

DWKS : A Local Descriptor of Deformations

Between Meshes and Point Clouds: Supplementary material

Robin Magnet
LIX, École Polytechnique
rmagnet@lix.polytechnique.fr

Maks Ovsjanikov
LIX, École Polytechnique
maks@lix.polytechnique.fr

Abstract

This document brings additional details on the spectral properties of Shape Difference Operators, among which a proof of the Theorem stated in the main paper. Additional qualitative and quantitative results are also given, comparing with more baselines for completeness.

1. Spectral Properties of Shape Difference Operators

1.1. Theoretical properties

Notations. Given two discrete shapes \mathcal{M} and \mathcal{N} , we note their respective Laplacian $L^{\mathcal{M}} = (A^{\mathcal{M}})^{-1}W^{\mathcal{M}}$ (resp. $L^{\mathcal{N}} = (A^{\mathcal{N}})^{-1}W^{\mathcal{N}}$) with A being a diagonal matrix filled with per-vertex areas, and W the stiffness matrix (e.g. the standard cotangent weight matrix [7]).

The discrete $L^2(\mathcal{S})$ inner product on shape \mathcal{S} is defined as

$$\langle f, g \rangle_{L^2(\mathcal{S})} = f^\top A^{\mathcal{S}} g \quad (1)$$

and the discrete $H_0^1(\mathcal{S})$ inner product as

$$\langle f, g \rangle_{H_0^1(\mathcal{S})} = \langle f, L^{\mathcal{S}} g \rangle_{L^2(\mathcal{S})} = f^\top W^{\mathcal{S}} g \quad (2)$$

Theorem 1 assumes the area-matrices $A^{\mathcal{M}}$ and $A^{\mathcal{N}}$ to be positive definite, which means no vertex has a 0 area. The assumption for a semi definite stiffness matrices however means the kernel of W only consists in the space of constant functions on the shape, which is the 0-set of the H_0^1 norm.

Proposition 1. *Under the assumptions of Theorem 1, the area-based shape difference operator is positive definite.*

Proof. The area-based shape difference operator V is defined implicitly via

$$\langle f, Vg \rangle_{L^2(\mathcal{M})} = \langle Ff, Fg \rangle_{L^2(\mathcal{N})} \quad \forall f, g \in L^2(\mathcal{M}) \quad (3)$$

Our discretized version of the L^2 inner product defines a positive-definite form since we suppose the area matrix A

to be non-degenerate. V is therefore self-adjoint operator with respect to the L^2 inner product.

Therefore, given $f \in L^2(\mathcal{M})$,

$$\langle f, Vf \rangle_{L^2(\mathcal{M})} = \|Ff\|_{L^2(\mathcal{N})}^2 \geq 0 \quad (4)$$

where the last inequality is an equality if and only if $f = 0$ in $L^2(\mathcal{M})$ since the functional map F is supposed to be non-degenerate.

V is therefore self-adjoint positive-definite, and the spectral theorem states its eigenvalues are real and positive. \square

Proposition 2. *Under the assumptions of Theorem 1, the conformal-based shape difference operator is positive semi-definite.*

Proof. The conformal-based shape difference operator R is defined implicitly via

$$\langle f, Rg \rangle_{H_0^1(\mathcal{M})} = \langle Ff, Fg \rangle_{H_0^1(\mathcal{N})} \quad \forall f, g \in H_0^1(\mathcal{M}) \quad (5)$$

Given non-degenerate matrices A and W , the H_0^1 inner product is positive definite (on H_0^1), and

$$\langle f, Rf \rangle_{H_0^1(\mathcal{M})} = \|Ff\|_{H_0^1(\mathcal{N})}^2 \geq 0 \quad (6)$$

for $f \in H_0^1(\mathcal{M})$ with equality if and only if $f = 0$ in $H_0^1(\mathcal{M})$ since F is supposed to be non-degenerate.

R is therefore self-adjoint positive definite on H_0^1 , and the spectral theorem states its eigenvalues are real and positive.

In practice, since operator R is extended to the entire L^2 space by setting it to 0 for constants, the operator is only positive semidefinite. \square

Proof of Theorem 1.

Proof. Under the assumption of Theorem 1, propositions 1 and 2 apply from which the result follows. \square

1.2. Practical Computation

In practice, given two discrete shapes \mathcal{M} and \mathcal{N} with a low dimensional functional map $\mathbf{C} \in \mathbb{R}^{k_{\mathcal{N}} \times k_{\mathcal{M}}}$ between them, the shape difference operators are computed in the spectral basis using the following formulas [9]:

$$\mathbf{V} = \mathbf{C}^\top \mathbf{C} \quad (7)$$

$$\mathbf{R} = (\Delta^{\mathcal{M}})^\dagger \mathbf{C}^\top \Delta^{\mathcal{N}} \mathbf{C} \quad (8)$$

where $\Delta^{\mathcal{M}}$ and $\Delta^{\mathcal{N}}$ are diagonal matrices of the first $k_{\mathcal{M}}$ (resp. $k_{\mathcal{N}}$) eigenvalues of the Laplace-Beltrami operator on \mathcal{M} (resp. \mathcal{N}).

The area-based shape difference operator \mathbf{V} being symmetric as shown by equation (7), its eigendecomposition is easily computed, while the eigen-decomposition of the conformal-based shape difference operator \mathbf{R} is obtained by solving a *generalized* eigenvalue problem

$$(\mathbf{C}^\top \Delta^{\mathcal{N}} \mathbf{C}) \Psi = \Delta^{\mathcal{M}} \Psi \Lambda \quad (9)$$

where $\mathbf{C}^\top \Delta^{\mathcal{N}} \mathbf{C}$ and $\Delta^{\mathcal{M}}$ are both symmetric and positive semi-definite.

1.3. Algebraic structure

In this section, we provide some insight on the discussion in section 4.1 of the main manuscript, namely about the claim that the spectrum of shape difference operators is better expressed using the log-scale.

By definition the shape difference operators between isometric shapes are identity operators, and all functions are preserved by these operators. Therefore the absence of deformations leads to 1 eigenvalues or 0 log-eigenvalues. More generally if a function f on \mathcal{M} only takes non-zero values on zones undergoing no deformation (regarding the one between \mathcal{M} and \mathcal{N}), then it will be preserved by the shape difference operator $D_{\mathcal{M},\mathcal{N}}$, which means it becomes a eigenvectors with a 1 eigenvalue.

It is straightforward to see the non-zero eigenvalues of $D_{\mathcal{N},\mathcal{M}}$ are the inverse of those from $D_{\mathcal{M},\mathcal{N}}$, which makes the log-eigenvalues opposite of each other.

The somewhat more complex case lies in the composition of shape difference operators. While for three shapes $\mathcal{M}, \mathcal{P}, \mathcal{N}$ the equality $D_{\mathcal{M},\mathcal{N}} = D_{\mathcal{M},\mathcal{P}} D_{\mathcal{P},\mathcal{N}}$ always holds, there is in general no simple relationship between the eigenvalues of each term. One scenario where the eigenvalues of the composition of the two operators are the product of the individual ones happens when one can find common eigenvectors for the two composed operators. This is known as codiagonalization of matrices and is only possible when the two operators commute.

2. Parameters

2.1. Optimization Objective

The optimization objective described in section 4.4 of the main manuscript consists in 5 different terms. Given a functional maps $\mathbf{C} \in \mathbb{R}^{k_{\mathcal{N}} \times k_{\mathcal{M}}}$, we discuss here the influence of each term in order to bring some intuition on how the hyperparameters can be tuned.

The first term $E_d(\mathbf{C}) = \|\mathbf{C}\mathbf{A} - \mathbf{B}\|^2$ simply enforces descriptor preservation with respect to the L^2 norm on the target shape \mathcal{N} , where descriptors are the columns of matrices \mathbf{A} and \mathbf{B} .

The second term $E_{dc}(\mathbf{C})$ was introduced in [5] in order to improve the descriptor preservation via the introduction of new operators built from individual descriptor functions.

The third term $E_l(\mathbf{C})$ corresponds to commutativity of the functional map with the two Laplace Beltrami operators, which enforces the functional map to represent a near-isometric map.

The fourth and fifth terms $E_c(\mathbf{C})$ and $E_a(\mathbf{C})$ introduced in [10], simply seek to preserve the action of the respectively conformal and area shape difference operators under the functional map.

2.2. Hyperparameters

DWKS. When computing DWKS descriptor, one has to chose a set of energy values (e_1, \dots, e_p) as well as a scale parameter σ . We advocate using moderate energy values to ignore extreme shape difference eigenvalues often created by noise in the functional correspondences. Using $e_1 = -\log 3$ and $e_p = \log 3$ has led to satisfying results on our side. The σ parameter describes how far an eigenvector will be spread on the energy-scale. To our knowledge there is no provably efficient procedure to fit this parameter for standard WKS descriptors [1], which is why we settle for a constant parameter across our experiments.

Optimization. Using the previous section, the optimization parameters are tuned so that the descriptors are well preserved as well as the action of the shape difference operators. Note that in our experiments using partial shapes, the near-isometric assumption doesn't hold which justifies the absence of the laplacian commutativity term in our experiments. Furthermore, because the DFAust collections contains many holes and cuts, we found it beneficial to set $\mu_a = 0$ and therefore ignore the area shape difference operators.

Projection to low dimension When projecting the computed pointwise map into a low-dimension functional map, we ignore the fraction α of vertices which are the hardest to match (as defined by the descriptor distance). This is due to

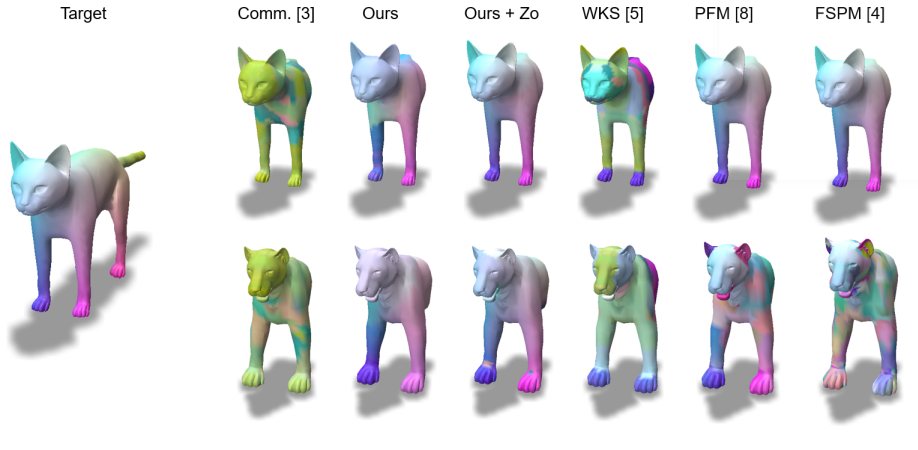


Figure 1. Results on the cats and lions dataset. Left mesh shows the base shape of the collection of complete cats. Right meshes display the computed pointwise maps in the case of a collection of partial cats (top row) and partial lions (bottom row).

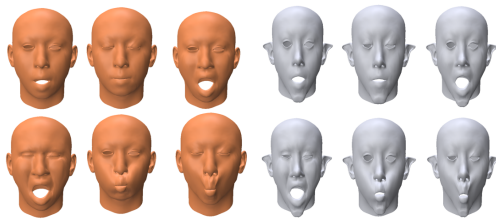


Figure 2. Visualization of the Synthetic faces dataset.

the fact shape difference operators and therefore DWKS descriptors can detect some intrinsic distortion near cuts and holes (usually due to noisy intra-collection maps). This explains why this parameter is set to a high value 20% when using partial shapes instead of the 5% for complete shapes.

Refinement. In the case of partial shape, we adapt the ZoomOut algorithm with step size $(1, \lfloor \frac{1}{\lambda} \rfloor)$ with λ the approximate ratio of area between the complete shape and the partial one, which eventually leads to a rectangular functional map.

3. Synthetic faces dataset

In the main manuscript, we refer to a synthetic faces dataset [9] shown on Figure 2, on which we compare our pipeline with the standard baseline [3] as well as the usual functional map pipeline [5] using WKS descriptors.

On this dataset, we show that our pipeline can ignore the area-based shape difference operators and still obtain great results unlike the pipeline from [3] as seen on Figure 3.

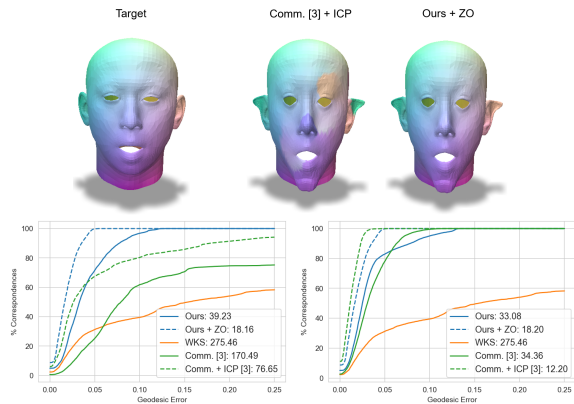


Figure 3. Results on the Synthetic Faces dataset. Top row shows pointwise maps after refinement when ignoring the area shape difference operators. Bottom row show the accuracy curves for this setting (left) and when using both type shape difference operators (right). Notice the stability of our method.

4. Comparison with partial spectral matching methods

The most direct competitor to our method is [3] extending the pipeline from [10], which leverages on intra-collection maps to compute cross-collection correspondences. We presented extensive comparisons with this approach in the main manuscript. However since these correspondences are only computed between the base shapes of each collection, any matching pipeline could be applied to this problem without using information given by the collections. Namely the standard functional correspondence computation problem [6] is known to be very efficient to compute correspondences between near-isometric shapes, and specific derivations have been obtained in [8, 4] in the

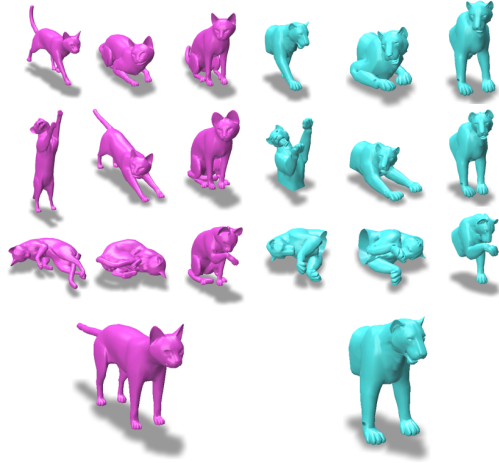


Figure 4. Visualization of the Cats and Lions dataset. The two bottom shapes are matched together using the deformations shown above them.

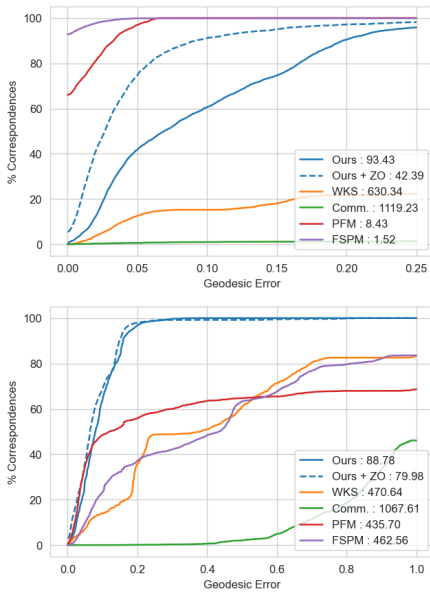


Figure 5. Quantitative results associated to results on Figure 1. Top graph displays accuracy results for the experiment matching a partial cat to a complete cat, bottom one those for the experiment marching a partial lion to a complete cat.

challenging scenario of partial matching.

In the following we present more quantitative and qualitative results on the two datasets presented in our paper, comparing to additional baselines [6, 8, 4].

Cats and Lions. We first focus on the Sumner dataset consisting in similar meshes of cats and lions as seen on Figure 4. Additionally to the experiment presented in the paper, we apply the same pipeline trying to match a collec-

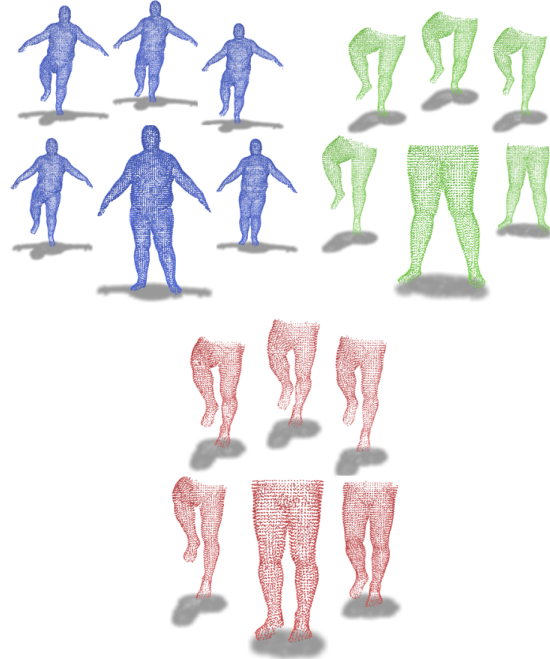


Figure 6. Visualization of the three collections corresponding to the jumping motion. The biggest shape represents the base shape of the collection and the smaller ones the used deformations.

tion of half-cat to the collection of complete cats. Qualitative results are displayed on Figure 1. Quantitative results associated to these two experiments are displayed on Figure 5 where the top graphs shows results in the case of a partial cat matched to a complete cat, and the bottom one those in the case of a partial lion matched to a complete cat. Note that in the first case, where the partial shape *is exactly isometric* to a subset of the complete shape, partial matching methods [8, 4] obtain excellent results, whereas performance drops very significantly in the second experiment in the case of near-isometry. Furthermore, the standard functional shape matching pipeline [6] fails to disambiguate left and right symmetry in both cases, and results from [3] are unable to obtain meaningful results. Our method however performs quite well in both cases, displaying an efficient use of common deformations to match vertices and disambiguate symmetries. We used the same parameters for both experiments, using the recommended values of [8, 4].

DFaust. We also compared our method to the same baselines on the DFaust [2] dataset. Note that since the dataset consists of pointclouds without normals, we replace the SHOT [11] descriptors in [4] by standard WKS descriptors [1], and remove comparison to [8] which requires information about faces.

Our datasets consists in scans of 3 individuals in two dif-

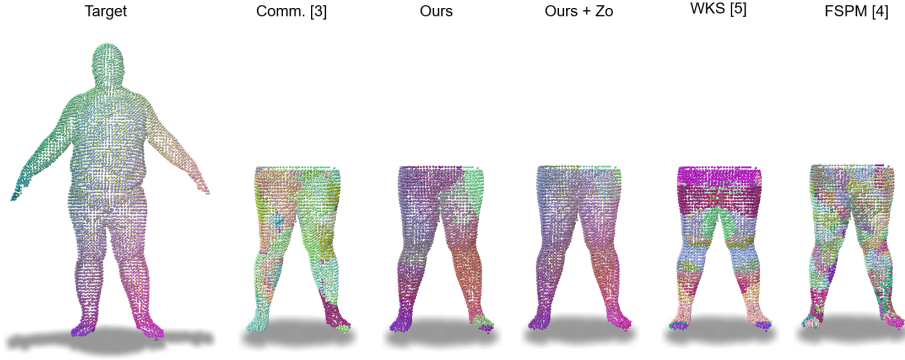


Figure 7. Examples of results on the DFaust dataset for the jumping motion. Left mesh shows the base shape of the first individual. Right meshes display the computed pointwise maps in the case of a collection of partial cats (top row) and partial lions (bottom row).

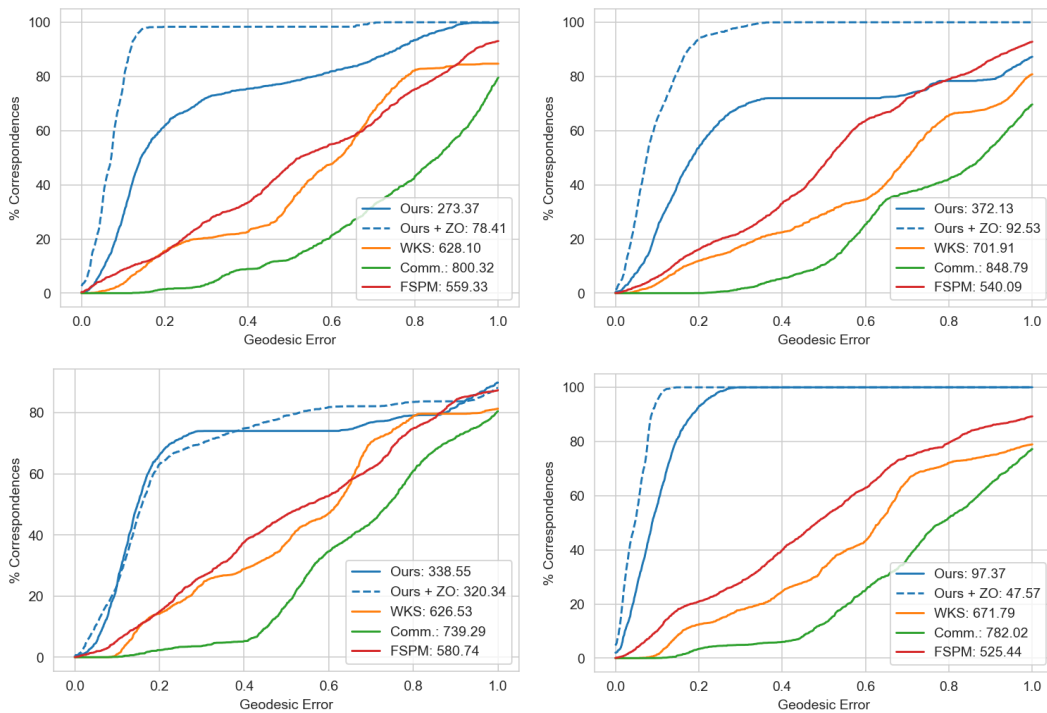


Figure 8. Results on the DFaust dataset. Each line show results for a given motion, and each column results for a given individual.

ferent motions, namely jumping on one leg and running on spot, which results in 6 collections of point clouds. Given a motion we match the related collection of the first individual, which consists in full scans, with the similar collections of the second and third individual, which consist in scans of the lower half of their bodies as shown on Figure 6.

Qualitative result using the second individual and the jumping motion are shown on Figure 7. Accuracy evaluations in all 4 cases are given on Figure 8, where the first line displays results when matching respectively the second (left) and third (right) individual to the first one using the

jumping motion, and the second line similar results but using the running motion. We remark again on Figure 7 that the standard functional map method [6] is unable to disambiguate left and right, and that in this case both partial matching [4] and the collection based method [3] can't to produce meaningful correspondences. On the other hand, our method clearly disambiguates the left-right symmetry, and strongly benefits from the ZoomOut refinement steps. This claim is reinforced by the accuracy results given on Figure 8, which show that our method significantly outperforms existing baselines.

References

- [1] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1626–1633, Barcelona, Spain, Nov. 2011. IEEE. 2, 4
- [2] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering Human Bodies in Motion. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5573–5582, Honolulu, HI, July 2017. IEEE. 4
- [3] Aharon Cohen and Mirela Ben-Chen. Robust Shape Collection Matching and Correspondence from Shape Differences. *Computer Graphics Forum*, 39(2):555–568, May 2020. 3, 4, 5
- [4] O. Litany, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Fully Spectral Partial Shape Matching. *Computer Graphics Forum*, 36(2):247–258, May 2017. 3, 4, 5
- [5] D. Nogneng, S. Melzi, E. Rodolà, U. Castellani, M. Bronstein, and M. Ovsjanikov. Improved Functional Mappings via Product Preservation. *Computer Graphics Forum*, 37(2):179–190, May 2018. 2, 3
- [6] Dorian Nogneng and Maks Ovsjanikov. Informative Descriptor Preservation via Commutativity for Shape Matching. *Computer Graphics Forum*, 36(2):259–267, May 2017. 3, 4, 5
- [7] Ulrich Pinkall and Konrad Polthier. Computing Discrete Minimal Surfaces and Their Conjugates. *Experimental Mathematics*, 2(1):15–36, Jan. 1993. 1
- [8] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial Functional Correspondence: Partial Functional Correspondence. *Computer Graphics Forum*, 36(1):222–236, Jan. 2017. 3, 4
- [9] Raif M. Rustamov, Maks Ovsjanikov, Omri Azencot, Mirela Ben-Chen, Frédéric Chazal, and Leonidas Guibas. Map-based exploration of intrinsic shape differences and variability. *ACM Transactions on Graphics*, 32(4):1, July 2013. 2, 3
- [10] Nitzan Shapira and Mirela Ben-Chen. Cross-Collection Map Inference by Intrinsic Alignment of Shape Spaces. *Computer Graphics Forum*, 33(5):281–290, Aug. 2014. 2, 3
- [11] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique Signatures of Histograms for Local Surface Description. In Kostas Daniilidis, Petros Maragos, and Nikos Paragios, editors, *Computer Vision – ECCV 2010*, Lecture Notes in Computer Science, pages 356–369, Berlin, Heidelberg, 2010. Springer. 4