

Isometric Multi-Shape Matching

Maolin Gao¹ Zorah Löhner^{1,3} Johan Thunberg² Daniel Cremers¹ Florian Bernard¹

¹ Technical University of Munich ² Halmstad University ³ University of Siegen




Figure 1. **Left:** We present a novel approach for isometric multi-shape matching based on matching each shape to a (virtual) universe shape (shown semi-transparent). Our formulation represents point-to-point correspondences between shapes i and j as the composition of the shape-to-universe permutation matrix P_i and the universe-to-shape permutation matrix P_j^T . By doing so, the pairwise matchings $P_{ij} = P_i P_j^T$ are by construction cycle-consistent. **Middle:** Our formulation successfully solves isometric multi-matching of partial shapes. **Right:** Due to the cycle-consistency we can use our correspondences to faithfully transfer textures across a shape collection.

Abstract

Finding correspondences between shapes is a fundamental problem in computer vision and graphics, which is relevant for many applications, including 3D reconstruction, object tracking, and style transfer. The vast majority of correspondence methods aim to find a solution between pairs of shapes, even if multiple instances of the same class are available. While isometries are often studied in shape correspondence problems, they have not been considered explicitly in the multi-matching setting. This paper closes this gap by proposing a novel optimisation formulation for isometric multi-shape matching. We present a suitable optimisation algorithm for solving our formulation and provide a convergence and complexity analysis. Our algorithm obtains multi-matchings that are by construction provably cycle-consistent. We demonstrate the superior performance of our method on various datasets and set the new state-of-the-art in isometric multi-shape matching.

1. Introduction

The identification of correspondences between 3D shapes, also known as the shape matching problem, is a longstanding challenge in visual computing. Correspondence

problems have a high relevance due to their plethora of applications, including 3D reconstruction, deformable object tracking, style transfer, shape analysis, or general data canonicalisation, e.g. to facilitate learning by establishing a common vector space representation.

There are certain problem formulations that cover generic correspondence problems involving different types of data and varying application scenarios. One example is the widely-studied quadratic assignment problem (QAP) [36]. Due to its NP-hardness [53], reasonably large QAPs cannot be solved satisfactorily in most settings. However, contrary to bringing generic objects (e.g. graphs) into correspondence, when considering 3D shapes it is possible to exploit particular structural properties in order to effectively solve the shape matching problem.

For example, it has been demonstrated that explicitly modelling the low-dimensional structure of shape matching problems often allows to find global optima for a wide range of shape matching formulations [5]. It was also shown that learning suitable feature representations from shapes improves the matching performance drastically compared to using hand-crafted features [40].

Moreover, when assuming (*near*)-isometries between shapes, efficient and powerful spectral approaches can be leveraged for shape matching [51]. Isometries describe

classes of deformable shapes of the same type but in different poses, *e.g.* humans or animals who are able to adopt a variety of poses. Potential applications for isometric shape matching include AR/VR or template matching. While (near)-isometric shape matching has been studied extensively for the case of matching a pair of shapes, the isometric multi-shape matching problem, where an entire collection of (near-isometric) shapes is to be matched, is less explored. Important applications of isometric multi-shape matching include learning low-dimensional shape space representations [82], motion tracking and reconstruction.

In principle, any pairwise shape matching method can be used for matching a shape collection. To do so, one can select one of the shapes as reference, and then solve a sequence of pairwise shape matching problems between each of the remaining shapes and the reference. However, a major disadvantage is that such an approach has a strong bias due to the choice of the reference. Alternatively, one could solve pairwise shape matching problems between all pairs of shapes in the shape collection. Although this way there is no bias, in general the resulting correspondences are not cycle-consistent. As such, matching shape A via shape B to shape C, may lead to a different correspondence than matching shape A directly to C.

In order to achieve cycle consistency, so-called *permutation synchronisation* methods can be used as post-processing [52]. A disadvantage of synchronisation-based multi-shape matching is that it is a two-stage procedure, where pairwise matchings are obtained in the first procedure, and synchronization is assured in the second. With that, the matching results are often suboptimal – even if one reverts to an alternating procedure using a soft coupling [63]. For *isometric* multi-shape matching this sequential procedure is particularly disadvantageous, since during the second stage the very strong prior about the isometric nature of the shapes is completely ignored.

Although multi-matchings obtained by synchronisation procedures are cycle-consistent, the matchings are often spatially non-smooth and noisy, as illustrated in Sec. 5. From a theoretical point of view, the most appropriate approach for addressing multi-shape matching is based on a unified formulation, where cycle consistency is assured already when the multi-matchings are computed. Although some approaches fit into this category [17, 8], none of the existing methods are tailored explicitly towards isometric matching in order to take full advantage in this setting.

In this work we fill this gap by introducing a generalisation of state-of-the-art isometric two-shape matching approaches towards isometric multi-shape matching. We demonstrate that explicitly exploiting the isometry property leads to a natural and elegant formulation that achieves improved results compared to previous methods. Our main contributions can be summarised as:

- A **novel optimisation formulation** for isometric multi-shape matching.
- An efficient and easy-to-implement algorithm with **provable convergence**.
- **Guaranteed cycle-consistency** without enforcing explicit constraints.
- **Improvements over the state-of-the-art** on various shape matching benchmarks.

2. Related Work

Assignment problems. Shape matching can be formulated as bringing points defined on one shape into correspondence with points on another shape. A simple mathematical formulation for doing so is the linear assignment problem (LAP) [49], where a linear cost function is optimised over the set of permutation matrices. The objective function defines the cost for matching points on the first shape to points on the second shape. In shape matching, the costs are typically computed based on feature descriptors, such as the heat kernel signature [13], wave kernel signature [2], or SHOT [61]. Despite the exponential size of the search space, there exist efficient polynomial-time algorithms to solve the LAP [10]. A downside of the LAP is that the geometric relation between points is not explicitly taken into account, so that the found matchings lack spatial smoothness. To compensate for this, a correspondence problem formulation based on the quadratic assignment problem (QAP) [35, 36, 53, 14, 41] can be used. In that case, in addition to linear point-to-point matching costs, quadratic costs for matching *pairs of points* on the first shape to pairs of points on the second shape are taken into account. Since pairs of points can be understood as edges in a graph, this corresponds to graph matching. Due to the NP-hardness of the QAP [53], there are no algorithms that can reliably find global optima efficiently for large (non-trivial) problem instances. In addition to exhaustive search algorithms that have exponential worst-case time complexity [4], there are various more efficient but non-optimal solution strategies. They include spectral methods [38, 18], convex relaxations [79, 23, 69, 33, 68], some of them relying on path-following [78, 80, 20, 6], as well as various non-convex formulations [37, 66, 72, 27]. For suitably defined matching costs the QAP is an appropriate formalism for *modelling* isometric shape matching. However, due to its NP-hardness the QAP is computationally very difficult to solve. Moreover, due to the generality of the formalism, it does not fully exploit the structural properties present in isometric shape matching problems, and is therefore a sub-optimal choice from a computational perspective.

Isometric shape matching. The near-isometric shape correspondence problem has been studied extensively in the literature, see [60] for a recent survey. Apart from methods tackling a QAP formulation (see previous paragraph),

there exist directions utilising other structural properties of isometries. The Laplace-Beltrami operator (LBO) [54], a generalisation of the Laplace operator on manifolds, as well as its eigenfunctions are invariant under isometries. Methods like [46, 47] directly incorporate this knowledge into the pipeline, or use descriptors based on these [2, 71, 13]. Functional maps [51] reformulate the point-wise correspondence problem as a correspondence between functions. The functional mapping is represented as a low-dimensional matrix for suitably chosen basis functions. The classic choice are the eigenfunctions of the LBO, which are invariant under isometries and predestined for this setting. Moreover, for general non-rigid settings learning these basis functions has also been proposed [43]. A wide variety of extensions to make functional maps more robust or more flexible have been developed. This includes orientation-preservation [56], image co-segmentation [73], denoising [22, 55], partiality [58], and non-isometries [21]. However, extracting a point-wise correspondence from a functional map matrix is not trivial [16, 57]. This is mainly because of the low-dimensionality of the functional map, and the fact that not every functional map matrix is a representation of a point-wise correspondence [51]. In [44], the authors simultaneously solve for point-wise correspondences and functional maps for non-rigid shape matching.

Due to their low-dimensionality and continuous representation, functional maps also serve as the backbone of many deep learning architectures for 3D correspondence. One of the first examples is FMNet [40], which has also been extended for unsupervised learning settings recently [26, 3, 59]. Other learning methods rely on a given template for each class [24] or local neighbourhood encoding to learn a compact representation [39]. The recently conducted SHREC correspondence contest on isometric and non-isometric 3D shapes [19] revealed that there is still room for improvement in both fields.

Generic multi-matching. The multi-matching problem is relatively well-studied for generic settings, e.g. for matching multiple graphs [77, 76, 65, 6, 67, 75], or matching key-points in image collections [74, 70, 42]. A desirable property of multi-matchings is cycle consistency (which we will formally define in Sec. 3.1). Establishing cycle consistency in a given set of pairwise matchings, known as permutation synchronisation, has been addressed extensively in the literature [50, 52, 29, 15, 81, 64, 70, 45, 62, 7].

Multi-shape matching. There are various works that particularly target the matching of multiple shapes. In [29, 32], semidefinite programming relaxations are proposed for the multi-shape matching problem. However, due to the employed lifting strategy, which drastically increases the number of variables, these methods are not scalable to large problems and only sparse correspondences are obtained. In [17], a game-theoretic formulation for establishing multi-




Figure 2. The method by Cosmo et al. [17] leads to extremely sparse multi-matchings (middle), whereas our method obtains dense matchings (right).

matchings is introduced. Due to the use of a sparse modelling approach, the method also has the disadvantage that only few points per shape are matched, see Fig. 2. [30] analyses similarity on large heterogeneous shape collections by imposing cycle consistency on given pairwise functional maps. In [28], tensor maps are introduced for synchronising heterogeneous shape collections using a low-rank tensor decomposition formulation. The work [25] presents a self-supervised learning approach for finding surface deformations. A higher-order projected power iteration approach was presented in [8], which was applied to various multi-matching settings, such as multi-image matching or multi-shape matching. A shortcoming when applying the mentioned multi-shape matching approaches to *isometric* settings is that they do not exploit structural properties of isometric shapes. Hence, they lead to suboptimal multi-matchings, which we experimentally confirm in Sec. 5. One exception is the recent work on spectral map synchronisation [31], which builds upon functional maps and is, in principal, well-suited for isometric multi-shape matching. However, although the authors take into account cycle consistency, respective penalties are only imposed on pairwise functional maps, rather than on the point-wise correspondences. In Sec. 5 we demonstrate that it leads to multi-matchings that have large cycle errors.

3. Background

In this section we introduce our representation for multi-matchings of 3D shapes, formalise the notion of cycle consistency, and provide a recap of functional maps.

3.1. Multi-Matching Representation

We are given a collection $\mathcal{X}_1, \dots, \mathcal{X}_k$ of k 3D shapes, where each shape is a triangular surface mesh that discretises a two-dimensional Riemannian manifold. The i -th shape \mathcal{X}_i is represented by a total of m_i vertices in 3D space. For any two non-negative integers s and t ,

$$\mathbb{P}_{st} = \{P \in \{0, 1\}^{s \times t} : P\mathbf{1}_t \leq \mathbf{1}_s, \mathbf{1}_s^\top P \leq \mathbf{1}_t^\top\}, \quad (1)$$

is the set of partial permutation matrices, where $\mathbf{1}_s$ is the s -dimensional column vector with each element equals to 1.

As such, correspondences between vertices of pairs of shapes \mathcal{X}_i and \mathcal{X}_j can be represented by using the partial permutation matrix $P_{ij} \in \mathbb{P}_{m_i m_j}$. To be more specific, if the element at position (u, v) in P_{ij} has the value 1, the u -th vertex of \mathcal{X}_i is said to be in correspondence with the v -th vertex of \mathcal{X}_j . We assume $P_{ij} = P_{ji}^\top$ and $P_{ii} = \mathbf{I}_{m_i}$, where \mathbf{I}_{m_i} denotes the identity matrix of size m_i .

Cycle consistency (pairwise). For bijective matchings, in which case the P_{ij} are full permutation matrices (the inequalities in (1) become equalities), cycle consistency [29] means that for all $i, j, \ell \in \{1, \dots, k\}$, it holds that

$$P_{ij}P_{j\ell} = P_{i\ell}. \quad (2)$$

Cycle consistency is a natural property and constitutes a necessary condition for the pairwise matchings to correspond to the ground truth. As such, cycle consistency can serve as additional constraint in order to better restrict the space of solutions in multi-matching problems.

Cycle consistency (universe). Instead of using the explicit cycle consistency constraints in (2), one can represent multi-matchings by using *shape-to-universe* matchings [52, 29, 70, 7]. In this case, cycle consistency holds implicitly without having to enforce the constraints (2) in the problem formulation, and without having to develop a customised solution strategy. The union of all *distinct* points across all k shapes are called *universe points*, and we use d to denote the total number of universe points. The shape-to-universe formulation of cycle consistency also applies to the case of partial multi-matchings, which is the setting we are interested in. The main idea of the shape-to-universe representation is that each point in each of the k shapes is brought into correspondence with exactly one of the universe points. Then, all points across the k shapes that are in correspondence with the same universe point are said to be in correspondence with each other [29]. Mathematically, let $P_i \in \mathbb{P}_{m_i d}$ be the partial permutation matrix that represents the matching of the i -th shape to the universe. Since each of the m_i points is assigned to exactly one universe point, we have $P_i \mathbf{1}_d = \mathbf{1}_{m_i}$. Pairwise matchings can be obtained from the shape-to-universe matchings via

$$P_{ij} = P_i P_j^\top. \quad (3)$$

The intuition is that the matching from i to j can be represented as matching i to the universe, followed by matching the universe to j , which is illustrated in Fig. 1.

For our later elaborations it will be convenient to stack all P_i 's into a tall block-matrix, which we define as

$$U = [P_1^\top, P_2^\top, \dots, P_k^\top]^\top. \quad (4)$$

The matrix U is $(m \times d)$ -dimensional, where $m = \sum_{i=1}^k m_i$. Moreover, we introduce the blockwise partial permutation

constraint notation $U \in \mathbb{P}$ (without subscript in \mathbb{P}) to indicate that for each block P_i in U it holds that $P_i \in \mathbb{P}_{m_i d}$ and $P_i \mathbf{1}_d = \mathbf{1}_{m_i}$. We emphasise that by representing multi-matchings in terms of the matrix U , the resulting pairwise matchings are, by definition, cycle-consistent.

3.2. Functional Maps

Functional Maps [51] formulate the correspondence problem as a linear mapping $\mathcal{C}_{ij} : L^2(\mathcal{X}_i) \rightarrow L^2(\mathcal{X}_j)$ between function spaces on the surfaces of $\mathcal{X}_i, \mathcal{X}_j$, rather than as a point-to-point correspondence between vertices. Let $\Phi_i \in \mathbb{R}^{m_i \times b}, \Phi_j \in \mathbb{R}^{m_j \times b}$ be the first b eigenfunctions of the Laplace-Beltrami operator (LBO) [54]. Then \mathcal{C}_{ij} transfers the function F represented in the basis Φ_i to the function G represented in the basis Φ_j , i.e.

$$\mathcal{C}_{ij}(\Phi_i^\dagger F) = \Phi_j^\dagger G. \quad (5)$$

Here, Φ_i^\dagger denotes the Moore-Penrose pseudoinverse of Φ_i . In particular, the optimal \mathcal{C}_{ij} will map compatible functions $F \in L^2(\mathcal{X}_i)$ and $G \in L^2(\mathcal{X}_j)$, e.g. descriptor functions or indicator functions on corresponding points, onto each other. We will use \mathcal{C} without subscripts to describe common properties of all \mathcal{C}_{ij} . Due to the linearity of \mathcal{C} , it can be written as a matrix. Orthogonality of \mathcal{C} is related to area preservation in the correspondence [51] which is also a property of isometries. Thus, we use orthogonality as a prior by projecting all \mathcal{C} 's onto the set of orthogonal matrices

$$\mathbb{O}_b = \{\mathcal{C} \in \mathbb{R}^{b \times b} : \mathcal{C}\mathcal{C}^\top = \mathbf{I}_b\}. \quad (6)$$

Similar to the previous section, we want to impose cycle consistency on the pairwise functional maps \mathcal{C}_{ij} . We do so by defining a shape-to-universe functional map \mathcal{C}_i from \mathcal{X}_i to a (virtual) universe shape. We achieve cycle consistency by composing each pairwise functional map using shape-to-universe functional maps, i.e.

$$\mathcal{C}_{ij} = \mathcal{C}_i \mathcal{C}_j^\top. \quad (7)$$

Analogously to (4), we stack all \mathcal{C}_i into a tall $(kb \times b)$ -dimensional block-matrix that we call

$$Q = [\mathcal{C}_1^\top, \mathcal{C}_2^\top, \dots, \mathcal{C}_k^\top]^\top. \quad (8)$$

In accordance with the definition of the permutation constraint, we define the stacked block-orthogonal constraint $Q \in \mathbb{O}$ (without subscript in \mathbb{O}) that indicates that every block $\mathcal{C}_i \in \mathbb{O}_b$.

4. Isometric Multi-Shape Matching

In this section, we introduce our matching formulation, the optimisation algorithm thereof, and provide a theoretical analysis. Our notation is summarised in Tab. 1.

Symbol	Meaning
k	total number of shapes to be matched
m_i	total number of points in shape i
$m = \sum_{i=1}^k m_i$	total number of points across all k shapes
d	universe size (total number of unique points across all shapes)
$P_i \in \mathbb{P}_{m_i d} \subset \mathbb{R}^{m_i \times d}$	shape-to-universe matching for shape i
$U \in \mathbb{P} \subset \mathbb{R}^{m \times d}$	stack of all shape-to-universe matchings
b	number of LBO basis functions
$\Phi_i \in \mathbb{R}^{m_i \times b}$	eigenfunction of the LBO of shape i
$\Phi \in \mathbb{R}^{m \times kb}$	block-diagonal matrix containing the eigenfunctions of all shapes
$\mathcal{C}_i \in \mathbb{R}^{b \times b}$	shape-to-universe functional map for shape i
$Q \in \mathbb{O} \subset \mathbb{R}^{kb \times b}$	stack of all shape-to-universe functional maps

Table 1. Overview of our notation.

4.1. Problem Formulation

The objective function of our isometric multi-matching formulation (that we will later maximise) reads

$$f(U, Q) = \sum_{i,j=1}^k \langle P_i^\top \Phi_i \mathcal{C}_i, P_j^\top \Phi_j \mathcal{C}_j \rangle \quad (9)$$

$$= \langle U^\top \Phi Q, U^\top \Phi Q \rangle, \quad (10)$$

where $\Phi = \text{diag}(\Phi_1, \dots, \Phi_k) \in \mathbb{R}^{m \times kb}$. The equality between the explicit summation formulation in (9) and the matrix formulation in (10) can be verified by expanding the matrix multiplications. When maximising the objective function, the inner product between the aligned basis functions Φ_i and Φ_j is maximised for all pairs i, j . For that purpose, P_i and P_j permute the vertices in terms of universe points, while \mathcal{C}_i and \mathcal{C}_j align the basis functions on the same universe points via an orthogonal transform. Rewriting each summand of (9) as $\text{tr}((P_i^\top \Phi_i) \mathcal{C}_i \mathcal{C}_j^\top (P_j^\top \Phi_j)^\top)$, we see each operation explicitly: $P_i^\top \Phi_i$ shuffles the vertices into consistent universe ordering, $\mathcal{C}_i \mathcal{C}_j^\top$ composes the (cycle-consistent) functional maps between i and j according to (7).

The overall optimisation is performed with respect to U and Q , with the constraints $U \in \mathbb{P}$ and $Q \in \mathbb{O}$. As such, our isometric multi-shape matching formulation reads

$$\begin{aligned} \max_{U, Q} \quad & \langle U^\top \Phi Q, U^\top \Phi Q \rangle \\ \text{s.t.} \quad & U \in \mathbb{P}, Q \in \mathbb{O}. \end{aligned} \quad (11)$$

4.2. Algorithm

In order to solve Problem (11), we propose a novel projection-based algorithm that we call ISOMUSH (Isometric Multi-Shape Matching). The optimisation alternates between updating U and Q . Each update step involves simple matrix multiplications, as well as the Euclidean projection onto the sets \mathbb{P} and \mathbb{O} . For permutations, as well as different objective functions, a similar strategy has been

Algorithm 1: ISOMUSH algorithm.

Input: Φ, ϵ (relative objective improvement)

Output: U, Q

Initialise: $t \leftarrow 0, U_0 \in \mathbb{P}, Q_0 \in \mathbb{O}$

```

1 repeat
2    $U_{t+1} \leftarrow \text{proj}_{\mathbb{P}}(\Phi Q_t Q_t^\top \Phi^\top U_t)$ 
3    $Q_{t+1} \leftarrow \text{proj}_{\mathbb{O}}(\Phi^\top U_{t+1} U_{t+1}^\top \Phi Q_t)$ 
4    $t \leftarrow t+1$ 
5 until  $\frac{f(U_t, Q_t)}{f(U_{t+1}, Q_{t+1})} \geq 1 - \epsilon$ 

```

proven effective in [71, 8]. We denote the Euclidean projections as $\text{proj}_{\mathbb{P}}(\cdot)$ and $\text{proj}_{\mathbb{O}}(\cdot)$. Each Euclidean projection returns the closest element in the constraint set according to the squared Frobenius norm. For the set \mathbb{O} , it is defined as

$$\begin{aligned} \text{proj}_{\mathbb{O}}(Q) &= \arg \min_{Y \in \mathbb{O}} \|Q - Y\|_F^2 \\ &= \arg \max_{Y \in \mathbb{O}} 2\langle Q, Y \rangle - \langle Y, Y \rangle = \arg \max_{Y \in \mathbb{O}} \langle Q, Y \rangle. \end{aligned} \quad (12)$$

The last equality arises from the orthonormality of all \mathcal{C}_i in Q . The projection onto the set \mathbb{P} is defined analogously, in which case the term $\langle Y, Y \rangle$ has the constant value m for $Y \in \mathbb{P}$ (since the term simply counts the total number of ones in Y , which has the fixed value m because $U \in \mathbb{P}$ implies $U \mathbf{1}_d = \mathbf{1}_m$). By U_t and Q_t we denote the values of U and Q at iteration t , respectively.

U-update. For $Z = \Phi Q_t Q_t^\top \Phi^\top$, the U -update step projects ZU_t onto \mathbb{P} . Hence, the U -update reads

$$U_{t+1} = \text{proj}_{\mathbb{P}}(ZU_t) = \arg \max_{U \in \mathbb{P}} \langle ZU_t, U \rangle \quad (13)$$

$$= \begin{bmatrix} \arg \max_{P_1 \in \mathbb{P}_{m_1 d}} \langle [ZU_t]_1, P_1 \rangle \\ \vdots \\ \arg \max_{P_k \in \mathbb{P}_{m_k d}} \langle [ZU_t]_k, P_k \rangle \end{bmatrix}, \quad (14)$$

where $[ZU_t]_i$ denotes the i -th block (of size $m_i \times d$) of ZU_t . Each block of U in (13) is independent, and consequently can be optimised separately, as written in (14). This reduces the projection into solving k independent (partial) linear assignment problems, which are solved by an efficient implementation [9] of the Auction algorithm [10].

Q-update. For $\bar{Z} = \Phi^\top U_{t+1} U_{t+1}^\top \Phi$, the Q -update step projects $\bar{Z}Q_t$ onto \mathbb{O} . It is given by

$$Q_{t+1} = \text{proj}_{\mathbb{O}}(\bar{Z}Q_t) = \arg \max_{Q \in \mathbb{O}} \langle \bar{Z}Q_t, Q \rangle \quad (15)$$

$$= \begin{bmatrix} \arg \max_{C_1 \in \mathbb{O}} \langle [\bar{Z}Q_t]_1, C_1 \rangle \\ \vdots \\ \arg \max_{C_k \in \mathbb{O}} \langle [\bar{Z}Q_t]_k, C_k \rangle \end{bmatrix}, \quad (16)$$

where $[\bar{Z}Q_t]_i$ denotes the i -th block (of size $b \times b$) of $\bar{Z}Q_t$. Similar as in the U -update, the result for each block of Q in (15) is independent, and can thus be optimised separately. Thus we can solve k independent singular value decompositions (SVDs), each for a small matrix of size $b \times b$.

4.3. Theoretical Analysis

In this section, the properties of the ISOMUSH algorithm is analysed. To this end, we prove that the algorithm converges, and present a complexity analysis.

4.3.1 Convergence

The convergence of our algorithm follows from the monotonicity of the individual updates. Here, we present the respective results, and refer readers to the supplementary material for the proofs.

Lemma 1 $\langle U_t^\top \Phi Q_t, U_{t+1}^\top \Phi Q_t \rangle \geq \langle U_t^\top \Phi Q_t, U_t^\top \Phi Q_t \rangle$ holds for any t .

Proposition 2 (Monotonicity of U -update)

The objective value cannot decrease through the U -update step (13), and $\langle U_{t+1}^\top \Phi Q_t, U_{t+1}^\top \Phi Q_t \rangle \geq \langle U_t^\top \Phi Q_t, U_t^\top \Phi Q_t \rangle$ holds.

Lemma 3 In each iteration t , $\langle U_{t+1}^\top \Phi Q_t, U_{t+1}^\top \Phi Q_{t+1} \rangle \geq \langle U_{t+1}^\top \Phi Q_t, U_{t+1}^\top \Phi Q_t \rangle$ holds.

Proposition 4 (Monotonicity of Q -update)

The objective value cannot decrease through the Q -update (15), and $\langle U_{t+1}^\top \Phi Q_{t+1}, U_{t+1}^\top \Phi Q_{t+1} \rangle \geq \langle U_{t+1}^\top \Phi Q_t, U_{t+1}^\top \Phi Q_t \rangle$ holds.

By combining these properties, and exploiting that U and Q are in compact sets, we obtain the following result:

Theorem 5 (Convergence)

The sequence $(f(U_t, Q_t))_{t=1,2,\dots}$ is monotonically increasing and convergent. Algorithm 1 terminates in finite time.

4.3.2 Complexity Analysis

The steps in the ISOMUSH algorithm comprises matrix multiplications and projections onto the sets \mathbb{P} and \mathbb{O} . In the following, we break down the complexity of each step:

Multiplications in U -update: The term $\Phi Q Q^\top \Phi^\top U$ can be computed as AB for $A = \Phi Q$ and $B = A^\top U$. Computing $A \in \mathbb{R}^{m \times b}$ has complexity $\mathcal{O}(b^2 m)$ (Φ is a block-diagonal matrix). Computing $B = A^\top U \in \mathbb{R}^{b \times d}$ has complexity $\mathcal{O}(bdk)$ (U is a sparse matrix with at most k nonzero elements per column). Finally, computing $AB \in \mathbb{R}^{m \times d}$ has complexity $\mathcal{O}(bdm)$. This results in an complexity of $\mathcal{O}(bm \cdot \max(d, b))$ for the U -step matrix multiplication.

Multiplications in Q -update: The term $\Phi^\top U U^\top \Phi Q$ can be computed as $C^\top D$ for $C = U^\top \Phi$ and $D = CQ$. Computing $C \in \mathbb{R}^{d \times kb}$ has complexity $\mathcal{O}(bdk)$ (U is a sparse matrix with at most k nonzero elements per column, and Φ is a block-diagonal matrix). Computing $D = CQ \in \mathbb{R}^{d \times b}$ has complexity $\mathcal{O}(b^2 dk)$. Computing $C^\top D \in \mathbb{R}^{kb \times b}$ has complexity $\mathcal{O}(b^2 dk)$. This results in an complexity of $\mathcal{O}(b^2 dk)$ for the Q -step matrix multiplication.

Projection onto \mathbb{P} : the projection onto \mathbb{P} is computed by k linear assignment problems of size $m_i \times d$. The auction algorithm has an average time complexity of $\mathcal{O}(d^2 \log(d))$, so that the overall projection has $\mathcal{O}(kd^2 \log(d))$.

Projection onto \mathbb{O} : the projection onto \mathbb{O} is computed by solving k independent projections onto \mathbb{O}_b . Using SVD, this amounts to a complexity of $\mathcal{O}(b^3)$.

5. Experiments

We show the effectiveness of our method on several datasets and compare against state-of-the-art approaches.

Error measure. We evaluate the accuracy of correspondences using the Princeton benchmark protocol [34]. Given the ground-truth correspondences (x_i, x_j^*) for each $x_i \in \mathcal{X}_i$, the error of the calculated match (x_i, x_j) is given by the normalised geodesic distance between x_j and x_j^*

$$e(x_i) = \frac{\text{dist}_{geo}(x_j, x_j^*)}{\text{diam}(\mathcal{X}_j)}, \quad (17)$$

where $\text{diam}(\cdot)$ denotes the shape diameter. We plot the accumulated errors smaller than a certain relative error, which is known as *percentage of correct keypoints* (PCK) curve. The perfect solution results a constant curve at 100%, which amounts to an area under the curve (AUC) of 1.

Cycle consistency. We quantify the cycle consistency of the methods in terms of the *cycle error*, which is the proportion of the number of cycle-consistency violations, divided by the total number of cycles. Without loss of generality we consider only triplet-cycles, see (2).

Methods. We compare our method against several recent state-of-the-art methods, including the pairwise matching approach ZOOMOUT [47], the two-stage approach ZOOMOUT+SYNC that performs synchronisation to achieve cycle consistency in the results produced by ZOOMOUT, as well as the multi-matching methods HIPPI [8] and CONSISTENTZOOMOUT [31]. Both HIPPI and CONSISTENTZOOMOUT utilise only a single representation in isolation, a point-wise and functional map representation, respectively. In contrast, ISOMUSH leverages both simultaneously, and we experimentally show that this is more robust and accurate.

Setup. We use results produced by ZOOMOUT to initialise all other methods. ZOOMOUT itself is initialised by the functional map solution [51] $\min_{C \in \mathbb{R}^{b \times b}} \|FC - G\|_F^2$




Figure 3. Percentage of correct keypoints (PCK) curves for TOSCA, FAUST and SCAPE. Our method leads to better PCK curves (also see the AUC in Tab. 2) than its competitors across all datasets. Dashed lines indicate methods that do not jointly optimise for multi-matchings.

		Ours	HiPPI	ZoomOut+Sync	ZoomOut	ConsistentZoomOut
TOSCA	AUC \uparrow	0.968	0.951	0.943	0.882	0.956
	time [s] \downarrow	28.3	95.2	305.9	164.6	79.9
	cycle error \downarrow	0	0	0	0.68	0.17
FAUST	AUC \uparrow	0.914	0.911	0.909	0.891	0.908
	time [s] \downarrow	23.2	82.8	170.6	122.8	52.9
	cycle error \downarrow	0	0	0	0.41	0.16
SCAPE	AUC \uparrow	0.940	0.938	0.925	0.884	0.922
	time [s] \downarrow	126.5	218.8	552.3	275.2	82.0
	cycle error \downarrow	0	0	0	0.58	0.25

Table 2. Quantitative evaluation in terms of the area under the PCK curve (AUC), the runtime (excluding initialisation), and the cycle error. All values are averaged over all instances for each dataset.

(without regularisers), where F and G are the concatenation of normalised Heat Kernel Signature [13] and SHOT [61]. The output of ZOOMOUT are pairwise correspondences $\{P_{ij}\}$ and pairwise functional maps $\{C_{ij}\}$ between all pairs of shapes. CONSISTENTZOOMOUT directly operates on the $\{C_{ij}\}$, so they are used for its initialisation. In contrast, HiPPI and our method require shape-to-universe representations. To obtain these, we use synchronisation to extract the shape-to-universe representation from the pairwise transformations. By doing so, we obtain the initial U and Q . We refer to this method of synchronising the ZOOMOUT results as ZOOMOUT+SYNC, which directly serves as initialisation for HiPPI and our method. Throughout this section we also report results of the initialisation methods ZOOMOUT and ZOOMOUT+SYNC. Further details can be found in the supplementary material.

5.1. Comparisons to State-of-the-Art Methods

TOSCA dataset. The TOSCA dataset [12] contains 76 shapes from 8 classes depicting different humans and creatures. We downsample all shapes to 2,000 faces. Our method shows state-of-the-art results and surpasses all competitors on this dataset, see Fig. 3 and Tab. 2. Exemplary

matchings of all competing methods can be found in Fig. 4.

FAUST dataset. The FAUST dataset [11] contains real scans of 10 different humans in different poses. We use the registration subset with 10 poses for each class and downsample each shape to 2,000 faces. Our method shows state-of-the-art results on this dataset, see Fig. 3 and Tab. 2. While the PCK curves between ours, ZOOMOUT+SYNC and HiPPI in Fig. 3 are close, the AUC in Tab. 2 shows that our performance is still superior by a small margin. Qualitative results can be found in the supplementary material.

SCAPE dataset. The SCAPE dataset [1] contains 72 poses of the same person, of which we choose 15 randomly and downsample them to 2,000 faces. Our method shows state-of-the-art results on this dataset, see Fig. 3 and Tab. 2. Exemplary matchings of all methods can be found in Fig. 5.

5.2. Multi-Matching of Partial Shapes

We demonstrate that our method applies to the difficult setting of matching partial shapes. As a proof-of-concept, we created a partial dataset by removing parts of shapes from the TOSCA dataset. Most partial matching pipelines include the full reference shape to resolve some of the complexity. Although our optimisation does not need any information about the complete geometry, we use a partiality-adjusted version of ZOOMOUT to obtain the shape-to-universe initialisation for ISOMUSH. In this case, the universe has the dimension of the full shape. Our method finds the correct correspondence among the partial shapes, while being cycle-consistent, see Fig. 1. Partial functional maps are rectangular and low-rank [58], and this experiment shows that our method can handle this more general case. See the supplementary material for details.

6. Discussion & Future Work

Deep learning. It was shown that deep learning is an extremely powerful approach for extracting shape correspondences [40, 26, 59, 25]. However, the focus of this work is on establishing a fundamental optimisation problem for

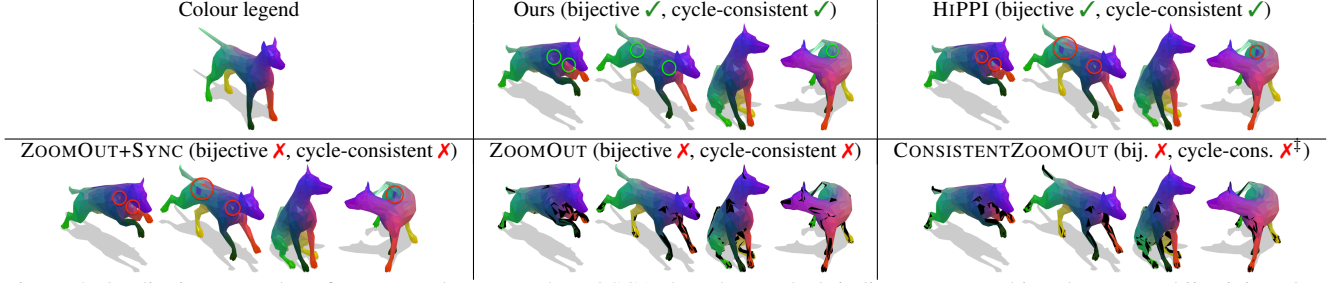


Figure 4. Qualitative examples of correspondences on the TOSCA dog class. Black indicates no matching due to non-bijection. Our method is cycle-consistent and improves upon the non-smooth and noisy correspondences of the two-stage initialisation obtained via ZOOMOUT+SYNC, whereas HiPPI does not (red circles). ZOOMOUT and CONSISTENTZOOMOUT have many unmatched points (black areas). [‡]CONSISTENTZOOMOUT obtains cycle-consistent C_{ij} , but not P_{ij} . (Best viewed magnified on screen)

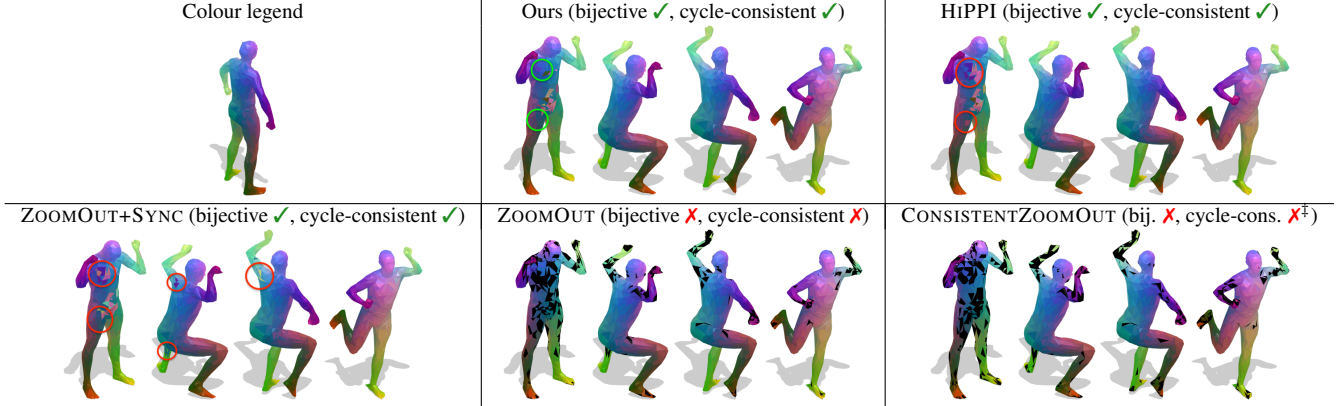


Figure 5. Qualitative examples of correspondences on SCAPE. Black indicates no matching due to non-bijection. As in Fig. 4, our results contain the least noise and are cycle-consistent, although there is one outlier shape where neither HiPPI nor our method could recover from a bad initialisation. [‡]CONSISTENTZOOMOUT obtains cycle-consistent C_{ij} , but not P_{ij} . (Best viewed magnified on screen)

mulation for cycle-consistent isometric multi-shape matching. As such, this work does not focus on learning methods per-se, but we believe that it has a strong potential to spark further work in this direction. In particular, our isometric multi-matching formulation can be integrated into an end-to-end learning framework via differentiable programming techniques [48]. Moreover in machine learning, an entire shape collection is typically used for training, so that our multi-matching setting is conceptually better-suited compared to the traditionally used pairwise matching methods.

Convergence. We have proven that the ISOMUSH algorithm is convergent in the objective $f(\cdot, \cdot)$. However, we did not establish convergence of the variables U and Q . In this context, we note that there are equivalence classes of U and Q that lead to the same objective value. To be more specific, for any (full) $d \times d$ permutation matrix P , and any $C \in \mathbb{O}_b$ we have $(UP) \in \mathbb{P}$, $(QC) \in \mathbb{O}$, and $f(U, Q) = f(UP, QC)$. The latter can be verified by plugging UP and QC into f while making use of the orthogonality of P and C . Although the ISOMUSH algorithm is convergent, and we have empirically verified that it improves upon the state-of-the-art for the isometric multi-shape matching problem, the investigation of stronger con-

vergence results is an interesting direction for future work.

7. Conclusion

We presented a novel formulation for the isometric multi-shape matching problem. Our main idea is to simultaneously solve for shape-to-universe matchings and shape-to-universe functional maps. By doing so, we generalise the popular functional map framework to multi-matching, while guaranteeing cycle consistency, both for the shape-to-universe matchings, as well as for the shape-to-universe functional maps. This contrasts the recent CONSISTENTZOOMOUT [31] method, which does not obtain cycle-consistent multi-matchings. Our algorithm is efficient, straightforward to implement, and monotonically increases the objective function. Experimentally we have demonstrated that our method outperforms recent state-of-the-art techniques in terms of matching quality, while producing cycle-consistent results and being efficient.

Acknowledgements. The authors gracefully acknowledge the support from the ERC Advanced Grant SIMULACRON, the Munich Center for Machine Learning and the Swedish Research Council (2019-04769).

References

- [1] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 408–416. ACM, 2005. 7
- [2] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *ICCV Workshops*, 2011. 2, 3
- [3] Mehmet Aygün, Zorah Löhner, and Daniel Cremers. Unsupervised dense shape correspondence using heat kernels. In *International Conference on 3D Vision (3DV)*, 2020. 3
- [4] M S Bazaraa and A N Elshafei. An exact branch-and-bound procedure for the quadratic-assignment problem. *Naval Research Logistics Quarterly*, 26(1):109–121, 1979. 2
- [5] Florian Bernard, Zeeshan Khan Suri, and Christian Theobalt. Mina: Convex mixed-integer programming for non-rigid shape alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [6] Florian Bernard, Christian Theobalt, and Michael Moeller. DS*: Tighter Lifting-Free Convex Relaxations for Quadratic Matching Problems. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 3
- [7] Florian Bernard, Johan Thunberg, Jorge Goncalves, and Christian Theobalt. Synchronisation of partial multi-matchings via non-negative factorisations. *Pattern Recognition*, 92:146–155, 2019. 3, 4
- [8] Florian Bernard, Johan Thunberg, Paul Swoboda, and Christian Theobalt. Hippie: Higher-order projected power iterations for scalable multi-matching. In *International Conference on Computer Vision (ICCV)*, pages 10284–10293, 2019. 2, 3, 5, 6
- [9] Florian Bernard, Nikos Vlassis, Peter Gemmar, Andreas Husch, Johan Thunberg, Jorge Goncalves, and Frank Hertel. Fast correspondences for statistical shape models of brain structures. In *Medical Imaging 2016: Image Processing*, volume 9784, page 97840R. International Society for Optics and Photonics, 2016. 5
- [10] Dimitri P Bertsekas. *Network Optimization: Continuous and Discrete Models*. Athena Scientific, 1998. 2, 5
- [11] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3794–3801, 2014. 7
- [12] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. 7
- [13] Michael M Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. 2, 3, 7
- [14] Rainer Burkard, Mauro Dell’Amico, and Silvano Martello. Assignment problems. 2009. 2
- [15] Yuxin Chen, Leonidas J Guibas, and Qi-Xing Huang. Near-Optimal Joint Object Matching via Convex Relaxation. 2014. 3
- [16] Etienne Corman, Maks Ovsjanikov, and Antonin Chambolle. Continuous matching via vector field flow. In *Eurographics Symposium on Geometry Processing (SGP)*, 2015. 3
- [17] Luca Cosmo, Emanuele Rodolà, Andrea Albarelli, Facundo Mémoli, and Daniel Cremers. Consistent partial matching of shape collections via sparse modeling. In *Computer Graphics Forum*, volume 36, pages 209–221, 2017. 2, 3
- [18] T Cour, P Srinivasan, and J Shi. Balanced graph matching. *Adv. Neural Inform. Process. Syst.*, 2006. 2
- [19] Roberto Dyke, Caleb Stride, Yu-Kun Lai, Paul L. Rosin, Mathieu Aubry, Amit Boyarski, Alex M. Bronstein, Michael M. Bronstein, Daniel Cremers, Matthew Fisher, Thibault Groueix, Daoliang Guo, Vladimir G. Kim, Ron Kimmel, Zorah Löhner, Kun Li, Or Litany, Tal Remez, Emanuele Rodolà, Bryan C. Russell, Yusuf Sahillioglu, Ron Slossberg, Gary K. L. Tam, Matthias Vestner, Z. Wu, and Jingyu Yang. Shrec’19: Shape correspondence with isometric and non-isometric deformations. In *Eurographics Workshop on 3D Object Retrieval (3DOR)*, 2019. 3
- [20] Nadav Dym, Haggai Maron, and Yaron Lipman. DS++ - A flexible, scalable and provably tight relaxation for matching problems. *ACM Transactions on Graphics (TOG)*, 36(6), 2017. 2
- [21] Marvin Eisenberger, Zorah Löhner, and Daniel Cremers. Smooth shells: Multi-scale shape registration with functional maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [22] Danielle Ezuz and Mirela Ben-Chen. Deblurring and denoising of maps between shapes. *Computer Graphics Forum (CGF)*, 36(5), August 2017. 3
- [23] Fajwel Fogel, Rodolphe Jenatton, Francis Bach, and Alexandre d’Aspremont. Convex Relaxations for Permutation Problems. In *Adv. Neural Inform. Process. Syst.*, 2013. 2
- [24] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. 3d-coded : 3d correspondences by deep deformation. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [25] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. Unsupervised cycle-consistent deformation for shape matching. *Computer Graphics Forum (CGF)*, 2019. 3, 7
- [26] Oshri Halimi, Or Litany, Emanuele Rodola, Alex M Bronstein, and Ron Kimmel. Unsupervised learning of dense shape correspondence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4370–4379, 2019. 3, 7
- [27] Benjamin Holzsuh, Zorah Löhner, and Daniel Cremers. Simulated annealing for 3d shape correspondence. In *International Conference on 3D Vision (3DV)*, 2020. 2
- [28] Qixing Huang, Zhenxiao Liang, Haoyun Wang, Simiao Zuo, and Chandrajit Bajaj. Tensor maps for synchronizing heterogeneous shape collections. *ACM Transactions on Graphics (TOG)*, 38(4):1–18, 2019. 3

- [29] Qi-Xing Huang and Leonidas Guibas. Consistent shape maps via semidefinite programming. In *Symposium on Geometry Processing*, 2013. 3, 4
- [30] Qi-Xing Huang, Fan Wang, and Leonidas Guibas. Functional Map Networks for Analyzing and Exploring Large Shape Collections. *ACM Trans. Graph.*, 2014. 3
- [31] Ruqi Huang, Jing Ren, Peter Wonka, and Maks Ovsjanikov. Consistent zoomout: Efficient spectral map synchronization. In *Computer Graphics Forum*, volume 39, pages 265–278. Wiley Online Library, 2020. 3, 6, 8
- [32] Itay Kezurer, Shahar Kovalsky, Ronen Basri, and Yaron Lipman. Tight relaxations of quadratic matching. *Computer Graphics Forum (CGF)*, 34(5), 2015. 3
- [33] Itay Kezurer, Shahar Z Kovalsky, Ronen Basri, and Yaron Lipman. Tight Relaxation of Quadratic Matching. *Comput. Graph. Forum*, 2015. 2
- [34] Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. In *ACM Transactions on Graphics (TOG)*, volume 30, 2011. 6
- [35] Tjalling C Koopmans and Martin Beckmann. Assignment Problems and the Location of Economic Activities. *Econometrica*, 25(1):53, Jan. 1957. 2
- [36] Eugene L Lawler. The quadratic assignment problem. *Management science*, 9(4):586–599, 1963. 1, 2
- [37] D Khuê Lê-Huu and Nikos Paragios. Alternating direction graph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4914–4922. IEEE, 2017. 2
- [38] Marius Leordeanu and Martial Hebert. A Spectral Technique for Correspondence Problems Using Pairwise Constraints. In *International Conference on Computer Vision (ICCV)*, 2005. 2
- [39] Isaak Lim, Alexander Dielen, Marcel Campen, and Leif Kobbelt. A simple approach to intrinsic correspondence learning on unstructured 3d meshes. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [40] Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 3, 7
- [41] Eliane Maria Loiola, Nair Maria Maia de Abreu, Paulo Oswaldo Boaventura Netto, Peter Hahn, and Tania Maia Querido. A survey for the quadratic assignment problem. *European Journal of Operational Research*, 176(2):657–690, 2007. 2
- [42] Jiayi Ma, Xingyu Jiang, Aoxiang Fan, Junjun Jiang, and Junchi Yan. Image matching from handcrafted to deep features: A survey. *International Journal of Computer Vision*, pages 1–57, 2020. 3
- [43] Riccardo Marin, Marie-Julie Rakotosaona, Simone Melzi, and Maks Ovsjanikov. Correspondence learning via linearly-invariant embedding. In *Adv. Neural Inform. Process. Syst.*, 2020. 3
- [44] Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, and Yaron Lipman. Point registration via efficient convex relaxation. *ACM Transactions on Graphics (TOG)*, 35(4):1–12, 2016. 3
- [45] Eleonora Maset, Federica Arrigoni, and Andrea Fusiello. Practical and Efficient Multi-View Matching. In *International Conference on Computer Vision (ICCV)*, 2017. 3
- [46] Diana Mateus, Radu Horaud, David Knossow, Fabio Cuzzolin, and Edmond Boyer. Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008. 3
- [47] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Makis Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 2019. 3, 6
- [48] Gonzalo Mena, David Belanger, Scott Linderman, and Jasper Snoek. Learning latent permutations with gumbel-sinkhorn networks. In *Int. Conf. Learn. Represent.*, 2018. 8
- [49] James Munkres. Algorithms for the Assignment and Transportation Problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, Mar. 1957. 2
- [50] Andy Nguyen, Mirela Ben-Chen, Katarzyna Welnicka, Yinyu Ye, and Leonidas J Guibas. An Optimization Approach to Improving Collections of Shape Maps. *Computer Graphics Forum*, 30(5):1481–1491, 2011. 3
- [51] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):1–11, 2012. 1, 3, 4, 6
- [52] Deepti Pachauri, Risi Kondor, and Vikas Singh. Solving the multi-way matching problem by permutation synchronization. In *Adv. Neural Inform. Process. Syst.*, 2013. 2, 3, 4
- [53] Panos M Pardalos, Franz Rendl, and Henry Wolkowicz. The Quadratic Assignment Problem - A Survey and Recent Developments. *DIMACS Series in Discrete Mathematics*, 1993. 1, 2
- [54] Ulrich Pinkall and Konrad Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1), 1993. 3, 4
- [55] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured regularization of functional map computations. In *Computer Graphics Forum (Proc. of Symposium on Geometry Processing (SGP))*, 2019. 3
- [56] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (ToG)*, 37(6), 2018. 3
- [57] Emanuele Rodolà, Michael Moeller, and Daniel Cremers. Regularized pointwise map recovery from functional correspondence. In *Computer Graphics Forum*, volume 36, pages 700–711. Wiley Online Library, 2017. 3
- [58] Emanuele Rodolà, Luca Cosmo, Michael Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. *Computer Graphics Forum (CGF)*, 2016. 3, 7
- [59] Jean-Michel Roufousse, Abhishek Sharma, and Maks Ovsjanikov. Unsupervised deep learning for structured shape matching. In *International Conference on Computer Vision (ICCV)*, 2019. 3, 7

- [60] Yusuf Sahillioglu. Recent advances in shape correspondence. *The Visual Computer*, 2019. 2
- [61] Samuele Salti, Federico Tombari, and Luigi Di Stefano. Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding*, 125:251–264, 2014. 2, 7
- [62] Michele Schiavinato and Andrea Torsello. Synchronization over the birkhoff polytope for multi-graph matching. In *International Workshop on Graph-Based Representations in Pattern Recognition*, pages 266–275, 2017. 3
- [63] F. R. Schmidt, E. Toeppe, D. Cremers, and Y. Boykov. Intrinsic mean for semimetric shape retrieval via graph cuts. In *Pattern Recognition (Proc. DAGM)*, volume 4713 of *LNCS*, pages 446–455, Heidelberg, Germany, sep 2007. Springer. 2
- [64] Yanyao Shen, Qixing Huang, Nati Srebro, and Sujay Sanghavi. Normalized Spectral Map Synchronization. In *Adv. Neural Inform. Process. Syst.*, 2016. 3
- [65] Xinchu Shi, Haibin Ling, Weiming Hu, Junliang Xing, and Yanning Zhang. Tensor power iteration for multi-graph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [66] Justin Solomon, Gabriel Peyré, Vladimir G Kim, and Suvrit Sra. Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)*, 35(4):1–13, 2016. 2
- [67] Paul Swoboda, Ashkan Mokarian, Christian Theobalt, Florian Bernard, et al. A convex relaxation for multi-graph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [68] Paul Swoboda, Carsten Rother, Hassan Abu Alhaija, Dagmar Kainmuller, and Bogdan Savchynskyy. A study of lagrangean decompositions and dual ascent solvers for graph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1607–1616, 2017. 2
- [69] L Torresani and V Kolmogorov. A dual decomposition approach to feature correspondence. 35(2):259–271, 2013. 2
- [70] Roberto Tron, Xiaowei Zhou, Carles Esteves, and Kostas Daniilidis. Fast Multi-Image Matching via Density-Based Clustering. In *International Conference on Computer Vision (ICCV)*, 2017. 3, 4
- [71] Matthias Vestner, Zorah Löhner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodola, Alex Bronstein, Michael Bronstein, Ron Kimmel, et al. Efficient deformable shape correspondence via kernel matching. In *International Conference on 3D Vision (3DV)*, 2017. 3, 5
- [72] Matthias Vestner, Roei Litman, Emanuele Rodolà, Alexander M Bronstein, and Daniel Cremers. Product Manifold Filter - Non-Rigid Shape Correspondence via Kernel Density Estimation in the Product Space. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [73] Fan Wang, Qixing Huang, and Leonidas J Guibas. Image co-segmentation via consistent functional maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 849–856, 2013. 3
- [74] Qianqian Wang, Xiaowei Zhou, and Kostas Daniilidis. Multi-Image Semantic Matching by Mining Consistent Features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [75] Runzhong Wang, Junchi Yan, and Xiaokang Yang. Graduated assignment for joint multi-graph matching and clustering with application to unsupervised graph matching network learning. *Adv. Neural Inform. Process. Syst.*, 33, 2020. 3
- [76] Junchi Yan, Minsu Cho, Hongyuan Zha, Xiaokang Yang, and Stephen M Chu. Multi-graph matching via affinity optimization with graduated consistency regularization. 38(6):1228–1242, 2015. 3
- [77] Junchi Yan, Jun Wang, Hongyuan Zha, Xiaokang Yang, and Stephen Chu. Consistency-driven alternating optimization for multigraph matching: A unified approach. *IEEE Trans. Image Process.*, 24(3):994–1009, 2015. 3
- [78] Mikhail Zaslavskiy, Francis Bach, and Jean-Philippe Vert. A Path Following Algorithm for the Graph Matching Problem. 31(12):2227–2242, 2009. 2
- [79] Q Zhao, S E Karisch, F Rendl, and H Wolkowicz. Semidefinite programming relaxations for the quadratic assignment problem. *Journal of Combinatorial Optimization*, 2(1):71–109, 1998. 2
- [80] Feng Zhou and Fernando De la Torre. Factorized Graph Matching. 38(9):1774–1789, 2016. 2
- [81] Xiaowei Zhou, Menglong Zhu, and Kostas Daniilidis. Multi-image matching via fast alternating minimization. In *International Conference on Computer Vision (ICCV)*, 2015. 3
- [82] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2