

Consistent Shape Matching via Coupled Optimization

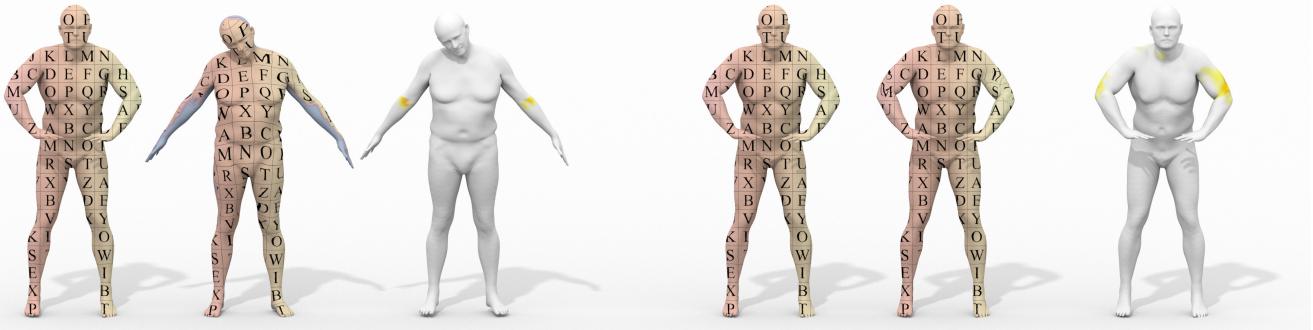
Omri Azencot^{1†}, Anastasia Dubrovina^{2†}, and Leonidas Guibas²¹University of California, Los Angeles
²Stanford University

Figure 1: Our method computes a point-to-point map and its inverse between a pair of shapes (left). Using the map, we transfer a texture from the source shape to the target, and we show the geodesic error obtained with our map, where white is no error and red is maximum error. In addition, we compose the map with its inverse, and use it to push a texture from the source shape to itself (right), obtaining accurately matched shapes, except for a negligible subset of points.

Abstract

We propose a new method for computing accurate point-to-point mappings between a pair of triangle meshes given imperfect initial correspondences. Unlike the majority of existing techniques, we optimize for a map while leveraging information from the inverse map, yielding results which are highly consistent with respect to composition of mappings. Remarkably, our method considers only a linear number of candidate points on the target shape, allowing us to work directly with high resolution meshes, and to avoid a delicate and possibly error-prone up-sampling procedure. Key to this dimensionality reduction is a novel candidate selection process, where the mapped points drift over the target shape, finalizing their location based on intrinsic distortion measures. Overall, we arrive at an iterative scheme where at each step we optimize for the map and its inverse by solving two relaxed Quadratic Assignment Problems using off-the-shelf optimization tools. We provide quantitative and qualitative comparison of our method with several existing techniques, and show that it provides a powerful matching tool when accurate and consistent correspondences are required.

1. Introduction

Estimating correspondences between three-dimensional (3D) shapes is an essential building block in many algorithms related to shape interpolation, deformation transfer and shape retrieval, to name a few. Different approaches usually focus either on pairs of shapes, or consider whole collections. The generated maps differ by their *density* (sparse vs. dense), *accuracy*, and *consistency*, i.e., whether compositions of maps along cycles lead to the identity

mapping. Often, correspondences between shape pairs are dense, less consistent and of moderate accuracy, whereas maps computed for shape collections are sparse, more consistent and of higher accuracy. In this work, we propose a novel mapping method for pairs of shapes, which utilizes the Quadratic Assignment Problem (QAP) formulation. It is solved using an efficient splitting method, yielding dense, accurate and approximately consistent correspondences.

To date, the majority of matching techniques designed for pairs of shapes concentrate on obtaining the map, but do not exploit or estimate its inverse. In practice, re-running the algorithm with exchanged roles for the source and target shapes will produce the

† These authors contributed equally to this work.

inverse map. However, high consistency is difficult to achieve, since the mappings are computed independently. Notable exceptions propose simultaneous optimization of low-dimensional matrices which represent the map and its inverse ([ERGB16], among others). Unfortunately, recovering dense maps in this setup is a challenging task in itself. Approaches guaranteeing bijective maps are also available and are based on compatible parametrization of shapes (e.g., [AL16]), or kernel density estimation directly in the space of permutation matrices [VLR^{*}17]. We opt to relax the bijectivity constraint, and perform joint optimization for the map and its inverse. Furthermore, since we directly solve for dense maps, a map recovery step is not needed.

Computing maps in shape collections was proven useful in Nguyen et al. [NBCW^{*}11] and following work. In particular, a requirement that compositions along cycles approximate the identity map strongly regularizes the problem, allowing to improve a given set of initial maps. However, the resulting minimization presents difficult computational challenges which can be mitigated by generating sparse correspondences [HG13], or by aligning reduced functional spaces [HWG14]. Nevertheless, these methods require some post-processing to produce dense maps. In contrast, we match a pair of shapes while promoting cycle-consistency for point-to-point correspondences. The information stored in the map and its inverse allows us to obtain good results on benchmark data, without the extra knowledge a complete collection provides.

We pose our consistent mapping method as a Coupled Quadratic Assignment Problem with a cycle-consistency constraint, replacing the usual bijectivity requirement. There are two main challenges in converting this methodology into a practical solution. First, this high-dimensional problem is computationally prohibitive even for estimating a one-directional map, let alone solving for the map and its inverse simultaneously. Second, efficient minimization is difficult, as these problems are NP-hard in general. To tackle the first difficulty, we propose an iterative splitting scheme, where we solve for each mapping direction separately *while* exploiting information stored in the maps from the previous iteration. In addition, we limit each point to improve its mapped location, considering a small subset of new candidate points. Finally, we further relax the obtained optimization tasks to linear programming problems and employ a state-of-the-art algorithm based on tree-reweighted message passing (TRW) [Kol06] to arrive at an effective mapping framework.

In this paper, we suggest an effective methodology for producing approximately consistent point-to-point mappings between pairs of shapes, whose accuracy significantly exceeds current state-of-the-art results. Our method does not assume anything about the spatial embedding of the shapes, it is not restricted to isometric surfaces and it is easy to implement and solve using available minimization tools. In general, our approach provides an efficient solution when high-quality correspondences are required, at the cost of relaxing the exact bijectivity constraint, while producing good map consistency. We evaluate the proposed method on a few benchmark datasets, and demonstrate the advantages of our approach for computing shape correspondences using several quantitative and qualitative error measures.

2. Related Work

The problem of computing mapping between surfaces has been extensively studied, and thus we refer the reader to recent surveys for a thorough review of previous geometry-based [VKZHCO11] and data driven [XKH^{*}16] approaches. Here, we focus our attention on the main different methodologies for shape matching, where we classify existing work based on their required input, as well as their accuracy and cycle-consistency or bijectivity features.

Distortion minimization. A common approach for generating maps between surfaces is via minimizing the distortion they produce, which can be estimated using pointwise surface descriptors, such as the Heat Kernel Signature [SOG09], pairwise measures, such as geodesic distance differences [BBK06], or a combination of the two [DK10]. One well-known approach utilizing this idea is Blended Intrinsic Maps (BIM) [KLF11], where a collection of maps are blended together with a set of weights, optimized to produce minimal distortion smooth maps. The distortion measures can be further plugged into an integer quadratic programming problem with bijectivity constraints known as the Quadratic Assignment Problem (QAP) [KB57].

Unfortunately, directly solving QAP is feasible only for a very small number of variables [LdABN^{*}07]. To deal with the inherent computational challenges, various relaxations have been proposed in the literature, including spectral techniques [LH05, ADK16], continuous quadratic optimization [BBK06], linear programming [RBA^{*}12], and semidefinite programming [KKBL15]. Other works [VLR^{*}17, CK15] employ discrete optimization relaxations, allowing to match a sparse or moderately size set of feature points. However, to extend their correspondences to a complete map, some up-sampling procedure is required. Thus, several multi-resolution approaches [SY11, RDK12] were suggested to alleviate this issue, producing dense point-to-point maps. Our approach can be also interpreted as a variant of QAP, but, as opposed to other techniques, we completely avoid the potentially problematic multi-resolution methodology.

Embedding into an intermediate domain. An alternative line of work for computing correspondences embeds the input shapes into a joint intermediate domain, where the mapping can be computed more efficiently. Examples of this approach consider domains such as a high dimensional Euclidean space [EK03], the space spanned by a reduced eigendecomposition of the Laplacian [MHK^{*}08], the complex plane [LF09], or a hyperbolic orbifold [AL16]. In the latter work, the resulting mappings are smooth and guaranteed to be cycle-consistent across a collection, showing outstanding results in practice on challenging data. However, their method requires a small subset of manually selected feature landmarks, limiting its applicability and affecting its overall behavior.

Generalized maps. Instead of encoding mappings as point-to-point correspondences, [OBCS^{*}12] proposed to align surfaces via their functional spaces. The resulting Functional Maps are real-valued matrices from which a dense map can be extracted [OCB^{*}17]. In a related work, Shtern and Kimmel [SK14] offered an automatic iterative algorithm to significantly improve noisy or sparse input maps, by aligning the spectral kernels of

the shapes. Finally, [SPKS16] introduces a generalized notion of a mapping through minimization of the Gromov—Wasserstein distance, matching probability distributions across shapes.

Consistent correspondences. When a whole collection of shapes is given, [NBCW¹¹] observed that maintaining cycle-consistency within the collection is advantageous to the matching process. However, while a few methods offered dense outputs (e.g., [HZG¹²]), most existing techniques produce sparse matchings. Other methods focused on optimizing the functional map and its *inverse* while promoting cycle-consistency using functional map composition [ERGB16, RO18] or an adjoint operator [HO17]. However, in this setup, extracting a dense map is challenging as the maps are encoded as low dimensional matrices. In comparison, our method directly optimizes for a dense map and thus we avoid the error prone task of map recovery from a given functional map. Recently, [ESBC19] proposed an optimization framework which is similar in nature to ours, solving jointly for the correspondence and its inverse. Their approach is intrinsic to the shape as our method, but they utilize a different energy functional and obtain a global optimization problem, whereas ours is local. Finally, Windheuser et al. [WSSC11] also include the inverse map in their linear programming relaxation. However, they facilitate only unary terms based on the bending energy, whereas our method utilizes both unary and pairwise terms.

Learning approaches. Recently, several learning approaches have emerged, employing random forests [RRBW¹⁴], graph-based neural networks [LRR¹⁷] and fully connected ResNets [RO18]. These methods require a training set of matched shapes to perform learning, whereas our approach does not require any training data.

3. Overview and Background

Our goal is to estimate high quality correspondences between a given pair of shapes, \mathcal{M}_1 and \mathcal{M}_2 . We denote by ϕ_{21} a discrete point-to-point map, associating every point in \mathcal{M}_1 with a point in \mathcal{M}_2 , and similarly, $\phi_{12} : \mathcal{M}_2 \rightarrow \mathcal{M}_1$. As discrete shapes are often represented with thousands of points, applying exhaustive search approaches is impractical due to the factorial number of possible matchings. One way to deal with the huge amount of possibilities is through a definition of an optimization problem whose minimizer is as close as possible to the correct map we expect to obtain. A common choice in this context is the Quadratic Assignment Problem (QAP) [KB57], which is a special quadratic programming problem with the challenging bijectivity and binary variables constraints. In Section 4, we describe how to extend the standard QAP to the case of estimating coupled correspondences.

Similar to QAP, our formulation also involves binary constraints, and thus is hard to solve in practice. Moreover, while the QAP point-of-view regularizes the space of solutions, dimensionality remains a challenge, hindering the construction of an efficient algorithm. To cope with these difficulties, in Section 5 we present a novel method to solve our coupled mapping problem. Specifically, we suggest to split the global coupled optimization to two simpler sub-problems of computing one-sided mappings while exploiting information encoded in the inverse map. In Fig. 2, we compare the

results of our method, highlighting the benefits of using the data from the inverse map vs. discarding it. Our splitting scheme leads to an iterative process where at each step we solve two modified QAPs. In addition, we construct the matrix components of each QAP such that only a *linear* amount of its entries are effective in the optimization. Combined together, we arrive at a formulation that naturally lends itself to efficient and provably convergent linear relaxation techniques, which we solve using standard tools.

4. Coupled Quadratic Assignment Problem

In what follows, we formulate our coupled mapping approach as a generalization of the Quadratic Assignment Problem, considering the map and its inverse as separate variables while expressing the bijectivity constraint as a cycle-consistency requirement (Subsection 4.1). We then describe the particular formulations we used for the unary and pairwise terms of our Coupled QAP (CQAP) (Subsection 4.2).

4.1. Problem formulation

We assume to be given two shapes represented by triangle meshes, \mathcal{M}_1 and \mathcal{M}_2 , each given by a uniform sampling of n points from

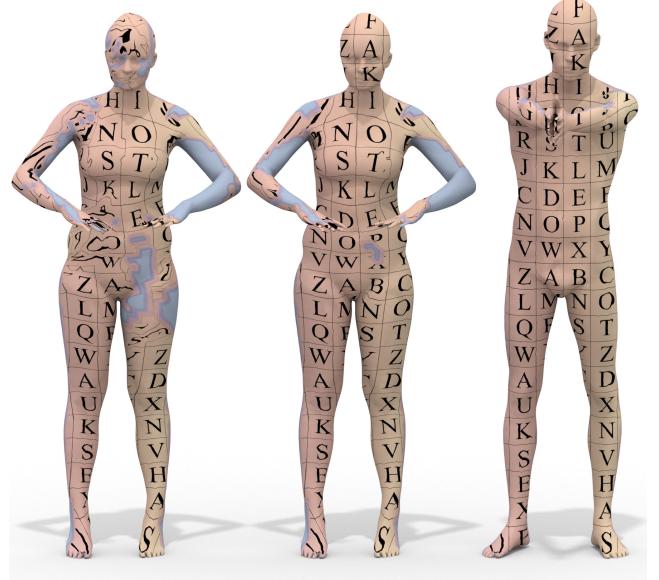


Figure 2: The inverse map is beneficial to the matching task. We map the woman (left, middle) to the man (right) using our method, without the inverse map (left) and with it (middle). The obtained correspondences allow us to transfer a texture from the target shape to the source shapes, where parts, invisible from the camera’s viewpoint, are omitted [APLJ4]. Indeed, without the information stored in the inverse map, the result is of overall poor quality, with many points (see the stomach and face) being incorrectly mapped. In contrast, our result (middle) shows a much better accuracy and succeeds in recovering most of the mismatched points.

the original shape. A standard approach for encoding a map ϕ_{ba} is using a binary matrix $X_{ba} \in \mathbb{N}^{n \times n}$, where $X_{ba}(k, i) = 1$ if vertex i in \mathcal{M}_a is mapped to vertex k in \mathcal{M}_b , and $X_{ba}(k, i) = 0$ otherwise. Standard QAP is formulated as a problem of finding a *bijective* map from e.g., \mathcal{M}_a to \mathcal{M}_b [LdABN*07]. Instead, we use an additional variable for the inverse directional map, and introduce a cycle-consistency constraint, leading to the following coupled quadratic assignment problem

$$\begin{aligned} \text{minimize}_{x_{21}, x_{12}} \quad & \left\{ \sum (x_{ba})^T P(x_{ba}) + \mu u^T x_{ba} \right\}, \\ \text{subject to} \quad & X_{ba} \cdot X_{ab} = I_n, \quad X_{ba} \in \{0, 1\}^{n \times n}, \end{aligned} \quad (1)$$

where the sum and the constraints above apply to $a = 1, b = 2$ and $a = 2, b = 1$. The binary vector $\{0, 1\}^{n^2} \ni x = \text{vec}(X)$ is a flattened version of the correspondence matrix X . The matrix $P \in \mathbb{R}^{n^2 \times n^2}$ is a *pairwise* term, accounting for the global distortion measures, and $u \in \mathbb{R}^{n^2}$ is a *unary* term, accounting for pointwise differences. The matrix I_n is the identity matrix of size n and μ is a scalar weight, scaling the effect of pointwise vs. pairwise penalties.

We emphasize that under the assumption that the shapes match, problem (1) should yield the same minimizer as traditional QAP. Namely, when the consistency constraints hold exactly, X_{21} and X_{12} are permutation matrices. Indeed, at first glance, our formulation (1) seems unnecessarily redundant and more computationally involved when compared to standard QAP. However, we opt for the above setup as it fits better to the relaxations and approximations we suggest next.

4.2. Pointwise and pairwise dissimilarity measures

We define the pointwise and pairwise dissimilarity measures using *intrinsic* surface descriptors. Given a point $i \in \mathcal{M}_a$, we define

$$F_a(i) = \left(w_1 f_a^1(i), \dots, w_{k_d} f_a^{k_d}(i) \right),$$

where $F_a(i)$ is a long vector of k_d stacked descriptor vectors $f_a^1(i), \dots, f_a^{k_d}(i)$, along with their respective weights w_1, \dots, w_{k_d} . For instance, to generate a particular mapping we may use a weighted combination of the Heat Kernel Signature (HKS) [SOG09] and Wave Kernel Signature (WKS) [ASC11], see Fig. 3 for an illustration. Using the above definition, the pointwise distortion measure is defined as

$$U(i, k) = |F_1(i) - F_2(k)|_2^2, \quad (2)$$

where $i \in \mathcal{M}_1$ and $k \in \mathcal{M}_2$. Finally, we define $u = \text{vec}(U)$ to be the flattened version of U .

The pairwise distortion quantifies the discrepancy introduced by the mapping to a local area around each point. Specifically, we use geodesic distances for that purpose. For each pair of points $i, j \in \mathcal{M}_1$ that are possibly matched to vertices $k, l \in \mathcal{M}_2$, respectively, we define the following distortion measure

$$P_{ijkl} = |d_1(i, j) - d_2(k, l)|, \quad (3)$$

where d_1 and d_2 are the geodesic distances computed for shapes \mathcal{M}_1 and \mathcal{M}_2 . We note that in general other measures could be employed for the pointwise and pairwise terms. However, since a significant portion of our work concentrates on the case of (nearly)

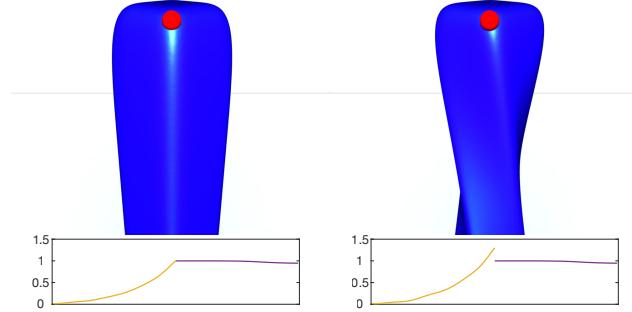


Figure 3: For each point on our shapes, we form a vector of intrinsic data. The example above shows the result for the marked red points on a bar and a twisted bar shapes (left and right, respectively). The plotted graphs correspond to WKS (yellow) and HKS (purple), with weights $1/45$ and 1 .

isometric shapes, we adopt the above formulations. Notice that our measures are symmetric with respect to the source and target shapes, leading to symmetric matrices, U and P . Thus, we can use the same u and P terms in Eq. (1), for both map directions.

The time complexity and memory requirements of our CQAP (1) imposes, make it infeasible in most cases, except for extremely low values of n [LdABN*07]. While various scalable relaxations for traditional QAP exist (see e.g., [DML17]), dimensionality remains a difficult challenge, commonly addressed via up-sampling techniques (see e.g., [VLR*17]). We proceed by describing the modifications we employ on the original problem, allowing us to directly work with high-resolution shapes.

5. Practical Coupled Matching

To mitigate the challenges discussed in Section 4, we propose to split the above CQAP (1) into two sub-problems, iteratively optimizing separately for x_{21} and x_{12} , while exploiting the mapping information available from the previous step. Inspired by [HZG*12], we devise a stopping condition based on the current maps quality. Specifically, given maps x_{21}^t and x_{12}^t from step t , we perform the following steps

1. solve for x_{21}^{t+1} , using x_{21}^t and $\tilde{x}_{21}^t := (x_{12}^t)^{-1}$,
2. solve for x_{12}^{t+1} , using x_{12}^t and $\tilde{x}_{12}^t := (x_{21}^{t+1})^{-1}$,

as long as the consistency constraint is improving, namely, as long as

$$\sum |X_{ba}^{t+1} \cdot X_{ab}^{t+1} - I_n|_F^2 \leq \sum |X_{ba}^t \cdot X_{ab}^t - I_n|_F^2, \quad (4)$$

where $|\cdot|_F$ is the Frobenius norm. Unfortunately, since we deal with approximate bijections, \tilde{x}_{21}^t and \tilde{x}_{12}^t are not well-defined. In particular, $\tilde{x}_{ba}^t(i)$ may be empty or include one or several points. In Section 6, we describe how \tilde{x}_{ba}^t can be uniquely defined, and what type of initial correspondences $\tilde{x}_{21}^0, \tilde{x}_{12}^0$, we use.

Estimating x_{21}^{t+1} and x_{12}^{t+1} is achieved via a relaxed version of the standard QAP, where we omit the bijectivity requirement. Formally,

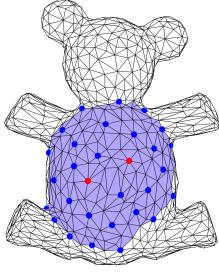


Figure 4: An illustration of the proposed candidate selection on a toy example. Red circles denote the locations of $x^t(i), x^d(i)$, the colored area denotes the unified 5-ring neighborhoods of the two vertices, and blue circles denote 30 vertices sampled using the farthest point sampling.

we consider the problem

$$\underset{x}{\text{minimize}} \quad x^T P x + \mu u^T x \quad \text{s.t.} \quad \mathbf{1}^T X = \mathbf{1}^T, X \in \{0, 1\}^{n \times n}, \quad (5)$$

where $\mathbf{1}$ is an all-ones vector of size n , and thus the constraint $\mathbf{1}^T X = \mathbf{1}^T$ means that every point on the source shape is mapped to a single (not necessarily unique) point on the target shape. From now on, we denote the above problem as RQAP (relaxed QAP).

To motivate our choice of relaxation, we stress that while the bijectivity condition is omitted in Eq. (5), the maps x_{21} and x_{12} are still coupled in our iterative procedure via their approximated inverses. In particular, the candidate nodes for e.g., x_{21} are taken so that a node which leads to a bijective map appears in the candidate list. In practice this heuristic significantly simplifies problem (1), while achieving good consistency results, see Fig. 11. Finally, we mention that the bijective constraints $X_{ba} \cdot X_{ab} = I_n$ are *non-convex* and thus require special handling. For instance, a similar setup to ours was recently proposed in the non-linear dynamics community [AYB19]. To deal with their difficult constraint, the authors propose to split the optimization using the Alternating Direction Method of Multipliers (ADMM). Computing a map in their minimization requires an estimate of the inverse map at each iteration, as is proposed in our approach. This example empirically supports our splitting methodology, and thus it provides a strong indication that our heuristic is effective.

In addition to our relaxation, one may consider continuous *linear* relaxations to problem (5), allowing to optimize for up to $n = 15 \times 10^3$ points [VLB*17]. However, high resolution shapes still require using some multiscale techniques. Instead, we aim at sparsifying u and P , thus avoiding the up-sampling step altogether. To this end, we consider two simplifying assumptions. First, instead of considering all combinations of point pairs (i, j) in the pairwise term (3), we focus on 1-ring relations. Namely, we limit the pairwise contributions to pairs of vertices sharing an edge in the triangulation of the source shape. Formally, it means that P_{ijkl} is ∞ for every combination (i, j) which is *not* an edge in the source shape. Second, we limit the candidate matches of a point $i \in \mathcal{M}_1$ to m uniformly sampled points from the r -radius neighborhoods of the pair $(x^t(i), x^d(i)) \in \mathcal{M}_2, \forall i$. That is, for each point on the source shape, the maps x and \tilde{x} provide a pair of candidate points

on the target shape. We consider the set of points that are within geodesic distance r from this pair, and sample m vertices from the entire set using farthest point sampling [HS85]. We illustrate our uniform sampling procedure in Fig. 4. Reducing the amount of possible candidate points to $m \ll n$ leads to m non-infinite entries in U , per point, and m^2 non-infinite entries in the matrix P , per pair of points, allowing, in practice, to operate on shapes with higher resolutions. For instance, we show in Fig. 5 our result on a pair of shapes with $\approx 25k$ vertices, where no post-processing up-sampling was required.

Overall, the above relaxations and approximations result in a *sparse* RQAP to be solved separately for each map, x_{21} and x_{12} . Still, as our matrices P are *not* positive semi-definite, a further relaxation of the binary constraints to real variables will leave us with non-convex quadratic programming problems, which are known to be NP-hard in general [PV91]. Alternatively, we show in Section 6 how our relaxed problem can be fitted into a Markov Random Field (MRF) formulation, allowing to use non-necessarily convex energy functionals, and can be subsequently solved via provably convergent tree-reweighted message passing method [Kol06].

6. Implementation Details

Sparse RQAP as a discrete MRF. At each iteration, we solve two sparse relaxed quadratic assignment problems, given by Eq. (5), where only a linear number of elements in P and U affect the optimization. For each point $i \in \mathcal{M}_a$, we denote by $\Phi(i)$ the set of m vertex indices from \mathcal{M}_b , obtained using the farthest point sampling described in Section 5. We denote by \mathcal{P} and \mathcal{U} the pairwise and unary terms consisting of only non-infinite entries from P and U , corresponding to the set $\Phi(i)$ for every i . In the language of discrete Markov Random Fields, the set $\Phi(i)$ represents the *labels* each vertex i can be associated with, m labels per vertex. Now, the objective function in problem (5) can be written as

$$E(\phi | \theta) = \sum_i \theta_i(k) + \sum_{(i,j)} \theta_{ij}(k, l), \quad (6)$$

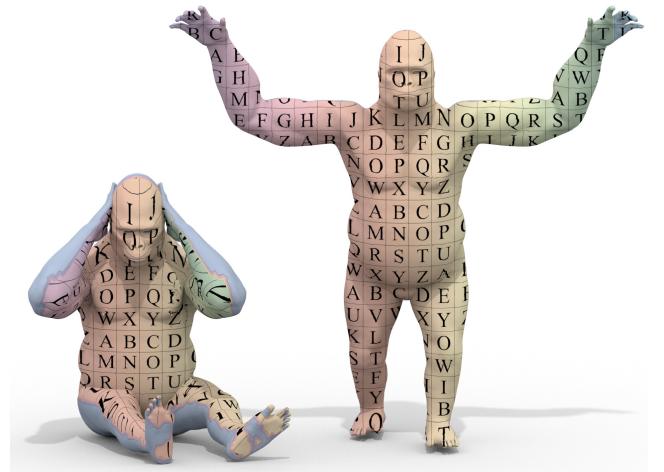


Figure 5: Texture mapping on TOSCA shapes with $\approx 25k$ vertices.

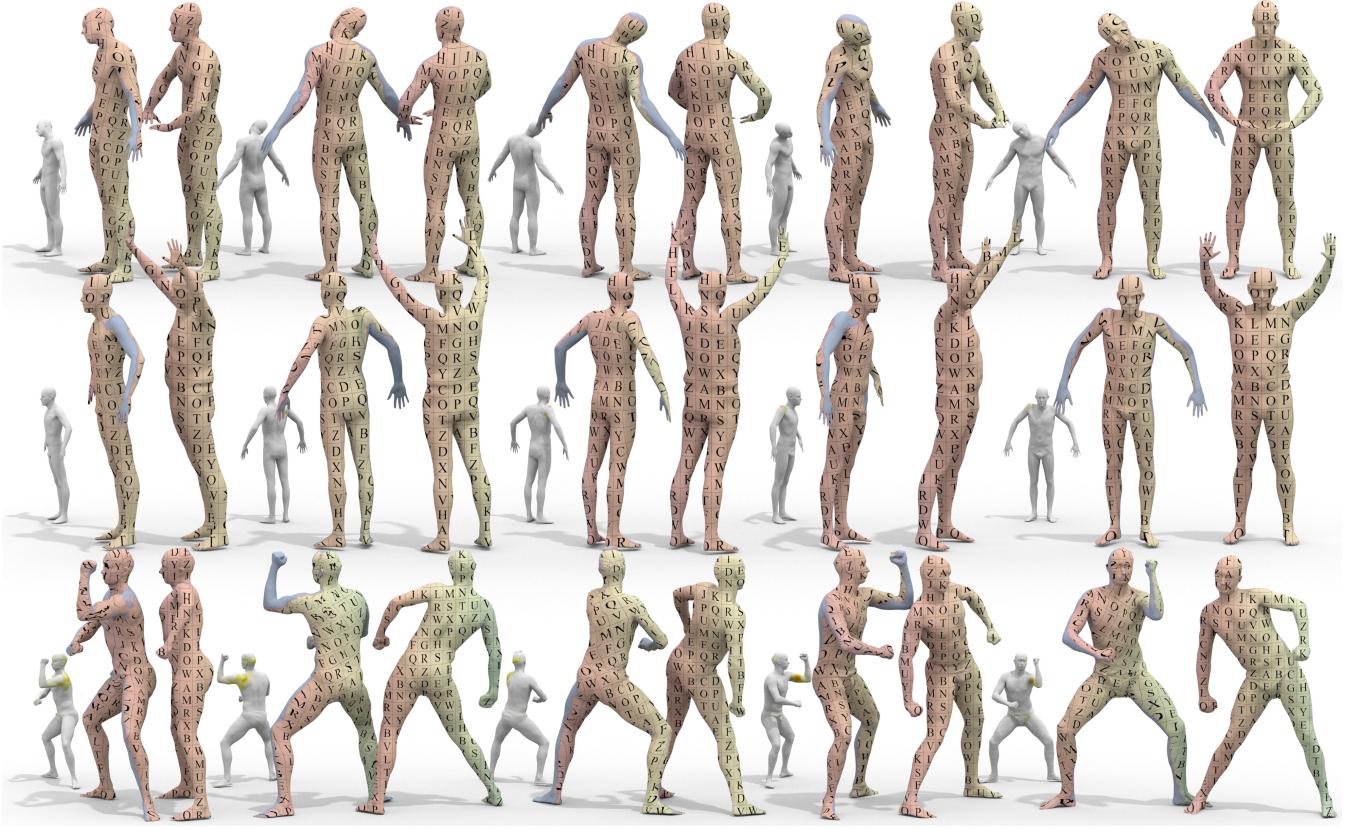


Figure 6: Correspondence quality illustration using texture transfer. In each row, we show a pair of shapes rendered from five different viewpoints, for which we computed a correspondence using our method. The maps are used to pull a texture from the target shape (right) to the source shape (left) for each of the viewpoints. In addition, we show the geodesic error our maps exhibit on the smaller source shapes. There, white color corresponds to zero error, and red to the maximal error. We note that for the above examples, more than 90% of the points were mapped to their ground truth locations.

where $\theta_i(k) = \mathcal{U}_{ik}$ and $\theta_{ij}(k, l) = \mathcal{P}_{ijkl}$ and $k \in \Phi(i), l \in \Phi(j)$. A minimum of E over the parameter set θ corresponds to a *maximum a-posteriori* (MAP) labeling ϕ [GG84] in a discrete MRF with unary and pairwise potentials. This energy is typically considered for matching problems arising in computer vision, such as image segmentation and stereo matching [SZS*08]. Finally, problem (6) can be solved directly using established discrete optimization tools. Specifically, we employ the TRW-S [Kol06] algorithm for this task.

Initial maps. Our method takes as input a pair of triangle meshes \mathcal{M}_1 and \mathcal{M}_2 along with two imperfect and, possibly, noisy initial point-to-point maps $\phi_{21}^0 : \mathcal{M}_1 \rightarrow \mathcal{M}_2$ and $\phi_{12}^0 : \mathcal{M}_2 \rightarrow \mathcal{M}_1$. As we opted for an as-automatic-as-possible method, we used initial maps computed with automatic techniques. Specifically, all the results we show in the paper were obtained by initializing the proposed method using either Blended Intrinsic Maps (BIM) [KLF11] or Iterative Closest Spectral Kernel Maps (ICSKM) [SK14]. We note that while we focused on input from these methods, any other matching technique could be considered instead. Given ϕ_{21}^0, ϕ_{12}^0 , and the set of parameters μ and w_1, \dots, w_{kd} , our algorithm is *fully automatic* and does not require any user interaction or matched fea-

ture points. We show a few results of our method in Fig. 6 given initial ICSKM maps, where we evaluate the obtained maps via texture transfer from the target shapes to the source shapes.

Inverse map computation. The candidate set selection, described in Section 5, requires having the map $\tilde{\phi}_{ba}$, where $\tilde{\phi}_{ba}$ denotes the inverse map of ϕ_{ab} . Naturally, inferring $\tilde{\phi}_{ba}$ for a non-bijective discrete map ϕ_{ab} is challenging, as $\tilde{\phi}_{ba}(i)$ could be undefined or include several points. To this end, we propose a simple approach to compute $\tilde{\phi}_{ba}$, which produces a well-defined map that contains a single match for every vertex i in \mathcal{M}_a . We denote by $\Psi(i)$ the set of vertices $k \in \mathcal{M}_b$ for which $i = \phi_{ab}(k)$. Then, the inverse map is defined as follows

$$\tilde{\phi}_{ba}(i) = \begin{cases} \bar{\Psi}(i), & \text{if } \Psi(i) \neq \emptyset \\ \phi_{ba}(i), & \text{otherwise} \end{cases}$$

where $\bar{\Psi}(i)$ is the vertex in \mathcal{M}_b , nearest, in terms of Euclidean distance, to the average of the set $\Psi(i)$. Namely, if we have several candidate points for the inverse, we simply take a vertex in \mathcal{M}_b which is closest to their average, and if we have no candidates, we use the opposite-directional map ϕ_{ba} . Finally, we emphasize that,

```

function COMAP( $x_{21}^0, x_{12}^0$ )
   $x_{21}^t \leftarrow x_{21}^0, x_{12}^t \leftarrow x_{12}^0$                                  $\triangleright$  Initialization
   $c^t = \sum |X_{ab}^t \cdot X_{ba}^t - I_n|^2_F$ 
  for  $j = 1, 2, \dots$  do
     $\tilde{x}_{21}^t \leftarrow x \left[ (\phi_{12}^t)^{-1} \right]$ 
     $x_{21}^t \leftarrow \text{MAP}(\tilde{x}_{21}^t, \tilde{x}_{12}^t, r_j)$                                  $\triangleright$  Infer new map  $x_{21}$ 
     $\tilde{x}_{12}^t \leftarrow x \left[ (\phi_{21}^t)^{-1} \right]$ 
     $x_{12}^t \leftarrow \text{MAP}(\tilde{x}_{12}^t, \tilde{x}_{21}^t, r_j)$                                  $\triangleright$  Infer new map  $x_{12}$ 
     $c \leftarrow \sum |X_{ab}^t \cdot X_{ba}^t - I_n|^2_F$ 
    if  $c > c^t$  then                                 $\triangleright$  Evaluate Eq. (4)
      break
     $c^t \leftarrow c$ 
  return  $x_{21}^t, x_{12}^t$ 

function MAP( $x^t, \tilde{x}^t, r$ )       $\triangleright$  Compute map  $x^{t+1} : \mathcal{M}_a \rightarrow \mathcal{M}_b$ 
  for all  $i = 1, \dots, n$  do
     $\Phi(i) \leftarrow \text{FPS}(x^t(i), \tilde{x}^t(i), r)$ 
     $\mathcal{U}_{ik} \leftarrow |F_a(i) - F_b(k)|^2$        $\triangleright$  Pointwise distortions, Eq. (2)
     $\mathcal{P}_{ijkl} \leftarrow |d_a(i, j) - d_b(k, l)|$      $\triangleright$  Pairwise distortions, Eq. (3)
     $x^{t+1} \leftarrow \text{TRW-S}(\mathcal{U}, \mathcal{P})$            $\triangleright$  Solve Eq. (5)
  return  $x^{t+1}$ 

```

Algorithm 1: An iterative splitting scheme for computing the map between the source and target shape as well as its inverse. The function FPS performs farthest point sampling of m points of the r_j -radius neighborhoods of its inputs. The function TRW-S is an implementation of the tree-reweighted message passing algorithm [Kol06]. See more details in Section 6.

while we took the above simple approach, other, more sophisticated tools for computing the inverse map could be employed, potentially providing an improved label set $\Phi(i)$.

Shape descriptors and geodesic distances. To construct the unary term \mathcal{U} , we employed three different intrinsic pointwise descriptors, such as the Heat Kernel Signature (HKS) [SOG09], Wave Kernel Signature (WKS) [ASC11], and Signature of Histogram of Orientations (SHOT) [TSDS10]. In Section 7, we provide the actual descriptor weights w_1, \dots, w_3 we used for each dataset. Geodesic distances were computed using the fast marching method [KS98]. For meshes with up 15×10^3 vertices, the entire distance matrices were precomputed and stored in memory. For larger meshes, we used the distance approximation method suggested in [SAZK15] using a truncated geodesic distance basis of size 50 with 100 samples. This approximation tends to produce larger relative errors for smaller distances, therefore, in addition to the truncated geodesic distance basis, we pre-computed and stored distances from all vertices to their 10-ring neighbors. These precise distances were used for the farthest point sampling at the candidate selection step.

Technical details. We implemented our algorithm in MATLAB; its pseudo-code is provided in Algorithm 1. The optimization was per-

formed using an implementation of the TRWS algorithm [Kol06]. Computation was performed on a PC equipped with a 3.5 GHz 6-Core Intel Xeon E5 CPU, and 64 GB RAM. Table 1 shows typical run times of our method for shape pairs of different resolutions. We report three computation times per pair of shapes, aggregated over three optimization iterations of our algorithm, which was applied using the setup described in Section 7.1. In the pre-processing step, we load geodesic distance matrices and compute three different shape descriptors (HKS, WKS and SHOT). The data preparation step includes computation of the candidate points, and forming the unary and the pairwise terms for the MRF. Lastly, we list the total computation time of the six TRWS calls. We note that for larger shapes, the farthest point sampling accounts for, roughly, 65% of the time spent during data preparation when choosing the candidate points from the union of 6-ring neighborhoods, and for 62% when using 4-ring neighborhoods. It is possible that other, less-time consuming, techniques may be used to produce candidate sets for the MRF computation. We leave further consideration of these ideas for future work.

7. Evaluation and Results

We compared our method on benchmark datasets, with known point-to-point mappings, including SCAPE [ASP*04], TOSCA [BBK08] and FAUST [BRLB14]. The data contain various human and animal shapes, allowing to experiment with isometric shapes within the same class, or consider the quasi-isometric case, taking shapes of different categories. We utilize several evaluation protocols to highlight the precision and consistency of our maps, such as the cumulative [KLF11] and average [CK15] geodesic errors, and information transfer between the shapes, including texture or color functions. We compare our method to state-of-the-art approaches such as Blended Intrinsic Maps (BIM) [KLF11], Functional Maps (FMAPS) [OBCS*12], Iterative Closest Spectral Kernel Maps (ICSKM) [SK14], Spectral Generalized Multi-Dimensional Scaling (SGMDS) [ADK16], Hyperbolic Orbifold Tutte embeddings (HOT) [AL16], Kernel Matching (KM) [VLB*17], and the work of Chen and Koltun (CGF) [CK15].

7.1. Benchmark results

We evaluate our method on pairs of shapes from the above datasets, and present the obtained cumulative geodesic error measures for SCAPE (Fig. 9), TOSCA (Fig. 10) and FAUST (Fig. 7). The evaluation is performed using the procedure described in [KLF11]. Namely, we calculate the distances between the mapped points and the ground truth correspondences, for each pair of shapes, and show

| # V | pre-processing (sec) | data preparation | TRW-S |
|-------|----------------------|------------------|-------|
| 4344 | 3.82 | 0.45 | 7.54 |
| 6890 | 7.36 | 1.3 | 8.52 |
| 12500 | 20.22 | 3.68 | 15.80 |
| 19248 | 18.27 | 8.55 | 26.45 |
| 27894 | 23.50 | 16.20 | 38.39 |

Table 1: Timings (in minutes, unless stated otherwise) of our algorithm tested on shapes with varying resolutions.

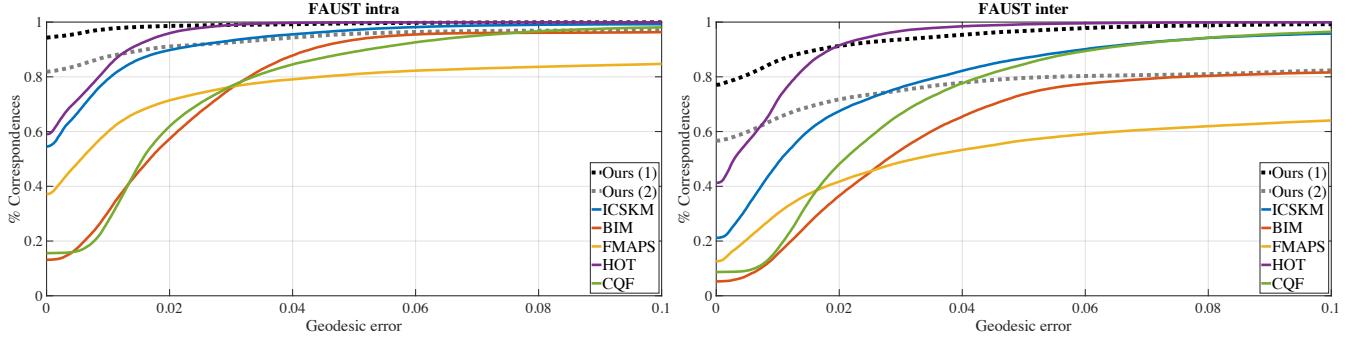


Figure 7: Cumulative geodesic error measures for FAUST, separated to same subject different pose class (intra, left) and different subjects (inter, right). See the text in Subsection 7.1 for more details.

the average aggregated percentage of points that are within a certain geodesic distance. For SCAPE and TOSCA we took the pairs used in BIM [KLF11], and for FAUST, we distinguish between the cases of same subject in different poses (intra) and different subjects (inter), using the same pairs as in [CK15]. We use “Ours (1)” to denote the results of our method initialized with maps produced by ICSKM, and “Ours (2)” the results when initializing with maps produced by BIM. We stress that compared to other state-of-the-art methods, BIM achieves relatively poor maps as we show in the cumulative geodesic error measures in Figs. 7,9,10. Using BIM as a baseline for refinement is thus a challenging setup, and it highlights the benefits of our matching procedure. To compare against methods whose output is a precise map, we project their results to the vertices of the target shape, yielding point-to-point maps.

The parameters we used for all pairs in a certain dataset were $\mu = 2$, $w = [0, 10, 1]$ (SCAPE), $\mu = 20$, $w = [20, 0, 1]$ (FAUST), and $\mu = .2$, $w = [0, 0, 1]$ (TOSCA), where the weights correspond to HKS, WKS, and SHOT, respectively. In our experience, the SHOT descriptors improved the results significantly and thus we them across all our experiments. The weights vector w is chosen so that HKS/WKS will share the same scale as SHOT. The parameter μ balances between our unary and pairwise terms. In cases where the geodesic distances are not expected to lead to good maps we take μ to be large as in the quasi-isometric scenario of FAUST inter. Finally, to obtain a consistent behavior of our algorithm across all shape pairs, we fix the Algorithm 1 to run 3 times with ranges that correspond to the 6-, 4- and 3-ring neighborhoods. This heuristic choice is reminiscent of a hierarchical approach where at each iteration the search space of vertices shrinks, as the method converges to a solution.

Overall, our approach matches more points to the ground truth compared to existing techniques with a large margin of $\approx 18\%$ for SCAPE, $\approx 28\%$ for TOSCA, $\approx 35\%$ for FAUST intra, and $\approx 36\%$ FAUST inter, measured from our best results, “Ours (1)”, to the best available results of other methods. We note that even when starting from relatively noisy maps provided by BIM, we manage to output high quality correspondences in “Ours (2)”, surpassing many of current state-of-the-art methods. Moreover, we indicate that some methods such as HOT perform better for longer-range distances. Finally, we show in Fig. 13 the average geodesic er-

ror obtained for each of the pairs in the benchmarks, where each method is sorted independently to the others, except for our results which are sorted with respect to ICSKM or BIM. We observe that our results significantly improve the average error of our starting point and is comparable to results of state-of-the-art techniques.

7.2. Consistency results

We evaluate the cycle-consistency of the maps obtained with our method. To this end, we generate the map and its inverse for a pair of shapes and compute their composition, yielding a map from the source shape to itself. We emphasize that for dense correspondences, such an experiment is extremely challenging as points can “drift” over cycles, thus exhibiting further error. We demonstrate our consistency results in Fig. 1 and Fig. 8, where we map a texture from the source shape to itself using the composed map. The source and target shapes correspond to the left and right smaller objects behind, respectively. In addition, we test these maps in the context of the benchmarks mentioned in Subsection 7.1 and we show the results in Fig. 11. Our results outperform all of the methods, except for approaches which produce bijective maps such as KM and HOT. We emphasize that HOT is fully cycle-consistent and thus the graph for their precise correspondences would just be 1, however, it is lower due to the projection we perform to obtain dense (i.e., point-to-point) mappings. Interestingly, map accuracy does not guarantee cycle-consistency – cf. FMAPS graphs for SCAPE in Figs. 9 and 11, where the obtained maps are relatively accurate but the average cycle-consistency constraint is of poor quality. The opposite also occurs – cf. BIM graphs for SCAPE in Figs. 9 and 11, where the average consistency is much higher than the map accuracy. These results further strengthen the observation in [NBCW*11] and other works regarding the regularizing effect of the cycle-consistency requirement.

7.3. Comparison with KM and HOT

We compare our results to KM and HOT in the challenging case of quasi-isometric shapes, considering *different* subjects from the FAUST dataset. Our method produces point-to-point (dense) maps, which are not guaranteed to be smooth, whereas KM outputs dense bijective correspondences and HOT generates smooth and *precise*

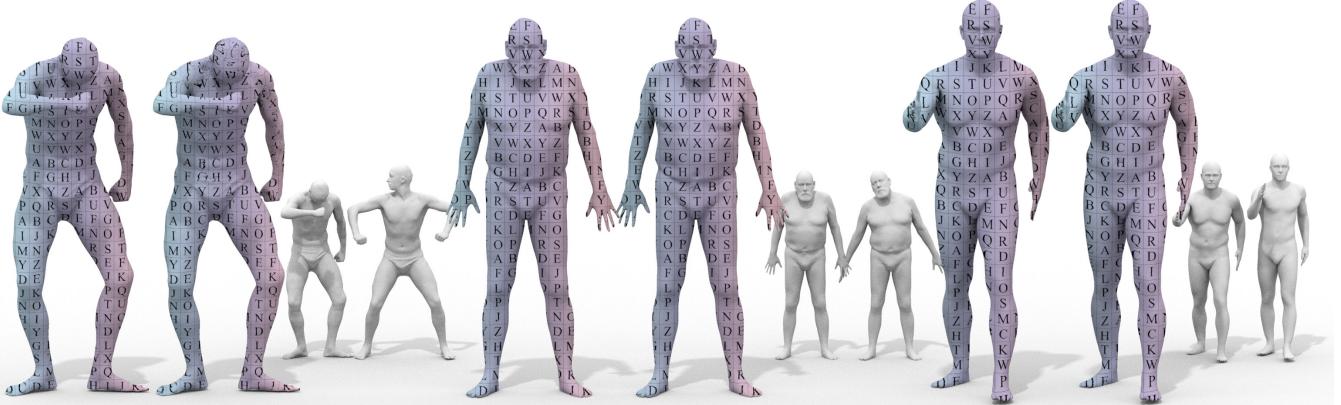


Figure 8: Texturing shapes over cycles is particularly challenging for point-to-point correspondences since errors are accumulated. Nevertheless, our method produces remarkable results on the examples above, revealing an insignificant amount of error.

bijections. Namely, in HOT, each point on the source shape is mapped to a point in one of the triangles of the target surface. As the triangle meshes in this example share the same connectivity, we projected HOT precise maps to the closest vertex, damaging their overall smooth nature, but allowing for a better precision evaluation. We show in Fig. 12 the results of mapping the human shapes, which appear to the left, to the right shape, and we transfer a texture from the target shape (right) using the ground-truth map (left), our result (middle-left), KM’s map (middle) and HOT’s projected matching (middle-right). Our output is the closest to the ground-truth in this example displaying minor incorrect matches, whereas the other methods display more error around the stomach and legs. We emphasize that while our method uses the geodesic distances in the pairwise term and thus assumes an underlying isometry, it nevertheless attains good results in the quasi-isometric case as we show for the FAUST dataset. However, if the shapes are non-isometric as in e.g., SHREC07 [GBP07], different unary and pairwise components may be required, see e.g., [EHA¹⁹]. We leave this direction for further investigation and future work.

7.4. Coupled vs. non-coupled matching

To quantify the effect of utilizing the inverse map in the matching process, we apply our method twice on the benchmarks: first, without using information from the inverse correspondences, and second, with this information. We show in Fig. 14 the graphs obtained from this experiment, where we demonstrate a clear advantage to using inverse map information. Specifically, the average cumulative geodesic error has improved by $\approx 1\%$ (SCAPE), $\approx 4\%$ (TOSCA), $\approx 3\%$ (FAUST intra), and $\approx 10\%$ (FAUST inter).

7.5. Matching with mesh topology changes

We additionally compared the performance of our method and other approaches on a more challenging scenario. All FAUST models share the same connectivity, which may affect the mapping results of different methods. In the following experiment, we *independently* re-meshed all the models in the FAUST dataset, to have around 5000 vertices, and re-computed the correspondences. The re-meshing was done in MeshLab using [GH97]. In Fig. 15 we evaluate the geodesic errors obtained by ICSKM, BIM, FMAPS, HOT, and our technique, initialized with ICSKM and BIM. We observe that in this case too, the proposed method significantly im-

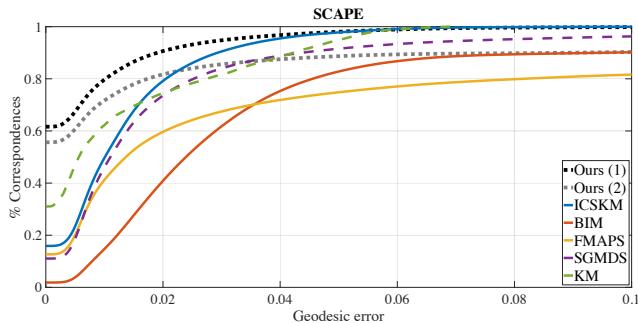


Figure 9: Cumulative geodesic error measures for SCAPE.

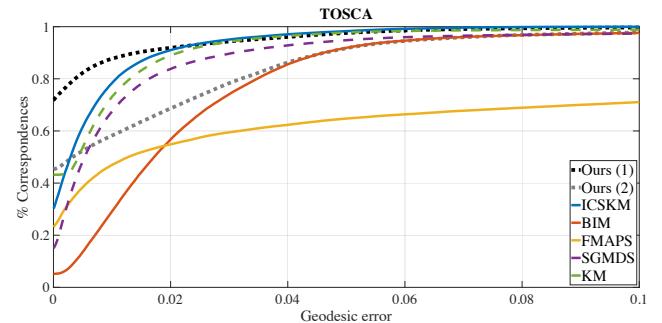


Figure 10: Cumulative geodesic error measures for TOSCA.

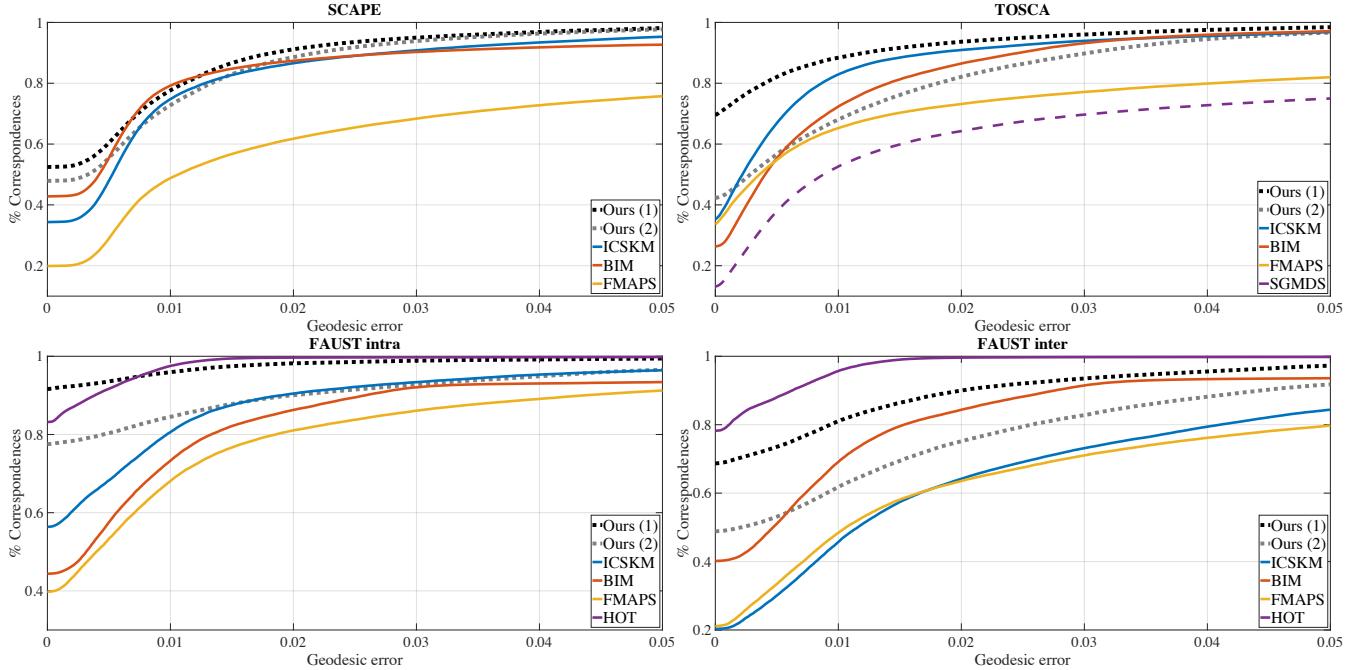


Figure 11: Cumulative geodesic error measures for the cycle-composed maps evaluated for each of the above methods on the benchmarks.

proves both BIM and ICSKM results, and outperforms the rest of the methods, especially for small range geodesic distance errors (< 0.005).

8. Limitations

One limitation of our method is that we assume to be given a pair of initial dense correspondences. While there exist automatic mapping tools for a large range of shapes (such as BIM and ICSKM), a stand-alone machinery is often preferred. In this context we note that our method works best when the correct matches are in close

proximity to the initial correspondences. A natural approach would then be to initially match between sub-sampled versions of the given shapes, followed by a drift step, allowing points to finalize their destination. Indeed, the work of [CK15] and [VLB*17] follow this general idea, however, we observe that mismatches during the initial step severely affect the overall behavior of a matching algorithm. Another disadvantage of our method involves its discrete nature as we output dense, point-to-point maps, where in many situations, precise, point-to-triangle correspondences allow greater flexibility. We leave further investigation of these shortcomings to future work.

9. Conclusions and Future Work

In this paper, we presented a new method for computing dense and approximately consistent maps between a pair of triangle meshes. Our method is fully intrinsic, it does not require input feature points, and it is applicable in the non-isometric case. Key to the robustness of our method is a novel iterative scheme which exploits information stored in the inverse map, while considering a linear number of candidate points per optimization iteration. Thus, we can work directly with full resolution shapes and avoid the potentially error-prone process of hierarchical matching. Our algorithm is simple to implement, can employ existing discrete optimization tools, and its results are more accurate and consistent when compared to state-of-the-art methods.

In the future, we would like to extend our method to include smoothness constraints, possibly generalizing our approach to yield point-to-triangle correspondences. In addition, we are interested in incorporating bijectivity as a hard constraint, guaranteeing exactly consistent maps. Finally, we believe that our approach could be also

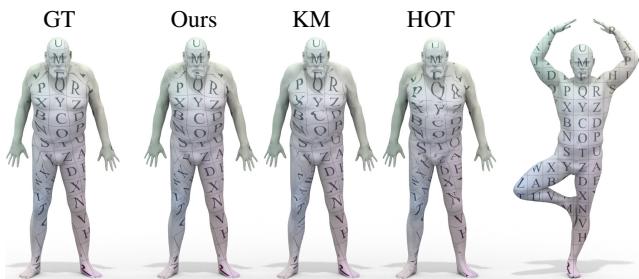


Figure 12: In the challenging setting of mapping quasi-isometric shapes, our method yields a result (middle-left) which is more accurate when compared to state-of-the-art techniques such as KM [VLB*17] (middle) and HOT [AL16] (middle-right) evaluated against the ground truth result (left). The target shape, from which the texture was pulled, is shown on the right.

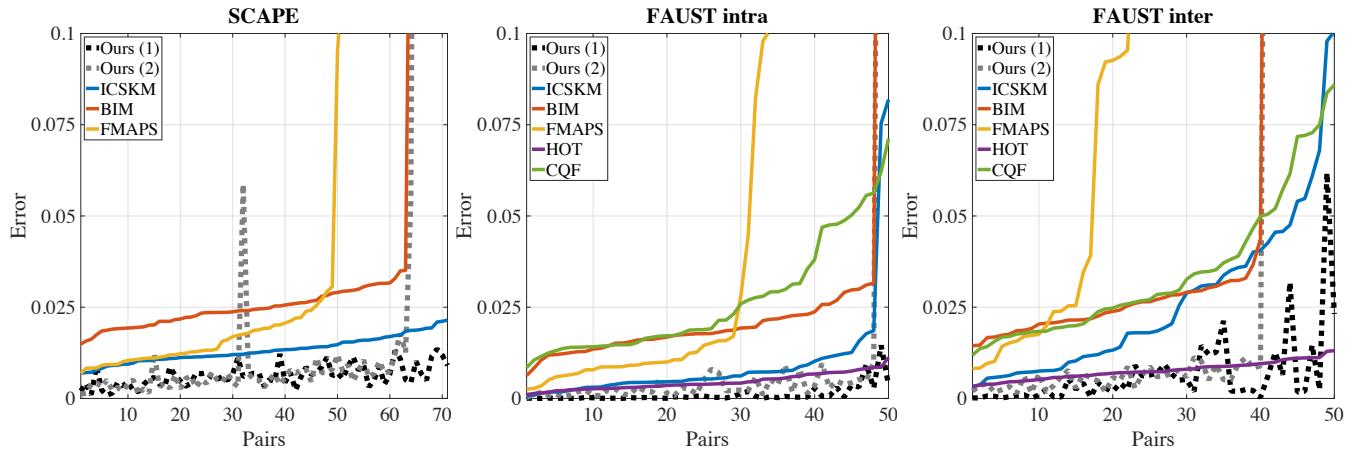


Figure 13: The average geodesic error is presented for each of the tested pairs in the benchmarks, where the results are sorted independently for each of the other methods, i.e., pair 40 for BIM and FMAPS may be different.

considered in the case of symmetric shapes, to produce maps which respect the underlying symmetry.

Acknowledgments. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 793800, a Zuckerman STEM Leadership Postdoctoral Fellowship, The Eric and Wendy Schmidt Postdoctoral Grant for Women in Mathematical and Computing Sciences, NSF grants IIS-1528025 and DMS-1546206, a Vannevar Bush Faculty Fellowship, and a gift from the Autodesk Corporation.

References

[ADK16] AFLAO Y., DUBROVINA A., KIMMEL R.: Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision (IJCV)* 118, 3 (2016), 380–392. 2, 7

[AL16] AIGERMAN N., LIPMAN Y.: Hyperbolic orbifold tutte embeddings. *ACM Trans. Graph.* 35, 6 (2016), 217–1. 2, 7, 10

[BBLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2014), pp. 3794–3801. 7

[CK15] CHEN Q., KOLTUN V.: Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision* (2015), pp. 2039–2047. 2, 7, 8, 10

[DK10] DUBROVINA A., KIMMEL R.: Matching shapes by eigendecomposition of the laplace-beltrami operator. In *Proc. 3DPVT* (2010), vol. 2. 2

[DML17] DYM N., MARON H., LIPMAN Y.: Ds++: A flexible, scalable and provably tight relaxation for matching problems. *arXiv preprint arXiv:1705.06148* (2017). 4

[EHA*19] EZUZ D., HEEREN B., AZENCOT O., RUMPF M., BEN-CHEN M.: Elastic correspondence between triangle meshes. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 121–134. 9

[EK03] ELAD A., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 10 (Oct. 2003), 1285–1295. 2

[ERGB16] EYNARD D., RODOLA E., GLASHOFF K., BRONSTEIN [APL14] AIGERMAN N., PORANNE R., LIPMAN Y.: Lifted bijections for low distortion surface mappings. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 69. 3

[ASC11] AUBRY M., SCHLICKERWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on* (2011), IEEE, pp. 1626–1633. 4, 7

[ASP*04] ANGUELOV D., SRINIVASAN P., PANG H.-C., KOLLER D., THRUN S., DAVIS J.: The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. *Advances in neural information processing systems 17* (2004), 33–40. 7

[AYB19] AZENCOT O., YIN W., BERTOZZI A.: Consistent dynamic mode decomposition. *arXiv preprint arXiv:1905.09736* (2019). 5

[BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences 103*, 5 (2006), 1168–1172. 2

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

[CK15] CHEN Q., KOLTUN V.: Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision* (2015), pp. 2039–2047. 2, 7, 8, 10

[DK10] DUBROVINA A., KIMMEL R.: Matching shapes by eigendecomposition of the laplace-beltrami operator. In *Proc. 3DPVT* (2010), vol. 2. 2

[DML17] DYM N., MARON H., LIPMAN Y.: Ds++: A flexible, scalable and provably tight relaxation for matching problems. *arXiv preprint arXiv:1705.06148* (2017). 4

[EHA*19] EZUZ D., HEEREN B., AZENCOT O., RUMPF M., BEN-CHEN M.: Elastic correspondence between triangle meshes. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 121–134. 9

[EK03] ELAD A., KIMMEL R.: On bending invariant signatures for surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 10 (Oct. 2003), 1285–1295. 2

[ERGB16] EYNARD D., RODOLA E., GLASHOFF K., BRONSTEIN

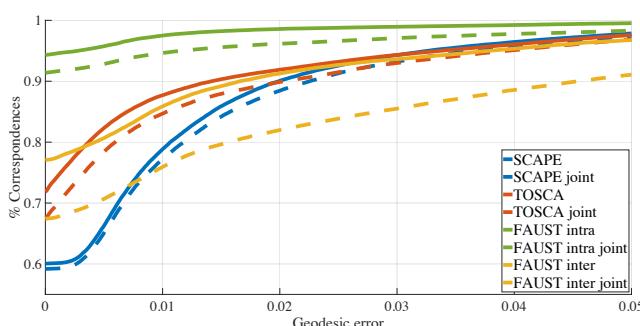


Figure 14: Utilizing inverse map information is advantageous to the matching process as we show in the benchmark results above. We computed maps using our method, applied with and without the inverse map, and we achieve improved results on all datasets, cf. dashed lines (without inverse) to continuous lines (with inverse).

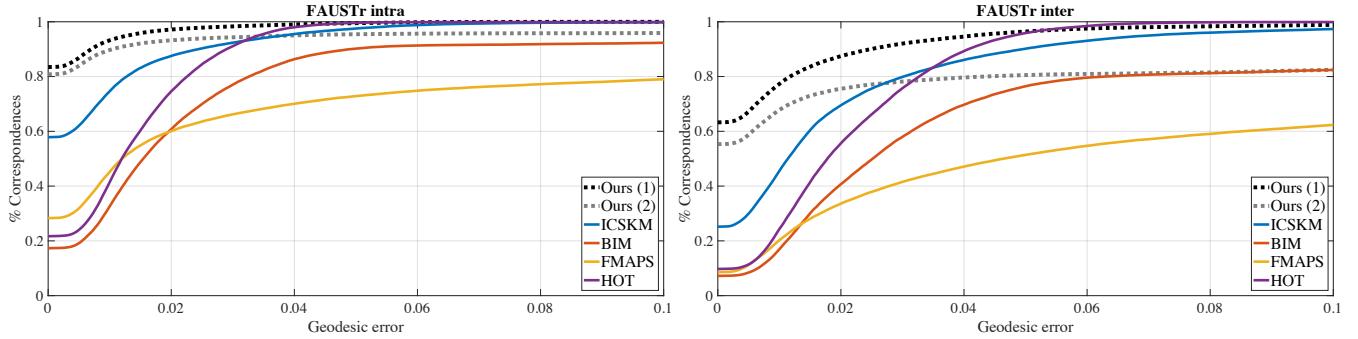


Figure 15: We re-mesh the FAUST dataset such that the resulting shapes do not share the connectivity any longer. We compare the above methods in this challenging scenario on the intra- and inter-class pairs. As can be seen in the geodesic error plots, our technique exhibits the best results by a large margin. We note that while 'Ours (1)' starts from ICSKM which achieves excellent results on its own, our second set of results, denoted by 'Ours (2)', uses BIM maps, and thus serves as a highly inaccurate initialization.

M. M.: Coupled functional maps. In *3D Vision (3DV), 2016 Fourth International Conference on* (2016), IEEE, pp. 399–407. [2](#) [3](#)

[ESBC19] EZUZ D., SOLOMON J., BEN-CHEN M.: Reversible harmonic maps between discrete surfaces. *ACM Transactions on Graphics (TOG)* 38, 2 (2019), 15. [3](#)

[GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: Shrec: shape retrieval contest: Watertight models track. *Online*: <http://watertight.geometri.cnr.it> 7 (2007). [9](#)

[GG84] GEMAN S., GEMAN D.: Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, 6 (1984), 721–741. [6](#)

[GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), ACM Press/Addison-Wesley Publishing Co., pp. 209–216. [9](#)

[HG13] HUANG Q.-X., GUIBAS L.: Consistent shape maps via semidefinite programming. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 177–186. [2](#)

[HO17] HUANG R., OVSJANIKOV M.: Adjoint map representation for shape analysis and matching. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 151–163. [3](#)

[HS85] HOCHBAUM D. S., SHMOYS D. B.: A best possible heuristic for the k-center problem. *Mathematics of operations research* 10, 2 (1985), 180–184. [5](#)

[HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 36. [2](#)

[HZG*12] HUANG Q.-X., ZHANG G.-X., GAO L., HU S.-M., BUTSCHER A., GUIBAS L.: An optimization approach for extracting and encoding consistent maps in a shape collection. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 167. [3](#) [4](#)

[KB57] KOOPMANS T. C., BECKMANN M.: Assignment problems and the location of economic activities. *Econometrica: journal of the Econometric Society* (1957), 53–76. [2](#) [3](#)

[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. In *ACM Transactions on Graphics (TOG)* (2011), vol. 30, ACM, p. 79. [2](#) [6](#) [7](#) [8](#)

[Kol06] KOLMOGOROV V.: Convergent tree-reweighted message passing for energy minimization. *IEEE transactions on pattern analysis and machine intelligence* 28, 10 (2006), 1568–1583. [2](#) [5](#) [6](#) [7](#)

[KS98] KIMMEL R., SETHIAN J. A.: Computing geodesic paths on manifolds. *Proceedings of the national academy of Sciences* 95, 15 (1998), 8431–8435. [7](#)

[LdABN*07] LOIOLA E. M., DE ABREU N. M. M., BOAVENTURA-NETTO P. O., HAHN P., QUERIDO T.: A survey for the quadratic assignment problem. *European journal of operational research* 176, 2 (2007), 657–690. [2](#) [4](#)

[LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. In *ACM Transactions on Graphics (TOG)* (2009), vol. 28, ACM, p. 72. [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. [2](#)

[LRR*17] LITANY O., REMEZ T., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Deep functional maps: Structured prediction for dense shape correspondence. *arXiv preprint arXiv:1704.08686* (2017). [3](#)

[MHK*08] MATEUS D., HORAUD R., KNOSSOW D., CUZZOLIN F., BOYER E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on* (2008), IEEE, pp. 1–8. [2](#)

[NBCW*11] NGUYEN A., BEN-CHEN M., WELNICKA K., YE Y., GUIBAS L.: An optimization approach to improving collections of shape maps. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 1481–1491. [2](#) [3](#) [8](#)

[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](#) [7](#)

[OCB*17] OVSJANIKOV M., CORMAN E., BRONSTEIN M., RODOLÀ E., BEN-CHEN M., GUIBAS L., CHAZAL F., BRONSTEIN A.: Computing and processing correspondences with functional maps. In *ACM SIGGRAPH 2017 Courses* (2017), ACM, p. 5. [2](#)

[PV91] PARDALOS P. M., VAVASIS S. A.: Quadratic programming with one negative eigenvalue is np-hard. *Journal of Global Optimization* 1, 1 (1991), 15–22. [5](#)

[RBA*12] RODOLÀ E., BRONSTEIN A. M., ALBARELLI A., BERGAMASCO F., TORSELLO A.: A game-theoretic approach to deformable shape matching. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on* (2012), IEEE, pp. 182–189. [2](#)

[RDK12] RAVIV D., DUBROVINA A., KIMMEL R.: Hierarchical match-

ing of non-rigid shapes. In *Scale Space and Variational Methods in Computer Vision*. Springer, 2012, pp. 604–615. 2

[RO18] ROUFOUSSE J.-M., OVSJANIKOV M.: Unsupervised deep learning for structured shape matching. *arXiv preprint arXiv:1812.03794* (2018). 3

[RRBW*14] RODOLÀ E., ROTA BULO S., WINDHEUSER T., VESTNER M., CREMERS D.: Dense non-rigid shape correspondence using random forests. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2014), pp. 4177–4184. 3

[SAZK15] SHAMAI G., AFLALO Y., ZIBULEVSKY M., KIMMEL R.: Classical scaling revisited. In *Proceedings of the IEEE International Conference on Computer Vision* (2015), pp. 2255–2263. 7

[SK14] SHTERN A., KIMMEL R.: Iterative closest spectral kernel maps. In *3D Vision (3DV), 2014 2nd International Conference on* (2014), vol. 1, IEEE, pp. 499–505. 2, 6, 7

[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. 2, 4, 7

[SPKS16] SOLOMON J., PEYRÉ G., KIM V. G., SRA S.: Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 72. 3

[SY11] SAHILLIOĞLU Y., YEMEZ Y.: Coarse-to-fine combinatorial matching for dense isometric shape correspondence. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 1461–1470. 2

[SZS*08] SZELISKI R., ZABIH R., SCHARSTEIN D., VEKSLER O., KOLMOGOROV V., AGARWALA A., TAPPEN M., ROTHER C.: A comparative study of energy minimization methods for markov random fields with smoothness-based priors. *IEEE transactions on pattern analysis and machine intelligence* 30, 6 (2008), 1068–1080. 6

[TSDS10] TOMBARI F., SALTÌ S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *European conference on computer vision* (2010), Springer, pp. 356–369. 7

[VKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. In *Computer Graphics Forum* (2011), vol. 30, Wiley Online Library, pp. 1681–1707. 2

[VLB*17] VESTNER M., LÄHNER Z., BOYARSKI A., LITANY O., SLOSSBERG R., REMEZ T., RODOLÀ E., BRONSTEIN A., BRONSTEIN M., KIMMEL R., ET AL.: Efficient deformable shape correspondence via kernel matching. *arXiv preprint arXiv:1707.08991* (2017). 5, 7, 10

[VLR*17] VESTNER M., LITMAN R., RODOLÀ E., BRONSTEIN A., CREMERS D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. *arXiv preprint arXiv:1701.00669* (2017). 2, 4

[WSSC11] WINDHEUSER T., SCHLICKERWEI U., SCHMIDT F. R., CREMERS D.: Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *Computer Vision (ICCV), 2011 IEEE International Conference on* (2011), IEEE, pp. 2134–2141. 3

[XKH*16] XU K., KIM V. G., HUANG Q., MITRA N., KALOGERAKIS E.: Data-driven shape analysis and processing. In *SIGGRAPH ASIA 2016 Courses* (2016), ACM, p. 4. 2