

SHREC'16: Matching of Deformable Shapes with Topological Noise

Z. Lähler^{1†}, E. Rodolà^{1,2†}, M. M. Bronstein^{2†}, D. Cremers^{2†},
O. Burghard³, L. Cosmo⁴, A. Dieckmann³, R. Klein³, Y. Sahillioglu⁵

¹TU Munich, Germany, ²University of Lugano, Switzerland, ³University of Bonn, Germany
⁴Università Ca' Foscari Venezia, Italy, ⁵Middle East Technical University, Turkey



Figure 1: With this benchmark we evaluate the accuracy of matching algorithms when presented with 3D deformable shapes undergoing topological changes. Shown in the figure is a subset of shapes from the proposed dataset. The red marks indicate parts of the shape where a topological “shortcut” takes place: there, the triangular mesh is modified so as to avoid any self-intersections, and parts that would otherwise be located in the interior of the shape are completely removed.

Abstract

A particularly challenging setting of the shape matching problem arises when the shapes being matched have topological artifacts due to the coalescence of spatially close surface regions – a scenario that frequently occurs when dealing with real data under suboptimal acquisition conditions. This track of the SHREC'16 contest evaluates shape matching algorithms that operate on 3D shapes under synthetically produced topological changes. The task is to produce a pointwise matching (either sparse or dense) between 90 pairs of shapes, representing the same individual in different poses but with different topology. A separate set of 15 shapes with ground-truth correspondence was provided as training data for learning-based techniques and for parameter tuning. Three research groups participated in the contest; this paper presents the track dataset, and describes the different methods and the contest results.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Shape Analysis

1. Introduction

Matching deformable 3D shapes is a widespread problem having applications in numerous fields, and as such has been actively investigated in the computer vision and graphics communities for more than a decade. To date, a wide variety of approaches have been proposed to tackle the problem of (nearly-)isometric shape

matching with different levels of robustness against deviation from isometry and geometric noise (see [vKZHC01] for a recent survey). Given the large availability of methods, however, standard data sets typically concentrate on the type of non-rigid deformation (*i.e.*, change in pose and in shape class) and only include minor factors of nuisance, such as small holes, sampling artifacts, and numerical noise. By contrast, the topological change that triangular meshes undergo due to scanning artifacts and to the coalescence of spatially close surface regions (see Fig. 2) have been

† Organizers




Figure 2: Inaccuracy of the scanning device and suboptimal measurement conditions often result in structural instability of the acquired data. This is usually manifested as the topological merging of neighboring parts of the surface mesh (marked above in red on simulated rangemaps) and missing parts [CRB*16].

much less investigated during the years; few datasets include such instances, and these are concentrated in small regions. For example, the SHREC'10 [BBC*10] and SHREC'11 [BBB*11] robust correspondence tracks included synthetic topological “shortcuts” of just a few triangles on high-resolution meshes; a more recent benchmark, FAUST [BRLB14], included real-world scanings of different individuals with missing parts (due to occlusions) and topological shortcuts (due to touching limbs), both limited to small areas. A parallel track of the SHREC'16 contest evaluating robustness to scanning artifacts, concentrates on meshes with missing parts only [CRB*16].

Among the different approaches, few have targeted the class of topological transformations. Ovsjanikov *et al.* [OSG08] investigated resilience to topological shortcuts in the context of intrinsic symmetry detection of deformable shapes. They observed that the GPS embedding [Rus07] of a shape is stable against topological noise, since small shortcuts induce changes only at the higher end of the Laplacian spectrum. In the area of shape matching, Wang *et al.* [WBBP11] took a similar perspective by considering the metrics induced by commute-time kernels as a more robust alternative to geodesic distance, within a minimum-distortion (isometric embedding) correspondence framework. Noting that the diagonal of such kernels provides the GPS embedding of the shape, this approach shares the same limits as [OSG08]. Sparse relaxations to this framework were proposed by Rodolà *et al.* [RBA*12,RTH*13] with moderate increase in accuracy. Boscaini *et al.* [BGB14] considered a different kernel to yield better shape embeddings under small topological perturbations, and van Kaick *et al.* [vKZH13] proposed bilateral maps as shape features with partial resilience to geometric shortcuts. Purely geometric approaches working in Euclidean space such as non-rigid ICP (*e.g.*, [LSP08]) are more oblivious to topological changes, but suffer from sensitivity to initialization and only work well under small change in pose. More recently, Chen and Koltun [CK15] reformulated the isometric embedding problem with a robust norm accounting for topological artifacts, and additionally considered an extrinsic term to regularize the matching. In the realm of geometric deep learning, Boscaini *et al.*

al. [BMR*16] recently proposed a CNN-based shape descriptor demonstrating empirical resilience to topological changes.

With this track, we introduce a new dataset with large topological artifacts, in some cases extending to the entire shape, and coming in two different resolutions. The dataset is designed to provide a challenging test bed for the top performing methods in shape matching, and measure quantitatively their performance under topological deformations – a crucial capability, required to deal with nowadays ever growing collections of acquired 3D models.

2. The dataset

The dataset includes modified shapes from the KIDS set [RRBW*14], together with new shapes produced specifically for the contest. The shapes were created with DAZ 3D Studio 4.9 by changing the pose of a fixed human template, hence allowing to retain the ground-truth pointwise correspondence across all poses. All shapes were then post-processed by removing interior parts and resolving self-intersections, resulting in a manifold mesh of the outer hull [CK10]. This procedure modifies the topology of each shape along several regions and to various extents, creating a diverse but realistic dataset.

A total of 25 shapes were created for the benchmark, 15 of which are provided as training data, and the remaining 10 are used for the contest. The training set comes with complete ground-truth correspondence as well as a left-right symmetric map for each shape, mapping each point to its symmetric counterpart; note that, as a result of the post-processing, some points may not have a symmetric match. Similarly, due to remeshing artifacts along the areas of self-intersection, not all vertices are associated with a reliable match. These vertices (covering 0.5-5% of the whole surface area) are excluded from the final evaluation, and are indicated by a binary map that is included in the dataset.

The entire dataset exists in *low resolution* ($\sim 10K$ vertices per shape) and *high resolution* ($\sim 50K-70K$ vertices) variants, which are evaluated separately. The dataset is available for download at <http://vision.in.tum.de/~laehner/shrec2016/>.

3. Evaluation measures

Each participating method is asked to retrieve vertex-to-vertex correspondences for each pair of the test set, amounting to 90 full matching problems in total. According to standard practice, symmetric solutions are accepted with no penalty. Both sparse and dense solutions are considered in the evaluation.

For the evaluation of the correspondence quality, we followed the Princeton benchmark protocol [KLF11]. Assume that a correspondence algorithm produces a pair of points $(x, y) \in \mathcal{M} \times \mathcal{N}$ between shapes \mathcal{M} and \mathcal{N} , whereas the ground-truth correspondence is (x, y^*) . Then, the inaccuracy of the correspondence is measured as

$$\varepsilon(x) = \frac{d_{\mathcal{N}}(y, y^*)}{\text{area}(\mathcal{N})^{1/2}}, \quad (1)$$

and has units of normalized length on \mathcal{N} (ideally, zero). Here $d_{\mathcal{N}}$ is the geodesic distance on \mathcal{N} . The value $\varepsilon(x)$ is averaged over all

pairs of shapes $(\mathcal{M}, \mathcal{N})$. We plot cumulative curves showing the percent of matches which have error smaller than a variable threshold. In separate tables, in order to distinguish between sparse and dense methods, we also show percentages of matched area.

4. Methods

Three methods were evaluated in this benchmark: the isometric embedding method of [SY12], an embedding technique based on the Green's function on manifolds, and a learning technique based on random forests [RRBW*14].

4.1. Isometric embedding (EM)

The method described in [SY12] seeks for a minimum-distortion correspondence among nearly isometric shapes. The algorithm starts by producing a consistent sampling on the two given shapes, and then seeks for an injective mapping that minimizes the metric distortion between the two point sets. In order to achieve robustness to topological changes, biharmonic distance [LRF10] is used as a metric in place of geodesic distance. The alignment process operates on the Euclidean embeddings of the two shapes, obtained by classical multidimensional scaling with respect to the chosen metric. Specifically, the problem is cast in a probabilistic setting to a log-likelihood maximization problem, which is then solved via expectation-maximization (EM). The optimization alternates between generating a one-to-one mapping via bipartite perfect matching, and iteratively refining this mapping by a greedy algorithm.

This method is evaluated in the sparse setting (~ 250 matches per pair of shapes) on both the low and high resolution benchmarks.

4.2. Green's function embedding alignment (GE)

Burghard, Dieckmann, and Klein propose to utilize the Green's function of the Laplace-Beltrami operator Δ to match shapes. Let \mathcal{M} be a manifold, ϕ_1, ϕ_2, \dots and $\lambda_1, \lambda_2, \dots$ be the spectral decomposition of the Laplacian Δ , and $\mathcal{L}^2(\mathcal{M})$ be the set of square integrable functions on \mathcal{M} . The Green's function is defined as the solution of $\Delta g_x = \delta_x$. Given a function f_0 the equation $\Delta g_x = f_0$ has a solution only if $\langle \phi_1, f_0 \rangle = 0$, and if g_x is a solution then there is a whole family of solutions, namely $(g_x + \mu \phi_1)_{\mu \in \mathbb{R}}$. One therefore seeks for the function g_x that is the solution to the two equations $\Delta g_x = \delta_x - \phi_1 \langle \phi_1, \delta_x \rangle$ and $\langle \phi_1, g_x \rangle = 0$. We now have a unique function $g_x \in \mathcal{L}^2(\mathcal{M})$ for each point $x \in \mathcal{M}$ and can identify each point with this function g_x :

$$\pi : \mathcal{M} \rightarrow \mathcal{L}^2(\mathcal{M}) \quad x \mapsto \sum_{i=2}^k \phi_i \frac{\phi_i(x)}{\lambda_i}$$

Note, that this embedding approximately preserves the metric because $\|\pi(x) - \pi(y)\|^2 = \sum_{i=2}^k \left(\frac{\phi_i(x) - \phi_i(y)}{\lambda_i} \right)^2$ are the biharmonic distances [LRF10]. Also the embedding is invariant to extrinsic deformations, because all involved quantities are intrinsic. For the purpose of establishing correspondences one can now align those embeddings. This is done in an iterative procedure resulting in solutions of increasing size. For this purpose, the method assumes to have a small set of fixed predefined sparse correspondences \mathcal{C}_0 (4 were used for the final evaluation, see Fig. 3) and a second set of

dense correspondences \mathcal{C}_1 , that resulted from the last iteration and are initially empty ($\mathcal{C}_1 = \emptyset$).




Figure 3: Examples of user-given sparse matches used by the GE method.

For two different shapes \mathcal{M} and \mathcal{N} , the embeddings are in different spaces, namely $\mathcal{L}^2(\mathcal{M})$ and $\mathcal{L}^2(\mathcal{N})$. Mapping between these spaces involves a functional map $T_F : \mathcal{L}^2(\mathcal{M}) \rightarrow \mathcal{L}^2(\mathcal{N})$ [OBBS*12]. Given two sets of sparse correspondences $\mathcal{C}_0, \mathcal{C}_1 \subset \mathcal{M} \times \mathcal{N}$, the method chooses G as the linear map which minimizes

$$\frac{1}{|\mathcal{C}_0|} \sum_{p,q \in \mathcal{C}_0} \|g_p^{\mathcal{M}} - G g_q^{\mathcal{N}}\|^2 + \frac{1}{|\mathcal{C}_1|} \sum_{p,q \in \mathcal{C}_1} \|g_p^{\mathcal{M}} - G g_q^{\mathcal{N}}\|^2 + \epsilon \|G\|^2$$

where both sets \mathcal{C}_0 and \mathcal{C}_1 have an equal weight and ϵ is generally a really small value (10^{-8}) required mostly in the first step where too few correspondences make the system underconstrained ($\mathcal{C}_1 = \emptyset$).

For each point on the source and each point on the target, a closest matching point is extracted and the correspondences are stored in \mathcal{C}_1 :

$$x \in \mathcal{M} \mapsto \arg \min_{y \in \mathcal{N}} \|g_x^{\mathcal{M}} - G g_y^{\mathcal{N}}\|$$

$$y \in \mathcal{N} \mapsto \arg \min_{x \in \mathcal{M}} \|g_x^{\mathcal{M}} - G g_y^{\mathcal{N}}\|$$

Finally, the following steps are iterated until convergence (40 iterations): (1) solve for G using \mathcal{C}_0 and \mathcal{C}_1 and (2) calculate new correspondences and store in \mathcal{C}_1 . In order to scale well with the number of points, both embeddings are subsampled to 1000 points via farthest point sampling. Correspondences in (2) are only retrieved from and to this subset. Only in the very last iterations the correspondences are retrieved over the dense point sets.

4.3. Random forests (RF)

A modified version of the learning-based technique described in [RRBW*14]. This approach uses the training set of 15 shapes to train an ensemble of decision trees. Each decision tree assigns, to each point of a test shape, a probability distribution defined on a discrete label set, where each label identifies a set of corresponding points from the training data (at most one point per training shape). The path along each tree is determined by means of binary decision functions that evaluate a prescribed point feature (taken to be the WKS [ASC11] in [RRBW*14]) with random parametrizations. This randomized feature selection allows to retain the full power of the prescribed intrinsic feature without resorting to a pre-defined parametrization, which might not be optimal for all points of the shape; at the same time it limits the correlation among trees, thus ensuring good generalization. The output of the forest is finally regularized by using the language of functional maps [OBBS*12].




Figure 4: Best and worst matching pairs (left and right column respectively) obtained on the high resolution dataset. Each row corresponds to a different method. Corresponding points have the same color, while grey shade denotes no match; note how GE tends to avoid areas where a topological change takes place – see for example the hands and nose on the first pair, and the face and knees on the second pair of shapes.

Figure 5: Best and worst matching pairs on the low resolution dataset. Note how certain shapes consistently show up as being either “easy” or “difficult” across different methods, regardless of the measure of error minimized by each algorithm.

In this implementation of [RRBW*14], the WKS feature is replaced with the HKS [SOG09], which is a *local* feature and, as such, less susceptible to topological changes. The feature is expressed as a combination of k Laplacian (squared) eigenfunctions over diffusion times t_i for $i = 1, \dots, T$; k and T constitute the parameters for the random forest, which is trained over 15 trees. As a regularization step, the landmark-based procedure followed in [RRBW*14] is substituted with a simple “low-pass” filtering of the forest prediction, obtained by projecting the predicted correspondence onto the first 90 Laplacian eigenfunctions, and recovering the underlying pointwise map by nearest neighbors in the spectral domain [OBCS*12].

This approach produces dense correspondences for all pairs of shapes, and is evaluated on the low and high resolution datasets.

5. Results

We evaluated the different methods quantitatively and qualitatively on the provided data. In Fig. 6 we show a comparison on the low and high resolution benchmarks separately, using the error measure defined in Section 3. Note that these plots include both sparse and dense solutions, and do not take into account the percentage of area

matched among each pair of shape. Such a measure is instead provided in Table 1, together with other statistics including the best and worst solutions obtained by each method.

In Fig. 7 we show a more detailed breakdown of each method. From these plots we note that GE and RF tend to yield more consistent results, with RF having a better worst-case behavior (see red curves in the third row), although never reaching 100% on any of the shape pairs. At the same time, EM yields $\sim 20\%$ (around 50 out of 250) *exact* matches on the best matching pairs, while RF reaches up to $\sim 40\%$ *exact* matches on the low resolution dataset and $\sim 30\%$ on high resolution data. These latter results reflect the nature of RF being essentially a classifier: The method attempts to maximize classification accuracy (*i.e.*, label assignment), which in this context is equivalent to the fraction of points matched with zero geodesic error. GE performs in line with RF on average, outperforming the latter in terms of maximum accuracy, and producing more dense solutions than both RF and EM.

Finally, qualitative examples of the best and worst solutions obtained by each method are shown in Figs. 4 and 5.

6. Conclusions

We evaluated a set of matching algorithms for establishing correspondences between deformable 3D shapes undergoing topological transformations. As a first observation, we note that none of the

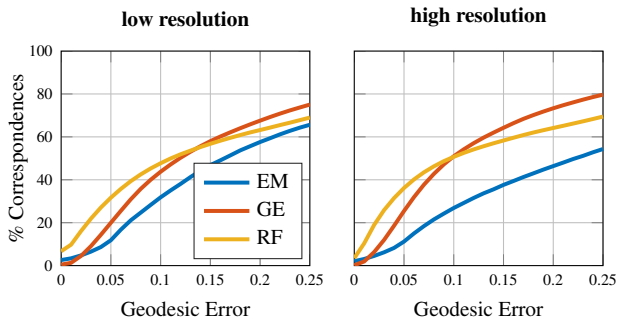


Figure 6: Comparisons among different methods on the low (left) and high resolution (right) benchmarks. The curves are averaged over all pairs of shapes. See Fig. 7 for a more detailed breakdown.

evaluated methods prevails over all the others, and all have similar accuracy on average. Learning-based techniques (RF) produce more consistent results across different matching instances, while methods based on minimizing the metric distortion (EM, GE) offer higher maximum accuracy.

A second key conclusion arising from this evaluation is that, despite the promising results displayed by the competing approaches, the problem of deformable shape matching under topological changes is still far from being solved. As a direct comparison, top performing approaches in nearly-isometric shape matching, including the ones evaluated in this track, achieve 100% correspondences within a geodesic error of 0.1 – whereas in this evaluation the top performing method attained $\sim 50\%$ at the same threshold, and did not reach 100% at all within the considered window (see Fig. 6).

Matching under topological changes is an open research problem that deserves to be taken into full consideration, for its huge practical value in dealing with today’s increasing amount of real-world 3D data. We hope that these results will renew the interest of the community in this challenging problem, and foster further research in this direction.

		% area matched	AUC (avg)	(best)	(worst)
low res	EM	2.48%	0.29	0.81	0.02
	GE	31.09%	0.45	0.79	0.02
	RF	28.59%	0.48	0.75	0.24
high res	EM	0.87%	0.20	0.83	0.04
	GE	49.04%	0.51	0.76	0.19
	RF	17.87%	0.49	0.80	0.22

Table 1: Comparisons between different methods in terms of total surface area matched on the target shape (averaged over all shape pairs) and normalized area-under-the-curve (ideally 1.0). The AUC refers to the curves described in section 3 and shown in Fig. 7. We show the average AUC over all shape pairs (second column), and the AUC obtained on the best and worst shape pairs (third and fourth columns). Best results are reported in bold.

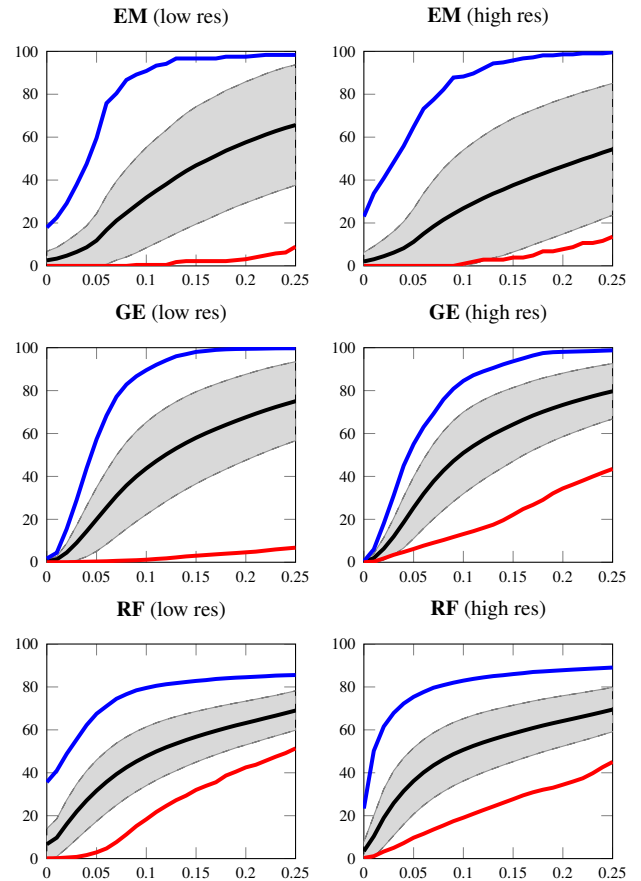


Figure 7: Performance comparison in terms of best matching pair (blue curve), worst pair (red curve), and average accuracy (black curve) with standard deviation (shaded area). Each row shows a different method, as indicated in the plot titles.

References

- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV Workshops* (2011). 3
- [BBB*11] BOYER E., BRONSTEIN A. M., BRONSTEIN M. M., ET AL.: SHREC 2011: Robust feature detection and description benchmark. In *Proc. 3DOR* (2011), pp. 71–78. 2
- [BBC*10] BRONSTEIN A. M., BRONSTEIN M. M., CASTELLANI U., ET AL.: SHREC 2010: robust correspondence benchmark. In *Proc. 3DOR* (2010), pp. 87–91. 2
- [BGB14] BOSCAINI D., GIRDIUŠAS R., BRONSTEIN M. M.: Coulomb shapes: Using electrostatic forces for deformation-invariant shape representation. In *Proc. 3DOR* (2014), pp. 9–15. 2
- [BMR*16] BOSCAINI D., MASI J., RODOLÀ E., BRONSTEIN M. M., CREMERS D.: Anisotropic diffusion descriptors. *Computer Graphics Forum* 35, 2 (2016). 2
- [BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3d mesh registration. In *Proc. CVPR* (2014). 2
- [CK10] CAMPEN M., KOBBELT L.: Exact and robust (self-)intersections for polygonal meshes. *Comput. Graph. Forum* 29, 2 (2010), 397–406. 2

- [CK15] CHEN Q., KOLTUN V.: Robust nonrigid registration by convex optimization. In *Proc. ICCV* (2015). 2
- [CRB*16] COSMO L., RODOLÀ E., BRONSTEIN M. M., ET AL.: SHREC'16: Partial matching of deformable shapes. In *Proc. 3DOR* (2016). 2
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T. A.: Blended intrinsic maps. *TOG* 30, 4 (2011), 79. 2
- [LRF10] LIPMAN Y., RUSTAMOV R., FUNKHOUSER T.: Biharmonic distance. *ACM Transactions on Graphics* 29, 3 (June 2010). 3
- [LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. In *Proc. SGP* (2008), pp. 1421–1430. 2
- [OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4 (July 2012), 30:1–30:11. 3, 4
- [OSG08] OVSJANIKOV M., SUN J., GUIBAS L.: Global intrinsic symmetries of shapes. In *Proc. SGP* (2008), pp. 1341–1348. 2
- [RBA*12] RODOLÀ E., BRONSTEIN A., ALBARELLI A., BERGAMASCO F., TORSELLO A.: A game-theoretic approach to deformable shape matching. In *Proc. CVPR* (June 2012), pp. 182–189. 2
- [RRBW*14] RODOLÀ E., ROTA BULÒ S., WINDHEUSER T., VESTNER M., CREMERS D.: Dense non-rigid shape correspondence using random forests. In *Proc. CVPR* (2014), pp. 4177–4184. 2, 3, 4
- [RTH*13] RODOLÀ E., TORSELLO A., HARADA T., KUNIYOSHI Y., CREMERS D.: Elastic net constraints for shape matching. In *Proc. ICCV* (December 2013), pp. 1169–1176. 2
- [Rus07] RUSTAMOV R. M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP* (2007), pp. 225–233. 2
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP* (2009). 4
- [SY12] SAHILLIOĞLU Y., YEMEZ Y.: Minimum-distortion isometric shape correspondence using EM algorithm. *IEEE Trans. Pattern Anal. Mach. Intell.* 34, 11 (2012), 2203–2215. 3
- [vKZH13] VAN KAICK O., ZHANG H., HAMARNEH G.: Bilateral maps for partial matching. *Comput. Graph. Forum* 32, 6 (Sept. 2013), 189–200. 2
- [vKZHCO11] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. 1
- [WBBP11] WANG C., BRONSTEIN M. M., BRONSTEIN A. M., PARAGIOS N.: Discrete minimum distortion correspondence problems for non-rigid shape matching. In *Proc. SSVM* (2011), pp. 580–591. 2