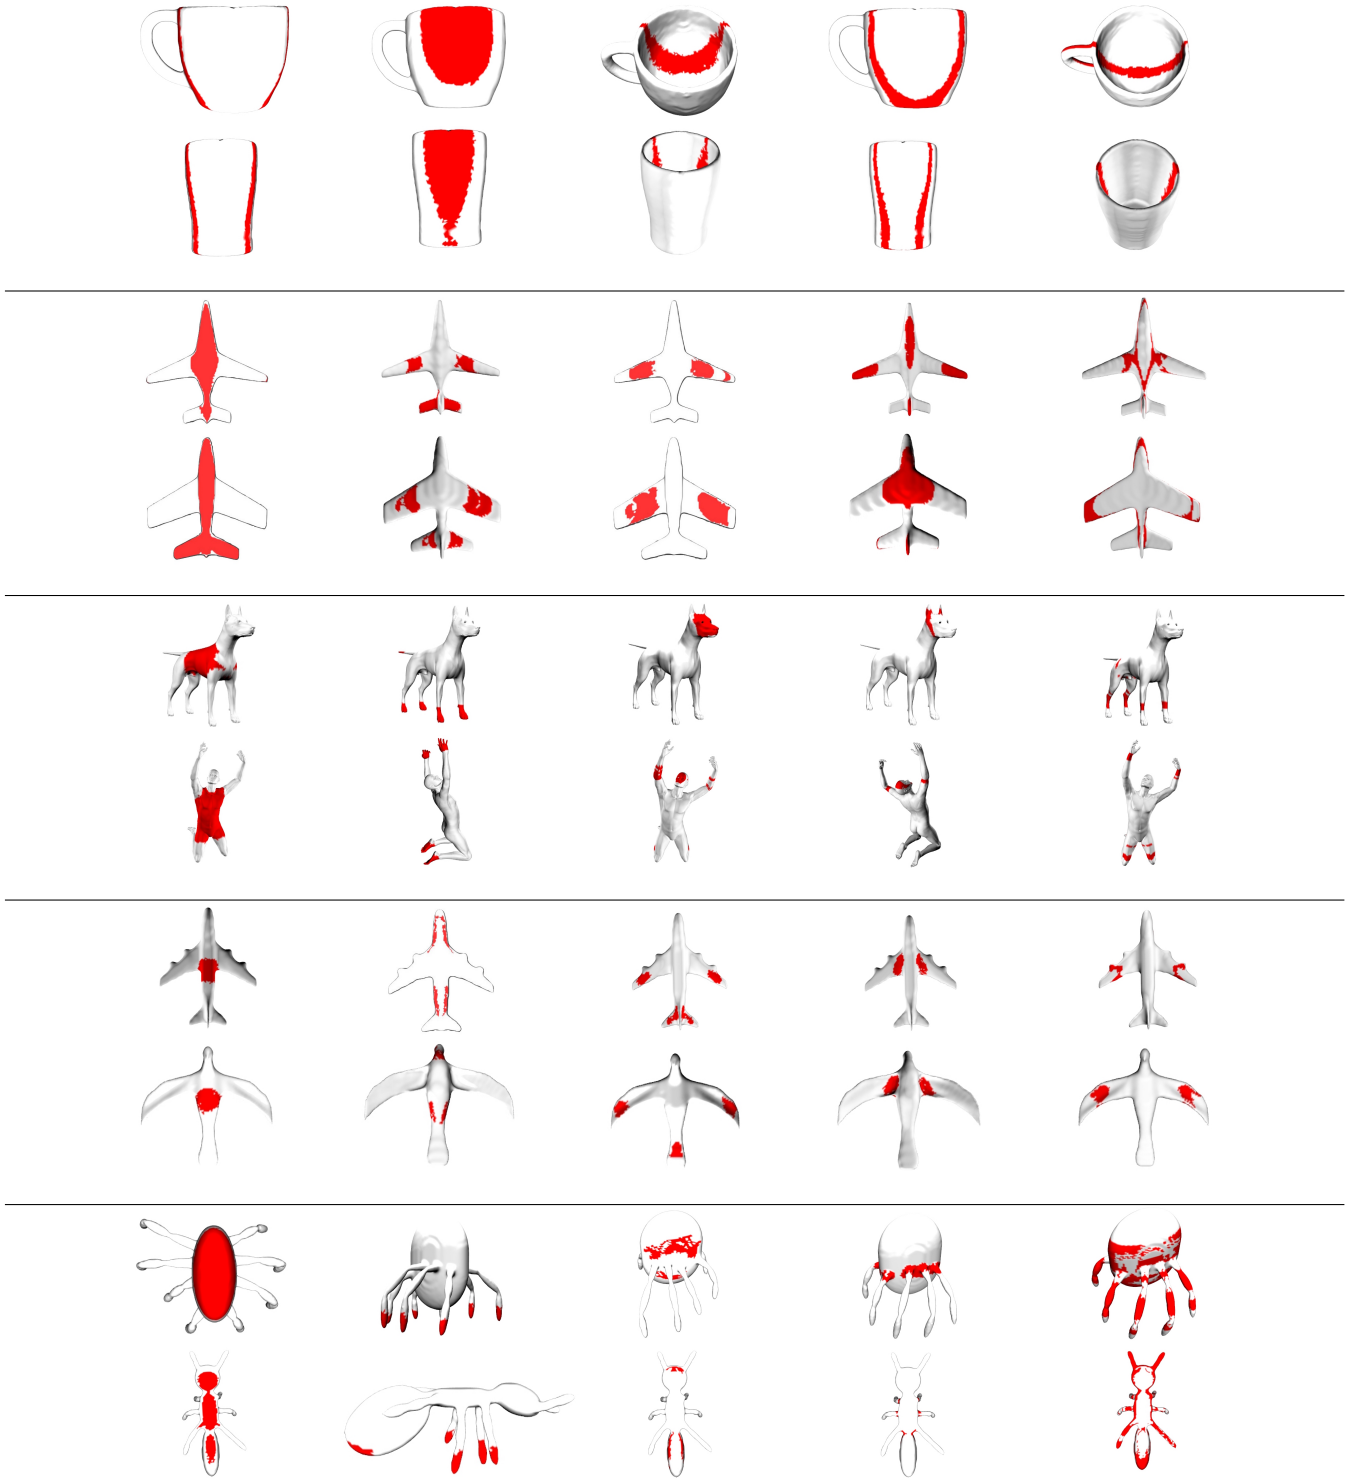

Figure 9: Every column refers to a certain stable region, while every row showcases the various stable regions of a certain shape. The stable regions have been constructed using the affinities of pairs of shapes on consecutive rows. The first two rows compute the correspondence between a pair of non-isometric cups in the Watertight dataset and the third and fourth rows compute the correspondence between a pair of non-isometric airplanes in the same dataset. The last three rows denote correspondences across different classes. The fifth and sixth rows denote the correspondences between the TOSCA shapes *dog* and *michael*. The seventh and eighth rows denote the correspondences between an airplane and a bird in the Watertight dataset and the ninth and tenth rows denote the correspondences between an ant and an octopus in the same dataset. The features used for the computation of the affinity matrix for all Watertight shapes are the features provided in the PSB dataset as described in [KHS10]. The features used for the computation of the affinity matrix for all TOSCA shapes are 100 HKS, 100 WKS and 100 mean curvature vectors.