

SHREC'16: Partial Matching of Deformable Shapes

L. Cosmo^{1†}, E. Rodolà^{2†}, M. M. Bronstein^{2†}, A. Torsello^{1†}, D. Cremers³, Y. Sahillioglu⁴

¹University of Venice, Italy, ²University of Lugano, Switzerland
³TU Munich, Germany, ⁴Middle East Technical University, Turkey




Figure 1: A subset of shapes from the benchmark. Participants are required to produce a point-wise correspondence (either sparse or dense) between a full template in a neutral pose (left) and its deformed versions with missing parts. The dataset includes 8 shape classes, for a total of 599 3D models. In this figure corresponding points have the same color, while shape boundaries are marked by a red contour.

Abstract

Matching deformable 3D shapes under partiality transformations is a challenging problem that has received limited focus in the computer vision and graphics communities. With this benchmark, we explore and thoroughly investigate the robustness of existing matching methods in this challenging task. Participants are asked to provide a point-to-point correspondence (either sparse or dense) between deformable shapes undergoing different kinds of partiality transformations, resulting in a total of 400 matching problems to be solved for each method – making this benchmark the biggest and most challenging of its kind. Five matching algorithms were evaluated in the contest; this paper presents the details of the dataset, the adopted evaluation measures, and shows thorough comparisons among all competing methods.

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

1. Introduction

Shape correspondence is a fundamental problem in computer vision and graphics, with a wide range of applications ranging from texture mapping to reconstruction [vKZHC011]. A particularly challenging setting arises when the shapes in question are allowed to undergo non-rigid deformations, which are typically assumed to be approximately isometric (such a model appears to be good for, e.g., human body poses). Even more challenging is *partial correspondence*, where one is shown only a subset of the shape and has to match it to a deformed full version thereof. Instances of this problem arise in numerous applications that involve real data ac-

quisition by 3D sensors, inevitably leading to missing parts due to occlusions or partial view.

In the rigid setting (e.g., for 3D scan completion), partial correspondence problems have been tackled by ICP-like approaches such as [AMCO08, ART15]. Attempts to extend these ideas to the non-rigid case [LSP08] had limited success due to sensitivity to initialization and to the underlying assumption of small deformations. In the non-rigid realm, several metric approaches centered around the notion of minimum distortion correspondence [BBK06] have been proposed. Bronstein *et al.* [BB08, BBBK09] combine metric distortion minimization with optimization over *regular* (i.e., contiguous) matching parts. Rodolà *et al.* [RBA*12, RTH*13] relaxed the regularity requirement by allowing sparse correspondences. Sahillioglu and Yemez [SY14] proposed a voting-based formula-

† Organizers




Figure 2: Example shapes from the cuts (top) and holes (bottom) datasets – the difference lies in the regularity of the missing parts.

tion to match shape extremities, which are assumed to be preserved by the partiality transformation. The aforementioned methods are based on intrinsic metric preservation and on the definition of spectral features, hence their accuracy suffers at high levels of partiality – where the computation of these quantities becomes unreliable due to boundary effects and meshing artifacts. More recent works include the alignment of tangent spaces [BWW*14] and the design of robust descriptors for partial matching [vKZH13]; in the context of collections of shapes, partial correspondence has been considered in [HWG14, CRA*16]. More recently, Masci *et al.* [MBBV15] introduced a deep learning framework for computing dense correspondences between deformable shapes, providing a generalization of the convolutional networks (CNN) to non-Euclidean manifolds. Later, Wei *et al.* [WHC*15] focused on matching *human* shapes undergoing changes in pose by means of classical CNNs, also tackling partiality transformations.

Previous benchmarks only partly investigated the robustness of matching methods to partiality transformations. The SHREC'11 track on robust correspondence [BBB*11] included “partial” and “view” classes to assess resilience to partiality. Relatively more focus has been put in the area of *retrieval* (as opposed to correspondence) from partial views of full *rigid* models, see, *e.g.*, the SHREC'13 track on shape-from-rangemap retrieval [SMB*13]. A parallel track of SHREC'16 [LRB*16] investigates correspondence between deformable shapes from simulated 3D acquisitions, but concentrates on the topological changes that the meshes incur (*i.e.*, topological “shortcuts”) rather than partiality.

With this benchmark, we investigate the robustness of deformable shape matching algorithms to partiality and missing parts. We do so on a big dataset (599 shapes) subdivided into 8 shape classes (humans and animals) undergoing near-isometric transformations. Partiality here assumes a *regular* and *irregular* structure, resulting in two sets of shapes over which the evaluation is performed separately (see Fig. 5). It is important to note that the present benchmark is focused on the *partial-to-full* scenario, namely matching a deformed partial shape to a full template in a neutral pose (see Fig. 1). The *partial-to-partial* setting is deferred to future investigation.

2. The dataset

The benchmark expands upon the datasets presented in [RCB*16]. As base models, we use shapes from the TOSCA dataset [BBK08], consisting of 76 nearly-isometric shapes subdivided into 8 classes (the class *gorilla* was removed due to non-manifold artifacts). Each class comes with a “null” shape in a standard pose (extrinsically bilaterally symmetric), which is used as the full template mesh to which partial shapes are matched during the evaluation. In order to make the datasets more challenging and avoid compatible triangulations, all shapes were remeshed to 10K vertices by iterative pair contractions [GH97]. After remeshing, missing parts are introduced in the following ways, giving rise to two different datasets:

Regular cuts. The template shape of each class was cut with a plane at 6 different orientations. The six cuts were then transferred to the remaining poses using the ground-truth correspondence, resulting in 320 partial shapes in total.

Irregular holes. Given a shape and an “area budget” determining the fraction of area to keep (40%, 70%, and 90%), we produced additional shapes by an erosion process applied to the surface. Seed holes were placed at 5, 25, and 50 Euclidean farthest samples over the shape; the holes were then enlarged to meet the specified area budget. The total number of shapes produced this way is 279.

See Fig. 5 for examples of shapes from the two datasets. Note that all shapes are composed by exactly one connected component. Shapes inside each dataset present different amounts of missing surface, ranging approximately from 10% to 60% of missing area.

Due to the remeshing process, ground-truth matches between each partial shape and the corresponding template are *sub-vertex*, and specified in barycentric coordinates on the mesh triangles. Each dataset is split into a training set (120 and 79 shapes) and a test set (200 shapes per dataset). For the former, sub-vertex ground-truth matches and symmetric left-right maps are provided for each shape; the latter is used for the final evaluation.

3. Evaluation measures

Each participating method is asked to retrieve *sub-vertex* point-to-point correspondences between each partial shape in the test set and the full template from the corresponding class, amounting to 400 matching problems in total. Following standard practice, intrinsically symmetric solutions are accepted with no penalty. Both sparse and dense solutions are considered in the evaluation.

Correspondence quality is measured according to the Princeton benchmark protocol [KLF11]. Assume that a correspondence algorithm produces a pair of points $(x, y) \in \mathcal{M} \times \mathcal{N}$ between partial shape \mathcal{M} and template \mathcal{N} , whereas the ground-truth correspondence is (x, y^*) . Then, the inaccuracy of the correspondence is measured by the *geodesic error*:

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

and has units of normalized length on \mathcal{N} (ideally, zero). Here $d_{\mathcal{N}}$ is the geodesic distance on the template shape \mathcal{N} . The value $\varepsilon(x)$ is averaged over all matching instances $(\mathcal{M}, \mathcal{N})$. We plot cumulative curves showing the percent of matches which have error smaller than a variable threshold.

4. Methods

Five methods were evaluated in the benchmark, namely: partial functional correspondence [RCB*16], the isometric embedding method of [SY12], game-theoretic matching [RBA*12], elastic net matching [RTH*13], and a learning technique based on random forests [RRBW*14]. None of the evaluated methods included sub-vertex correspondences, *i.e.*, all produced vertex-to-vertex solutions.

4.1. Partial functional maps (PFM)

The matching technique proposed by Rodolà *et al.* [RCB*16] as an extension to the functional maps framework [OBCS*12] in order to deal with partial shapes.

In this framework, shape correspondence is modeled as a linear operator $T_F : L^2(\mathcal{M}) \rightarrow L^2(\mathcal{N})$, mapping functions on shape \mathcal{M} to functions on shape \mathcal{N} via the composition $T_F(f) = f \circ T^{-1}$, where $T : \mathcal{M} \rightarrow \mathcal{N}$ is a bijective mapping between the two shapes. Assume the two function spaces are equipped with orthogonal bases $\{\varphi_i\}_{i \geq 1} \in L^2(\mathcal{M})$ and $\{\psi_j\}_{j \geq 1} \in L^2(\mathcal{N})$. Then, since T_F is a linear operator, it can equivalently be represented by a matrix \mathbf{C} with coefficients $c_{ij} = \langle T_F(\varphi_i), \psi_j \rangle$. Seeking a functional correspondence among the two shapes then amounts to determining the unknown \mathbf{C} that better preserves certain pointwise features or other mapping constraints [OBCS*12].

As a convenient choice for the aforementioned bases, it has been proposed to adopt the eigenfunctions of the respective Laplacians on the two shapes. In particular, the manifold Laplacian yields an eigen-decomposition $\Delta_{\mathcal{M}}\varphi_i = \lambda_i\varphi_i$ for $i \geq 1$, with eigenvalues $0 = \lambda_1 < \lambda_2 \leq \dots$ and eigenfunctions $\{\varphi_i\}_{i \geq 1}$ forming an orthonormal basis of $L^2(\mathcal{M})$. In case the two shapes to be compared are related by a near-isometry, the equality $\psi_i = \pm\varphi_i \circ T^{-1}$ holds (approximately) for all $i \geq 1$, which leads to the matrix representation \mathbf{C} of the functional map being diagonal, $c_{ij} = 0$ if $i \neq j$.

In case one of the two shapes has holes or missing parts, the functional representation of the correspondence still has a meaningful structure. Let \mathcal{M} be a partial shape, \mathcal{N} a full shape, and let $\mathcal{N}' \subset \mathcal{N}$ be the region of \mathcal{N} corresponding to \mathcal{M} under a near-isometry $T : \mathcal{M} \rightarrow \mathcal{N}'$. Then, for each eigenfunction φ_i of \mathcal{M} there exists an eigenfunction ψ_j of \mathcal{N} for some $j \geq i$, such that $\psi_j \approx \pm\varphi_i \circ T^{-1}$ [RCB*16]. In other words, the eigenfunctions of the Laplacian are still compatible under partiality, but some eigenfunctions of the full shape do not have a corresponding counterpart on the partial shape. This results in a matrix \mathbf{C} manifesting a slanted diagonal structure (see inset), with an angle depending on the area ratio of the two surfaces [RCB*16]. Using this knowledge as a prior, the method alternates between optimizing for the correspondence \mathbf{C} and for the matching part on the full shape (see Fig. 3). As a result, the matching algorithm produces dense correspondences for all shapes.





Figure 3: The partial functional maps (PFM) approach of [RCB*16] alternates between dense correspondence (bottom row) and matching part (top row) until convergence.

of the two shapes to be matched is a nearly isometric part of the other up to an arbitrary scale. The isometrically deformed partial shape and its complete version are first sampled [HS85] into point sets S and T , respectively, where $|S| = |T| = 10$. The method then seeks for an optimal partial map $\phi : S \rightarrow T$ with minimum distortion. Since two shapes are never perfectly isometric, even partly, due to imperfections of the modeling process and geometry discretization errors, it is not usually possible to find a zero distortion mapping. Hence the goal becomes minimization of the following *scale-invariant* metric distortion function:

$$D_{\text{iso}}(\phi) = \frac{1}{|\phi|} \sum_{(s_i, t_j)} d_{\text{iso}}(s_i, t_j), \quad (2)$$

where $d_{\text{iso}}(s_i, t_j)$ is the contribution of the individual correspondence (s_i, t_j) to the overall distortion:

$$d_{\text{iso}}(s_i, t_j) = \frac{1}{\binom{|\phi'|}{2}} \sum_{((s_a, t_b), (s_c, t_d)) \in \mathcal{C}(\phi')} |\rho(s_i, t_j; s_a, t_b) - \rho(s_i, t_j; s_c, t_d)| \quad (3)$$

with $\phi' = \phi - \{(s_i, t_j)\}$ and $\mathcal{C}(\phi')$ denoting the set of all pairwise combinations from ϕ' . The ratio function $\rho(s_i, t_j; s_k, t_l)$ is then written in terms of raw geodesic distances, for a given $(s_k, t_l) \in \phi$:

$$\rho(s_i, t_j; s_k, t_l) = \max \left(\frac{g(s_i, s_k)}{g(t_j, t_l)}, \frac{g(t_j, t_l)}{g(s_i, s_k)} \right) \quad (4)$$

where $g(\dots)$ is the raw geodesic distance between two surface points. This definition of metric distortion is based on the observation that the ratios between geodesic distances on a surface remain unchanged under scaling and isometric deformations. Hence if S and T are sampled consistently from the given arbitrarily scaled (partially) isometric shapes, one can find an optimal mapping ϕ^* such that $D_{\text{iso}}(\phi^*) = 0$ (Fig. 4 left). To make the problem tractable,




Figure 4: Left: Demonstration of the scale-invariant distortion measure D_{iso} . The ratios between geodesic distances remain invariant under scaling and isometric deformation: $\rho(s_i, t_j; s_a, t_b) = \rho(s_i, t_j; s_c, t_d)$. Right: Overview of the combinatorial matching algorithm from [SY12].

$M = 5$ evenly spaced vertices are sampled from S , and D_{iso} is computed for $M! \binom{M}{M}$ possible one-to-one mappings, promoting the mapping that yields the minimum distortion as the optimal mapping ϕ^* (Fig. 4 right). Finally, this sparse correspondence ϕ^* is extended into a dense one of size 100 by filling a cost matrix based on ϕ^* and running a minimum-weight perfect matching on it [Kol09].

4.3. Game-theoretic matching (GT)

The game-theoretic matching technique proposed by Rodolà *et al.* [RBA*12] estimates sparse correspondences between two shapes X and Y by minimizing an intrinsic measure of distortion. Given two pairs of corresponding points (x, y) and (x', y') , the quality of the correspondences can be quantified in terms of the intrinsic metric d_X and d_Y by measuring to which extent the distance between x and x' measured on X matches the distance between the corresponding points y and y' measured on Y ,

$$\varepsilon(x, y, x', y') = |d_X(x, x') - d_Y(y, y')|. \quad (5)$$

Here the authors adopt a relaxed notion of correspondence defined in terms of an indicator function $u : X \times Y \rightarrow [0, 1]$ such that, for every measurable subsets $A \subseteq X$ and $B \subseteq Y$,

$$\int_A \int_Y u dy dx = \int_A dx; \quad \int_B \int_X u dx dy = \int_B dy. \quad (6)$$

Using this relaxed notion of correspondence, one obtains a class of distortion metrics, known as *Gromov-Wasserstein* metrics, defined as

$$D(X, Y) = \frac{1}{2} \inf_u \|\varepsilon\|_{L^p(u \times u)}, \quad (7)$$

where $1 \leq p \leq \infty$, and

$$\|\varepsilon\|_{L^p(u \times u)}^p = \int_{(X \times Y)^2} \varepsilon^p(x, y, x', y') u(x, y) u(x', y') dx dy dx' dy'. \quad (8)$$

Given this formalization, the estimation problem is cast into an evolutionary game-theoretic framework where strategies are modeled as candidate assignments $(x, y) \in X \times Y$ based on some measure of pointwise similarity among the surface points. To this end, Rodolà and coauthors [RBA*12] use HKS [SOG09] descriptors with the standard Euclidean distance since they demonstrate good resilience to a variety of deformations. Then, they consider the assignment constraint $\mathbf{u} \in \Delta$ where $\mathbf{u} \equiv \text{vec}\{\mathbf{U}\}$ is the correspondence vector, constrained to lie in the standard mn -simplex

$$\Delta = \{\mathbf{u} \in \mathbb{R}^{mn} : \mathbf{u}^T \mathbf{1} = 1 \text{ and } \mathbf{u} \geq 0\}. \quad (9)$$

Finally, the correspondence estimation is turned into the maximization of the mutual similarity between correspondence, captured by a $mn \times mn$ matrix \mathbf{A} whose elements measure the similarity between pairs of correspondences which, in terms of the Gromov-Wasserstein metric with $p = 2$, results in setting $a_{(ij)(i'j')} = \exp(-\alpha \varepsilon_{ij i'j'}^2)$. This gives rise to the maximization problem

$$\max \mathbf{u}^T \mathbf{A} \mathbf{u} \text{ s.t. } \mathbf{u} \in \Delta. \quad (10)$$

In this framework, the matching problem is better interpreted

as an inlier selection problem in which matches form a coherent group according to the given pairwise distortion metric. In this scenario players pre-programmed according to a fixed strategy are repeatedly selected from a common population to play a symmetric two-player game. As the game is repeated, players adopting strategies that yield larger payoffs are positively selected, resulting in a selection process where inconsistent hypotheses are driven to extinction. This gives rise to sparse correspondences between the involved shapes.

4.4. Elastic net (EN)

The elastic net matching framework by Rodolà *et al.* [RTH*13] is a direct generalization of spectral matching [LH05] and of the game-theoretic matching technique [RBA*12]. As in the GT matching framework, sparse correspondences between two shapes X and Y are estimated by minimizing the intrinsic distortion of a set of correspondences represented in terms of the same fuzzy indicator function $u : X \times Y \rightarrow [0, 1]$ used in the previous approach. The distortion is, thus, measured in terms of the Gromov-Wasserstein metrics (7) resulting in the relaxed Quadratic Assignment Problem (QAP)

$$\begin{aligned} \min_{\mathbf{U}} \text{vec}\{\mathbf{U}\}^T \mathbf{A} \text{vec}\{\mathbf{U}\} \\ \text{s.t. } \mathbf{U} \mathbf{1} \leq 1, \mathbf{U}^T \mathbf{1} \leq 1, \mathbf{U} \geq 0, \end{aligned} \quad (11)$$

where $\text{vec}\{\mathbf{U}\}$ is the $m^2 n^2$ -dimensional column-stack vector representation of the correspondence matrix \mathbf{U} , \mathbf{A} is the $mn \times mn$ non-negative symmetric cost matrix registering the pairwise distortion terms $a_{(ij)(i'j')} = \exp(-\alpha \varepsilon_{ij i'j'}^2)$, and \succeq denotes element-wise inequality.

In order to incorporate a notion of stability into the matching process, the author cast the problem as one of model-fitting, seeking a good approximation of the true relationship between the two shapes, *i.e.*, the optimal correspondence x^* , with deviation measured in the Gromov-Wasserstein distance. Problems of this kind are often studied with the tools of regression analysis, where the interest shifts to the extraction of the relations connecting the variables underlying the possible assignments $\{x_i\}_{i=1..n}$. Here candidate matches act as explanatory variables, while we seek to find the combination that best describes the data in the minimal distortion sense. Under this view, these variables may be correlated, and it can be of interest to determine groups of highly correlated predictors, as they will likely form consistent matches.

In this view, spectral matching can be directly related to ridge regression, whose L^2 penalty is known to generally improve conditioning of the problem, yet always keeping all the predictors in the model. Similarly, the game-theoretic technique is equivalent to lasso regression, where the sparsity-inducing L^1 regularizer (9) performs continuous shrinkage and automatic variable selection.

The elastic net framework attempts to strike a balance between these two behaviors. This is obtained by substituting the L^1 and L^2 constraints with a family of constraints known as elastic net [ZH05]. This regularization technique shares with the lasso the ideal property of performing automatic variable selection, and most notably it is able to select entire groups of highly correlated variables. The elastic net criterion is defined as a convex combination

of the lasso and ridge penalties:

$$(1 - \alpha)\|x\|_1 + \alpha\|x\|_2^2, \quad \alpha \in [0, 1]. \quad (12)$$

It becomes ridge regression for $\alpha = 1$, and the lasso for $\alpha = 0$. This penalty function is singular at 0 and *strictly* convex (differently from the lasso) for $\alpha > 0$, thus possessing the characteristics of both penalties.

This leads to the following family of relaxations for the QAP:

$$\begin{aligned} \min_x \quad & x^T A x \\ \text{s.t.} \quad & (1 - \alpha)\|x\|_1 + \alpha\|x\|_2^2 = 1, \quad x \succeq 0, \end{aligned} \quad (13)$$

with $\alpha \in [0, 1]$ regulating the trade-off between size of the correspondence and matching error, allowing to fine tune the model complexity ranging from the highly selective pure lasso for $\alpha = 0$ to the more tolerant ridge behavior for $\alpha = 1$.

The value $\alpha = 0.8$ was used in these evaluations; being essentially a variant of the GT approach, the method produces sparse correspondences.

4.5. Random forests (RF)

A modified version of the learning-based technique described in [RRBW*14]. This approach uses the training set of 199 shapes to train a collection of random forests, one forest per shape class. Each decision tree in the forest 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. The path along each tree is determined by means of binary decision functions that evaluate a prescribed point feature with random parametrizations. This randomized feature selection allows to retain the full power of the 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.

In this implementation of [RRBW*14], the WKS [ASC11] feature is replaced with the HKS [SOG09], which is a *local* feature and, as such, less susceptible to the boundary effects induced by partiality. The feature is a T -dimensional vector per point, where each dimension is expressed as the nonlinear combination:

$$f_j(x) = \sum_{i=1}^k e^{-\lambda_i t_j} \varphi_i(x)^2, \quad (14)$$

where λ_i and φ_i for $i = 1, \dots, k$ are respectively the Laplacian eigenvalues and eigenfunctions of the shape, and t_j for $j = 1, \dots, T$ are diffusion times. k and T constitute the parameters of the random forest, which is trained over 15 trees. A separate forest is trained for each class of the dataset (for a total of 8 forests), and is subsequently applied to the corresponding class during the test phase.

As a regularization step, the landmark-based procedure followed in [RRBW*14] is substituted with a simple “low-pass” filtering of the forest prediction: The predicted correspondence is converted into a functional map by using the first 90 Laplacian eigenfunctions on both shapes (see Section 4.1), and the underlying point-wise map is then recovered by nearest neighbors in the spectral domain [OBCS*12]. This approach produces dense correspondences for all shapes in the benchmark.



Figure 5: Some examples of good solutions obtained by the competing methods on the cuts (left pairs) and holes (right pairs) benchmarks.

5. Results

In Fig. 6 we show quantitative comparisons of all matching algorithms using the error measure defined in Section 3, while in Fig. 7 we evaluate the behavior of each method across increasing amounts of partiality. Qualitative examples of the solutions obtained by each method are given in Fig. 5, and average numbers of matches are reported in Table 1.

An immediate result that is evident from the plots is that the

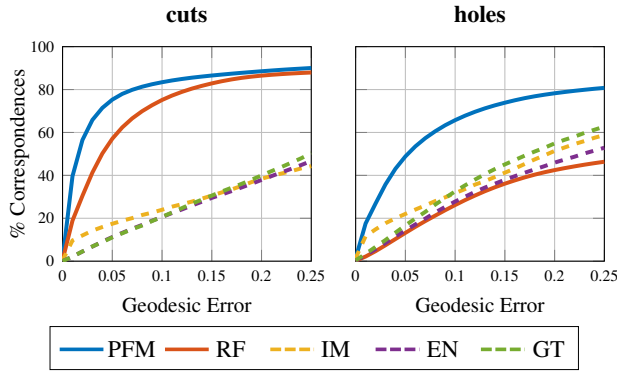


Figure 6: Comparisons on the cuts (left) and holes (right) datasets. Each curve (solid for dense methods, dashed for sparse) is averaged over all shapes in all classes (200 shapes per dataset).

two considered datasets yield wildly different results, with *holes* being the more challenging of the two. This is especially evident on PFM, where the performance drops by 10%, and RF, where the drop in accuracy is as big as 40%. The latter technique in particular becomes the worst performing method among all competitors on the *holes* dataset. The reason for this behavior is probably to be found in the fact that the two approaches make use of spectral quantities: In particular, RF suffers from the negative effect that the presence of boundaries exerts on the chosen point descriptors, which increases with the length of the boundary. This effect is less pronounced in the case of PFM, since the latter method makes use of *extrinsic*, local descriptors as data (see [RCB*16] for details), whose calculation is only marginally influenced by the presence of boundaries.

In contrast, methods based on minimizing the metric distortion (IM, EN, GT) give results of comparable quality on both datasets. However, average accuracy is not very high due to the distortion that the considered metrics undergo in the presence of missing surface regions.

6. Discussion and conclusions

Compared to the traditional “full-to-full” counterpart, the problem of partial shape matching has received surprisingly limited focus from the community. However, in this era of 3D data acquisition, the problem is gaining more and more practical relevance and is one of the key challenges that need to be tackled.

With this benchmark we explored some of the current approaches, and compared them across a big dataset and at various amounts of partiality. Overall, methods based on minimizing met-

	PFM	RF	IM	EN	GT
<i>cuts</i>	dense	dense	61.3	87.8	51.0
<i>holes</i>	dense	dense	78.2	112.6	76.4

Table 1: Average number of matches obtained by each method on the two datasets.




Figure 7: Correspondence quality of each method at increasing levels of partiality (measured as percentage of missing area).

ric distortion (IM, GT, EN) seem to have limited success; this is probably due to the instability of the metrics under the topological changes induced by partiality on one hand, and to the sensitivity of the point features used to simplify the problem on the other. Methods based on machine learning (RF) demonstrate more resilience due to the presence of (dense) training data, which serves as an example for the kind of transformations that the matching system is likely to encounter. This suggests a possible avenue for further research in the direction of machine learning techniques applied to shape analysis. Finally, the method based on partial functional correspondence (PFM) yields the best results thanks to the strong prior on the structure of the map relating partial to full shape. However, how to extend these results to the case of partial-to-partial matching remains an open question, and an interesting direction of future research.

References

- [AMCO08] AIGER D., MITRA N. J., COHEN-OR D.: 4-points congruent sets for robust pairwise surface registration. *TOG* 27, 3 (2008), 85. 1
- [ART15] ALBARELLI A., RODOLÀ E., TORSELLO A.: Fast and accurate surface alignment through an isometry-enforcing game. *Pattern Recognition* 48 (2015), 2209–2226. 1
- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV Workshops* (2011). 5
- [BB08] BRONSTEIN A. M., BRONSTEIN M. M.: Not only size matters: regularized partial matching of nonrigid shapes. In *Proc. NORDIA* (2008). 1
- [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
- [BBBK09] BRONSTEIN A., BRONSTEIN M., BRUCKSTEIN A., KIMMEL R.: Partial similarity of objects, or how to compare a centaur to a horse. *IJCV* 84, 2 (2009), 163–183. 1
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS* 103, 5 (2006), 1168–1172. 1
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 2

- [BWW*14] BRUNTON A., WAND M., WUHRER S., SEIDEL H.-P., WEINKAUF T.: A low-dimensional representation for robust partial isometric correspondences computation. *Graphical Models* 76, 2 (2014), 70–85. 2
- [CRA*16] COSMO L., RODOLÀ E., ALBARELLI A., MÉMOLI F., CREMERS D.: Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum* (2016). 2
- [GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proc. SIGGRAPH* (1997), pp. 209–216. 2
- [HS85] HOCHBAUM D., SHMOYS D.: A best possible heuristic for the k-center problem. *Mathematics of Operations Research* 10 (1985), 180–184. 3
- [HWG14] HUANG Q., WANG F., GUIBAS L. J.: Functional map networks for analyzing and exploring large shape collections. *TOG* 33, 4 (2014), 36. 2
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T. A.: Blended intrinsic maps. *TOG* 30, 4 (2011), 79. 2
- [Kol09] KOLMOGOROV V.: Blossom V: A new implementation of a minimum cost perfect matching algorithm. *Mathematical Programming Computation (MPC)* (2009), 43–67. 4
- [LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *Proc. CVPR* (2005), vol. 2, pp. 1482–1489. doi:10.1109/ICCV.2005.20. 4
- [LRB*16] LÄHNER Z., RODOLÀ E., BRONSTEIN M. M., ET AL.: SHREC'16: Matching of deformable shapes with topological noise. In *Proc. 3DOR* (2016). 2
- [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. 1
- [MBBV15] MASCI J., BOSCAINI D., BRONSTEIN M. M., VANDEREGHEYNST P.: Geodesic convolutional neural networks on riemannian manifolds. In *Proc. 3dRR* (2015). 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, 5
- [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. 1, 3, 4
- [RCB*16] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum* (2016). 2, 3, 6
- [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. 3, 5
- [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. 1, 3, 4
- [SMB*13] SİPIRAN I., MERUANE R., BUSTOS B., ET AL.: SHREC'13 track: Large-scale partial shape retrieval using simulated range images. In *Proc. 3DOR* (2013), pp. 81–88. 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, 5
- [SY12] SAHILLIOĞLU Y., YEMEZ Y.: Scale normalization for isometric shape matching. *Computer Graphics Forum (Proc. Pacific Graphics)* 31, 7 (2012). 3
- [SY14] SAHILLIOĞLU Y., YEMEZ Y.: Partial 3-d correspondence from shape extremities. *Computer Graphics Forum* 33, 6 (2014), 63–76. 1
- [vKZH13] VAN KAICK O., ZHANG H., HAMARNEH G.: Bilateral maps for partial matching. *Computer Graphics Forum* 32, 6 (2013), 189–200. 2
- [vKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. 1
- [WHC*15] WEI L., HUANG Q., CEYLAN D., VOUGA E., LI H.: Dense human body correspondences using convolutional networks. *CoRR abs/1511.05904* (2015). URL: <http://arxiv.org/abs/1511.05904>. 2
- [ZH05] ZOU H., HASTIE T.: Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B* 67 (2005), 301–320. 4