

# ΣIGMA: Scale-Invariant Global Sparse Shape Matching

Maolin Gao<sup>1,4</sup> Paul Roetzer<sup>3</sup> Marvin Eisenberger<sup>1,4</sup>  
Zorah Löhner<sup>2</sup> Michael Moeller<sup>2</sup> Daniel Cremers<sup>1,4,5</sup> Florian Bernard<sup>3</sup>

<sup>1</sup> Technical University of Munich    <sup>2</sup> University of Siegen    <sup>3</sup> University of Bonn

<sup>4</sup> Munich Center for Machine Learning    <sup>5</sup> University of Oxford

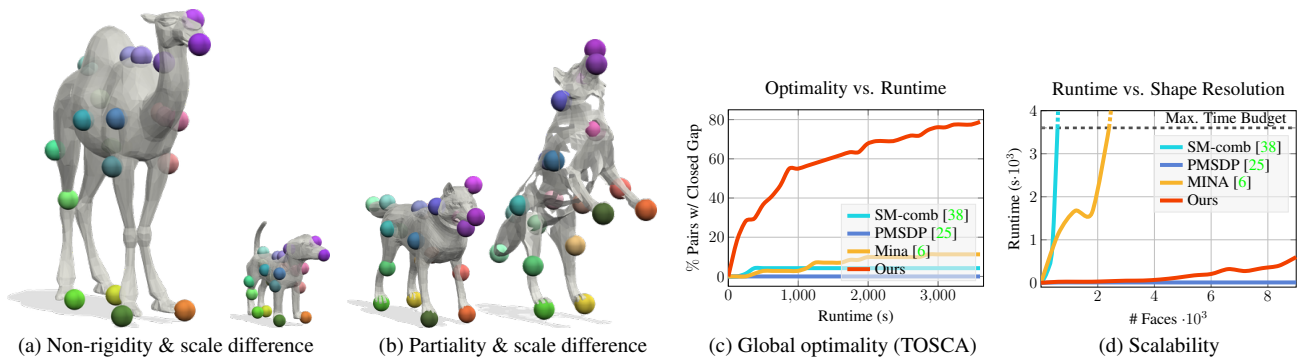


Figure 1: We propose a sparse non-rigid shape matching approach that is provably invariant to rigid transformations and global scaling, can (often) be solved to optimality, and scales linearly with mesh resolution. (a, b) Our matchings for non-rigid shapes with drastically different scale and partiality. (c, d) Our method is the only one that is able to both solve the majority of pairs to global optimality within a time budget of 1h, as well as scaling up to high mesh resolutions.

## Abstract

We propose a novel mixed-integer programming (MIP) formulation for generating precise sparse correspondences for highly non-rigid shapes. To this end, we introduce a projected Laplace-Beltrami operator (PLBO) which combines intrinsic and extrinsic geometric information to measure the deformation quality induced by predicted correspondences. We integrate the PLBO, together with an orientation-aware regulariser, into a novel MIP formulation that can be solved to global optimality for many practical problems. In contrast to previous methods, our approach is provably invariant to rigid transformations and global scaling, initialisation-free, has optimality guarantees, and scales to high resolution meshes with (empirically observed) linear time. We show state-of-the-art results for sparse non-rigid matching on several challenging 3D datasets, including data with inconsistent meshing, as well as applications in mesh-to-point-cloud matching.

## 1. Introduction

Finding correspondences or matchings between parts of 3D shapes is a well-studied problem that lies at the heart

of many tasks in computer vision, computer graphics and beyond. Example tasks related to such 3D shape matching problems include generative shape modelling [49], 3D shape analysis [34], motion capture [46], or model-based image segmentation [5], which are relevant for applications in autonomous driving, robotics, and biomedicine.

While 3D shape matching is traditionally addressed in terms of an optimisation problem formulation, in recent years learning-based approaches became more popular. Specifically, with the advent of geometric deep learning we have witnessed a dramatic improvement in the performance of data-driven methods for 3D shape matching, even for difficult settings such as unsupervised learning for partial shapes [13, 3]. Yet, such learning-based approaches lack theoretical guarantees, both regarding the (global) optimality and often also regarding structural properties of obtained solutions. While optimality guarantees are crucial in safety-critical domains (e.g. autonomous driving, or computer-aided surgery), desirable structural properties may be related to geometric consistency, such as smoothness or the continuity of matchings [48]. From a technical point of view it is typically straightforward to impose respective structural properties within an optimisation-

Method	Init.-free	Solver	% of optimal pairs <sup>†</sup>	Scale Inv.	Rigid Inv.	Data
PMSDP [25]	✓	convex solver	5.0%	✓	✓	point cloud
MINA [6]	✓	MIP solver	11.5%	✗	(✓)*	mesh & point cloud
SM-comb [38]	✗	custom solver	23.5% <sup>‡</sup>	✗	✓	mesh
<b>Ours</b>	✓	MIP solver	73.0%	✓	✓	mesh & point cloud

Table 1: **Overview of axiomatic shape matching approaches** that have a global optimisation flavour. Our approach is the only one that at the same time is initialisation-free, scale- and rigid motion-invariant, and works on both meshes and point clouds. Moreover, ours achieves the best proportion of globally optimal solutions across all experiments. (<sup>†</sup>summarised across all datasets. <sup>‡</sup>w/o runtime budget. \*MINA explicitly optimises for a global rotation using an  $SO(3)$  discretisation.)

based framework. However, resulting formulations are high-dimensional and non-convex, so that finding ‘good’ solutions is a major challenge – many resulting problems lead to large-scale integer linear programming formulations, e.g. when imposing discrete diffeomorphism constraints [48], or when considering the NP-hard quadratic assignment problem [21, 30, 35, 36, 17, 47].

Yet, 3D shape matching typically involves profound structural characteristics (e.g. related to geometric constraints), so that, despite the complexity and high dimensional nature of respective optimisation problems, several formalisms have been proposed that can be efficiently solved in practice. Often they are built on a low dimensional matching representation, e.g. in the spectral domain [27], or in terms of a sparse set of discrete control points that give rise to a dense correspondence [6]. In this work we build upon the latter and propose a novel formalism that has a range of desirable properties, including scale invariance, rigid motion invariance and a substantially better scalability (compared to [6]), which directly leads to a major improvement of the matching quality. In Tab. 1 we provide an overview of the properties of our method and its direct competitors.

**Contributions.** We propose a novel mixed-integer programming formulation for **Scale-Invariant Global sparse shape Matching (SIGMA)** based on a projected Laplace-Beltrami operator, which combines intrinsic and extrinsic geometric information and is applicable to a variety of challenging non-rigid shape matching problems. In summary, our main contributions are:

- A novel initialisation-free mixed-integer programming formulation for sparse shape matching, which can be solved to global optimality for many practical instances and (empirically) scales linearly with the mesh resolution.
- Our method is provably invariant to rigid transformations and global scaling, thus eliminating the extrinsic alignment required by many shape matching pipelines.
- We propose the use of the *projected* Laplace-Beltrami operator for geometry reconstruction, which combines intrinsic and extrinsic geometric information, while still being invariant under actions of the Euclidean group  $E(3)$ .

- We obtain state-of-the-art results on multiple challenging non-rigid shape matching datasets.

## 2. Related Work

Shape matching is a widely studied topic and providing an in-depth survey would be beyond the scope of this paper. We refer the reader to [40] for an overview of non-rigid matching and only mention directly related literature here.

**Quadratic Assignment Problem.** The shape matching problem can be formulated as a quadratic assignment problem (QAP) [23]. The respective QAP aims to find the optimal permutation between (vertices of) two shapes preserving a predefined pairwise property. For non-rigid shape matching this property is often chosen as the geodesic distance between pairs of vertices. However, the QAP is NP-hard [35], so that most methods rely on convex relaxations [41, 12, 7, 19], or heuristics [22, 47]. In consequence, these methods cannot guarantee to find the global optimum of the original non-convex problem and can thus result in arbitrarily bad solutions. Instead of considering a QAP formulation that seeks to find a permutation between all shape vertices, our approach uses a low-dimensional, discrete matching representation that is combined with dense geometry reconstruction, thereby requiring only the computation of a permutation between a small number of keypoints.

**Matching with Optimality Guarantees.** Even though computing global optima is intractable for general QAP formulations, some existing works that consider more specialised shape matching models can be solved efficiently. This is the case for contour-to-contour [42], contour-to-image [15, 43], and contour-to-surface [20, 37] matchings. An analogous approach has been proposed for surface-to-surface matching in [48]. However, it does not scale to high resolutions and global optimality cannot be guaranteed. In [38] an efficient decomposition of [48] has been introduced which scales to higher resolutions but still cannot guarantee optimality. The approach of [6] uses a convex mixed-integer alignment model to find correspondences between two surfaces. While its worst-case time complexity is exponential, it converges to a global optimum in reasonable time in many practical settings.

Another direction was taken by [27] in which point-wise correspondences are transformed to functional correspondences which can then be reduced to a low-dim, continuous problem via the Laplace-Beltrami eigenfunctions. The optimisation for this so-called functional map can be done to global optimality depending on the energy in the spectral domain, but it is hard to guarantee properties like bijectivity for the underlying point-wise correspondence, despite several recent improvement attempts [28, 32].

**Sparse Matching.** Sparse matching methods only compute correspondences for a small subset of vertices of the input. This allows to impose constraints that are otherwise intractable for full resolution shapes, and therefore leads to few but often high quality matches. Many dense correspondence methods can only find local optima [47, 14], and hence rely on a meaningful initialisation, which can f.e. be given by a robust sparse set of initial correspondences [39, 50, 29]. The sparsity of the solution in [36] is achieved by utilising an L1-norm formulation, directly promoting sparsity. In [25] the problem of non-rigid matching is posed as high-dimensional procrustes matching problem which is convexified so that the relaxed problem can be solved to global optimality. The MINA approach of [6] defines a deformation model based on sparse keypoints, and eventually optimises via mixed-integer programming. While conceptually [6] resembles our approach, it explicitly solves for a rotation matrix. In contrast, ours does not require this and is thus significantly more efficient, so that we reduce optimality gaps substantially faster, see Fig. 1c.

**Scale Invariant Shape Analysis.** Shapes with different global scaling are ubiquitous, and techniques such as shape normalisation and scale-invariant feature descriptors have been developed [4, 10]. A scale-invariant metric incorporating Gaussian curvature as the adaptive normalisation factor has been proposed in [1]. Based on that, an intrinsic scale-invariant Laplacian-Beltrami Operator (LBO) can be constructed. In contrast, our projected LBO is also invariant to rigid motions and blended with extrinsic information, which is found useful for shape reconstruction.

### 3. Global Sparse Shape Matching

We address the task of sparse non-rigid shape matching, which we summarise in the following. For each problem instance, we consider a pair of surfaces  $\mathcal{X}$  and  $\mathcal{Y}$  which are given as triangle mesh discretisations of Riemannian 2-manifolds embedded in 3D space. These are defined as the triplets  $\mathcal{X} = (\mathbf{X}, \mathbf{F}^{(\mathcal{X})}, \mathcal{I})$  and  $\mathcal{Y} = (\mathbf{Y}, \mathbf{F}^{(\mathcal{Y})}, \mathcal{J})$ , where the first element denotes the vertices, the second element the faces, and the third element the indices of keypoints (*i.e.* a subset of vertex indices). Tab. 2 summarises our notation.

Our aim is to determine optimal correspondences between the keypoint vertices  $\mathbf{X}_{\mathcal{I}}$  and  $\mathbf{Y}_{\mathcal{J}}$  (of  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively), which we represent as a permutation ma-

Symbol	Description
$\mathcal{X} = (\mathbf{X}, \mathbf{F}^{(\mathcal{X})}, \mathcal{I})$	shape $\mathcal{X}$
$ \mathbf{X} ,  \mathbf{F}^{(\mathcal{X})} $	number of vertices and faces in $\mathcal{X}$
$\mathbf{X} \in \mathbb{R}^{ \mathbf{X}  \times 3}$	vertices of shape $\mathcal{X}$
$\mathbf{F}^{(\mathcal{X})} \in \mathbb{N}^{ \mathbf{F}^{(\mathcal{X})}  \times 3}$	faces of shape $\mathcal{X}$
$\mathcal{I} \in \mathbb{N}^n$	keypoint indices on shape $\mathcal{X}$
$\mathbf{X}_{\mathcal{I}} \in \mathbb{R}^{n \times 3}$	keypoint coordinates on shape $\mathcal{X}$
$\tilde{\mathbf{X}} \in \mathbb{R}^{ \mathbf{X}  \times 4}$	vertices of shape $\mathcal{X}$ in homogeneous coordinates (appended with 1)
$\hat{\mathbf{X}} \in \mathbb{R}^{ \mathbf{X}  \times 3}$	reconstruction of $\mathbf{X}$ with connectivity of $\mathcal{X}$ and pose of $\mathcal{Y}$ (same vertex ordering as $\mathbf{X}$ )
$\Pi^{(\mathcal{X})} \in \mathbb{R}^{ \mathbf{X}  \times  \mathbf{X} }$	projection onto the null space of $\tilde{\mathbf{X}}$
$\mathbf{h}^{(\mathcal{X})} \in \mathbb{R}^{ \mathbf{X} }$	orientation-aware features of $\mathcal{X}$
$\mathbf{h}_{\mathcal{I}}^{(\mathcal{X})} \in \mathbb{R}^n$	orientation-aware features of $\mathcal{X}$ 's keypoints
$\mathcal{Y} = (\mathbf{Y}, \mathbf{F}^{(\mathcal{Y})}, \mathcal{J})$	shape $\mathcal{Y}$
$\vdots$	(analogous as above)
$\mathbf{P} \in \{0, 1\}^{n \times n}$	permutation of keypoints

Table 2: Summary of our notation.




Figure 2: We **utilise synergies between shape matching and shape reconstruction**. To this end, rigid motion-invariant geometric information of shape  $\mathcal{X}$ , encoded in the projected LBO  $\Delta_{\text{proj}}^{(\mathcal{X})}$ , and the keypoint coordinates  $\mathbf{Y}_{\mathcal{J}}$  of shape  $\mathcal{Y}$ , are combined to reconstruct shape  $\mathcal{X}$  in the pose of shape  $\mathcal{Y}$ . Our motivation is that high quality correspondences will lead to better reconstruction and vice versa.

trix  $\mathbf{P} \in \{0, 1\}^{n \times n}$ . In order to find this permutation, we consider a regularisation that utilises synergies between shape matching and shape reconstruction. To this end, we utilise *rigid motion-invariant* geometric information of shape  $\mathcal{X}$ , and then use keypoints of shape  $\mathcal{Y}$  in order to reconstruct shape  $\mathcal{X}$  in the pose of shape  $\mathcal{Y}$  (and vice-versa with swapped roles of  $\mathcal{X}$  and  $\mathcal{Y}$ ), see Fig. 2.

For the rigid motion-invariant geometry encapsulation we introduce the projected Laplace-Beltrami operator in Sec. 3.1 and our sparse matching formalism in Sec. 3.2.

#### 3.1. Projected Laplace-Beltrami Operator

To encapsulate rigid motion-invariant geometric information, we propose a variant of the Laplace-Beltrami operator that we call *Projected Laplace-Beltrami operator* (PLBO). For  $\Delta_{\text{stiff}}^{(\mathcal{X})}$  being the stiffness matrix component of

the Laplacian [31], we define the PLBO as

$$\Delta_{\text{proj}}^{(\mathcal{X})} := (\mathbf{\Pi}^{(\mathcal{X})})^\top \Delta_{\text{stiff}}^{(\mathcal{X})} \mathbf{\Pi}^{(\mathcal{X})}, \quad (1)$$

where  $\mathbf{\Pi}^{(\mathcal{X})}$  denotes the projection matrix

$$\mathbf{\Pi}^{(\mathcal{X})} := \mathbf{I} - \tilde{\mathbf{X}}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{X}})^{-1} \tilde{\mathbf{X}}^\top, \quad \text{with } \tilde{\mathbf{X}} := (\mathbf{X} \ \mathbf{1}), \quad (2)$$

and  $\mathbf{1} \in \mathbb{R}^{|\mathbf{X}|}$  is a vector of all ones. To obtain an operator that is scale-invariant, the definition of Eqn. (1) is area-normalised, *i.e.* based only on the stiffness matrix component  $\Delta_{\text{stiff}}^{(\mathcal{X})}$  of the Laplacian. Due to the use of the coordinate function, the PLBO is not purely intrinsic but also incorporates extrinsic information. We found this mix of intrinsic and extrinsic information beneficial for accurate correspondence computation, and even though the PLBO uses the extrinsic coordinate function, it is still invariant under transformations in the Euclidean group  $E(3)$ :

**Lemma 1.** Let  $\Delta(\mathbf{X}) := \Delta_{\text{proj}}^{(\mathcal{X})} \in \mathbb{R}^{|\mathbf{X}| \times |\mathbf{X}|}$  be the projected Laplace-Beltrami operator for the vertices  $\mathbf{X}$ , defined in Eqn. (1). For any rigid body transformation

$$\begin{pmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{pmatrix} \in E(3), \quad \text{with } \mathbf{R} \in O(3), \mathbf{t} \in \mathbb{R}^3, \quad (3)$$

it holds that  $\Delta(\mathbf{X}) = \Delta(\mathbf{X}\mathbf{R}^\top + \mathbf{1}\mathbf{t}^\top)$ .

See the supplementary material for the proof.

While there are several popular deformation models in the literature [8, 16, 44, 45], we find that existing formulations often depend on the extrinsic pose of the input surfaces. Therefore, such methods require rigidly aligned poses since applying a rotational offset  $\mathbf{X}\mathbf{R}^\top$  can alter the results for any  $\mathbf{R} \in O(3)$ . Compared to the standard LBO operator  $\Delta^{(\mathcal{X})} \in \mathbb{R}^{|\mathbf{X}| \times |\mathbf{X}|}$ , the PLBO  $\Delta_{\text{proj}}^{(\mathcal{X})}$  projects the original coordinate function  $\mathbf{X}$  onto its null space, so that when used as a regulariser (as we introduce in Eqn. (6)) only components outside the null space are penalised. In practice, this reduces oversmoothing and leads to more accurate reconstructions, see Fig. 3 for an illustration.

### 3.2. Sparse Matching

Our final optimisation problem comprises a reconstruction term  $\mathcal{E}_{\text{rec}}$ , a deformation term  $\mathcal{E}_{\text{def}}$  and an orientation-aware term  $\mathcal{E}_{\text{ori}}$ , and reads

$$\begin{aligned} \min_{\mathbf{P}, \hat{\mathbf{X}}, \hat{\mathbf{Y}}} \quad & \mathcal{E}_{\text{rec}}(\mathbf{P}, \hat{\mathbf{X}}, \hat{\mathbf{Y}}) + \lambda_{\text{def}} \mathcal{E}_{\text{def}}(\hat{\mathbf{X}}, \hat{\mathbf{Y}}) + \lambda_{\text{ori}} \mathcal{E}_{\text{ori}}(\mathbf{P}) \\ \text{s.t.} \quad & \mathbf{P} \in \{0, 1\}^{n \times n}, \mathbf{P}^\top \mathbf{1}_n = \mathbf{1}_n, \mathbf{P} \mathbf{1}_n = \mathbf{1}_n. \end{aligned} \quad (4)$$

Our formulation is provably invariant under global scaling and rigid-body motions of the inputs:

**Lemma 2.** Let  $(\mathbf{P}, \hat{\mathbf{X}}, \hat{\mathbf{Y}})$  be a global optimiser of Eqn. (4).




Figure 3: Comparison of **reconstructed shapes** using the standard **LBO** and our proposed **projected LBO**. From the definition in Eqn. (1) we can show that  $\Delta_{\text{proj}}^{(\mathcal{X})} \mathbf{X} = \mathbf{0}$ , *i.e.* this yields a deformation prior that incurs no cost on exact reconstructions of the initial geometry  $\mathbf{X}$ . Thus, the projected LBO is able to better preserve local geometry and leads to more realistic reconstructions. See the supplementary material for details.

- Let  $\mathcal{X}' := (s\mathbf{X}, \mathbf{F}^{(\mathcal{X})}, \mathcal{I})$  be a rescaled input shape  $\mathcal{X}$ , where a scalar factor  $s > 0$  is applied to the vertex coordinates. Then  $(\mathbf{P}', \hat{\mathbf{X}}', \hat{\mathbf{Y}}') := (\mathbf{P}, \hat{\mathbf{X}}, s\hat{\mathbf{Y}})$  is a global optimiser of Eqn. (4) between  $\mathcal{X}'$  and  $\mathcal{Y}$ .
- Let  $\mathcal{X}'' := (\mathbf{X}\mathbf{R}^\top + \mathbf{1}\mathbf{t}^\top, \mathbf{F}^{(\mathcal{X})}, \mathcal{I})$  be a rigidly transformed version of  $\mathcal{X}$  with  $\mathbf{R} \in SO(3), \mathbf{t} \in \mathbb{R}^3$ . Then  $(\mathbf{P}'', \hat{\mathbf{X}}'', \hat{\mathbf{Y}}'') := (\mathbf{P}, \hat{\mathbf{X}}, \hat{\mathbf{Y}}\mathbf{R}^\top + \mathbf{1}\mathbf{t}^\top)$  is a global optimiser of Eqn. (4) between  $\mathcal{X}''$  and  $\mathcal{Y}$ .

Lemma 2 holds analogously for rescaling and rigid transformations of  $\mathcal{Y}$  due to the symmetry of the formulation. We provide proofs in the supplementary material.

The **reconstruction term**  $\mathcal{E}_{\text{rec}}$  encourages the reconstructed keypoints  $\hat{\mathbf{X}}_{\mathcal{I}}$  to be best aligned with the given keypoints of  $\mathcal{Y}$  after they have been reordered via the permutation  $\mathbf{P}$  (we apply the same in a symmetric manner for reconstructing keypoints  $\hat{\mathbf{Y}}_{\mathcal{J}}$  from keypoints of  $\mathcal{X}$  reordered via  $\mathbf{P}^\top$ ):

$$\mathcal{E}_{\text{rec}} = \frac{1}{nd_{\mathcal{Y}}} \|\hat{\mathbf{X}}_{\mathcal{I}} - \mathbf{P}\mathbf{Y}_{\mathcal{J}}\|_F + \frac{1}{nd_{\mathcal{X}}} \|\hat{\mathbf{Y}}_{\mathcal{J}} - \mathbf{P}^\top \mathbf{X}_{\mathcal{I}}\|_F, \quad (5)$$

where  $d_{\mathcal{X}}$  and  $d_{\mathcal{Y}}$  denote the diameter of  $\mathcal{X}$  and  $\mathcal{Y}$ , *i.e.* the maximum among the geodesic distances between all pairs of vertices for each shape, respectively.

The **deformation term**  $\mathcal{E}_{\text{def}}$  favours reconstructions  $\hat{\mathbf{X}}$  which have a similar local geometry to  $\mathcal{X}$  as it was encoded in its PLBO  $\Delta_{\text{proj}}^{(\mathcal{X})}$  (again, we apply the same in a symmetric manner for favouring reconstructions  $\hat{\mathbf{Y}}$  with a similar local geometry to  $\mathcal{Y}$  as encoded by  $\Delta_{\text{proj}}^{(\mathcal{Y})}$ ):

$$\mathcal{E}_{\text{def}} = \frac{1}{|\mathbf{X}|d_{\mathcal{Y}}} \|\Delta_{\text{proj}}^{(\mathcal{X})} \hat{\mathbf{X}}\|_F + \frac{1}{|\mathbf{Y}|d_{\mathcal{X}}} \|\Delta_{\text{proj}}^{(\mathcal{Y})} \hat{\mathbf{Y}}\|_F. \quad (6)$$



Fig. 3 shows how the use of the PLBO affects  $\hat{\mathbf{X}}$ .

The **orientation-aware term**  $\mathcal{E}_{\text{ori}}$  preserves the surface orientation in the correspondence through extrinsic orientation-aware feature maps  $\mathbf{h}^{(\mathcal{X})}$  and  $\mathbf{h}^{(\mathcal{Y})}$ :

$$\mathcal{E}_{\text{ori}} = \frac{1}{n} \|\mathbf{h}_{\mathcal{I}}^{(\mathcal{X})} - \mathbf{P}\mathbf{h}_{\mathcal{J}}^{(\mathcal{Y})}\|_F. \quad (7)$$

Even though the PLBO uses extrinsic information, it is still agnostic to intrinsic symmetries, such as left-right flips, due to the in-built invariance under the group  $E(3)$ , which includes mirror-symmetries. Thus, we propose the orientation-aware regularisation term to disambiguate such intrinsic symmetries. Similar solutions were developed previously for the functional maps framework [33]. Our orientation-aware feature map is defined as

$$\mathbf{h}^{(\mathcal{X})} = \begin{pmatrix} \langle (\nabla \mathbf{f}^{(\mathcal{X})})_1 \times (\nabla \mathbf{g}^{(\mathcal{X})})_1, \mathbf{n}_1^{(\mathcal{X})} \rangle \\ \vdots \\ \langle (\nabla \mathbf{f}^{(\mathcal{X})})_{|\mathcal{X}|} \times (\nabla \mathbf{g}^{(\mathcal{X})})_{|\mathcal{X}|}, \mathbf{n}_{|\mathcal{X}|}^{(\mathcal{X})} \rangle \end{pmatrix}, \quad (8)$$

where  $\mathbf{n}_i^{(\mathcal{X})} \in \mathbb{R}^3$  are the unit outer normals at the  $i$ -th vertex of  $\mathcal{X}$ , and  $\mathbf{f}^{(\mathcal{X})}$  and  $\mathbf{g}^{(\mathcal{X})}$  are two distinct scalar fields (see Sec. 4.2 and supplementary material for details). The outer product of their normalised gradient fields convey their orientation through the right-hand rule. Thus, the resulting feature field  $\mathbf{h}^{(\mathcal{X})}$  implicitly encodes such orientation information. Comparing the features on both surfaces  $\mathcal{X}$  and  $\mathcal{Y}$  allows us to disambiguate intrinsic symmetries.

At the same time,  $\mathbf{h}^{(\mathcal{X})}$  maintains (proper) rotation invariance. This follows directly from the fact that the gradient fields  $\nabla \mathbf{f}^{(\mathcal{X})}$ ,  $\nabla \mathbf{g}^{(\mathcal{X})}$  and normal field  $\mathbf{n}^{(\mathcal{X})}$  are both rotation-equivariant, leaving their inner product unchanged.

## 4. Experiments

In this section, we qualitatively and quantitatively evaluate our proposed method on several shape matching datasets (including TOSCA [9], SMAL [51], SHREC20 Non-Isometric [11] and the DeformingThings4D-Matching dataset [24]), and study its performance in terms of accuracy, global optimality and global scaling dependency.

### 4.1. Evaluation Metric

**Accuracy.** We evaluate the correspondence accuracy according to the Princeton benchmark protocol [18], which reports the *percentage of correct keypoints (PCK)* for the given sparse control points. For dense methods [38], we only consider this subset of sparse keypoints to enable a direct comparison of accuracy.

**Optimality.** We use the *relative optimality gap* to measure the degree of optimality. It is defined as

$$\text{gap} = \left| \frac{\overline{\text{obj}} - \text{obj}}{\overline{\text{obj}}} \right|, \quad (9)$$

where  $\overline{\text{obj}}$  is the best upper bound, and  $\text{obj}$  the best lower bound of the objective.

### 4.2. Implementation Details

**Optimisation.** The optimisation of all MIP-based methods was performed using MOSEK 10.0.35 [2] on a desktop computer with AMD Ryzen 9 5950X 16-core processor and 64GB RAM. A time budget of 1h is set for each problem instance and its global optimum is considered to be found if the relative gap is reduced below  $10^{-2}$ .

**Parameters.** We set  $\lambda_{\text{def}} = 5$  and  $\lambda_{\text{ori}} = 0.025$  in Eqn. (4) and choose two entries with different frequencies from the scale-invariant wave kernel signature [4] as  $\mathbf{f}^{(\mathcal{X})}$  and  $\mathbf{g}^{(\mathcal{X})}$  for the orientation-aware feature defined in Sec. 3.2. We coin the full formulation in Eqn. (4) as *Ours* and additionally show results without the orientation term  $\mathcal{E}_{\text{ori}}$  as *Ours w/o Ori.* Moreover, we note that dropping either  $\mathcal{E}_{\text{def}}$  or  $\mathcal{E}_{\text{rec}}$  will lead to a linear assignment problem for the permutation  $\mathbf{P}$  and a failed reconstruction in  $\hat{\mathbf{X}}$ ,  $\hat{\mathbf{Y}}$ .

**Solution Pruning.** We employ the same solution pruning strategy as done in [6] to reduce the search space. For each keypoint the pruning process keeps  $k$  potential matches based on the similarity of the histogram of geodesics at each keypoint, see [6] for details. In our experiments, we apply this pruning with  $k = 11$  to all methods that admit an initial search space reduction, namely to our method, MINA and PMSDP.

### 4.3. Competing Methods

We compare against state-of-the-art matching methods which exhibit a global optimisation flavour.

**SM-comb** is a combinatorial solver for the elastic shape matching formalism [48] proposed by Roetzer et al. [38]. It solves an integer linear program to find triangle-triangle matchings. This approach heavily relies on the discretisation of both shapes and might fail to produce a feasible solution entirely in cases of large triangulation discrepancies.

**PMSDP** is a convex semidefinite programming relaxation approach proposed by Maron et al. [25]. It uses the LBO eigenbases as a high-dimensional feature embedding in the procrustes matching problem and is therefore an intrinsic and global scaling invariant formalism. Its implementation uses a pruning strategy to rule out unlikely correspondences using the average geodesic distance (AGD) descriptor [18]. We found this to be too aggressive and often pruning the correct ground truth matches, thus, leading to poor performance. It is reported in Fig. 4 under *PMSDP vanilla*. Hence, we replace its pruning strategy with ours, as discussed in Sec. 4.2, while keeping all other default options, and coin it as *PMSDP tuned*.

**MINA.** This method proposed by Bernard et al. [6] is the closest to ours among all baselines. Although MINA also employs a MIP model for sparse deformable shape

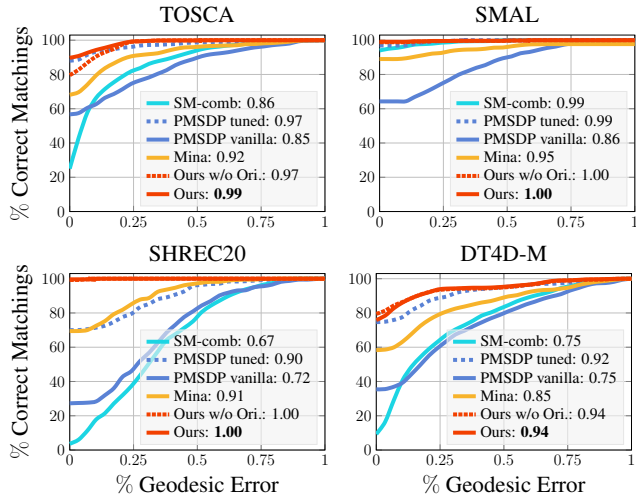


Figure 4: **PCK curves** on datasets TOSCA [9], SMAL [51], SHREC20 [11] and DT4D-M [24]. The values in the legends are *Area-Under-the-Curves* (AUC  $\uparrow$ ). Across all datasets our method consistently outperforms all other approaches. For SMAL and SHREC20 our method produces nearly perfect results. The orientation-aware term improves our performance in most cases.

matching problems, the key difference is that their approach is extrinsic due to its piece-wise affine deformation of the source to the target shape that is combined with a global rigid transformation, which in turn makes it computationally more expensive. Note that MINA originally uses an aggressive search space pruning which we found to prune a large portion of correct matchings. Thus, we use a more conservative pruning by allowing twice as many potential matches as MINA.

#### 4.4. Isometric Shape Matching

**TOSCA** [9] is a popular dataset for the task of isometric shape matching. It contains shapes of 9 different categories and for each category, multiple non-rigidly deformed shapes. In our experiment, the sparse keypoints provided by [18] with ground truth labels are considered for matching, leading to 71 different isometric pairs with 21 to 46 keypoints. We run all experiments with a time budget of 1h<sup>1</sup>.

Fig. 4 summarises the PCK for the keypoints and Fig. 5 (left) unveils the statistics of the final relative optimality gaps. Fig. 6 shows qualitative matching results. Our proposed method outperforms all baselines and is able to produce the lowest relative gaps. More specifically, ours successfully certifies 56 out of the total 71 matching pairs whereas MINA can only certify 8 pairs (cf. Fig. 1c, Tab. 3). Note that these numbers are different to the ones reported in

<sup>1</sup>SM-comb runs till its end due to the lack of time limitation.




Figure 5: (Left) **Relative optimality gap statistics** on the TOSCA. The Inf gap reflects cases where SM-comb [38] fails to find a solution. Our proposed method produces the lowest and most concentrated relative optimality gap. (Right) Mean geodesic error for shape pairs with **different shape scales**. PMSDP [25] and our method are the only ones which perform consistently across different shape scales, with ours being more accurate.

MINA [6] due to their substantially more aggressive search space pruning. PMSDP computes sparse correspondences very fast due to its convex (relaxed) formulation of the original non-convex problem. However, it requires an additional post-processing step to project onto the space of permutations and the final results have larger relative optimality gaps and are less accurate. SM-comb’s final estimation also has larger optimality gaps, which results in bad matching performance. Consequently, it is only able to provide optimality certification for 13 out of all 71 pairs (cf. Tab. 3).

**SMAL** [51] is a 3D animal dataset containing near-isometric shapes, such as foxes and dogs. For our experiments we use furthest point sampling (FPS) to select 25 sparse keypoints, consistently subsample the shapes to 5k faces, and then transfer the selected keypoints to the lower resolution by nearest neighbour search. This introduces noise into the keypoint positions while leaving the mesh topology consistent. Furthermore, the shapes are not pre-aligned and used ‘as is’ in our experiments.

As shown in Fig. 4, our proposed method works well in this setting and solves all matching instances to global optimality within the 1h budget (cf. Tab. 3). SM-comb [38] can produce high quality matchings due to the consistent meshing, which, however, is often not available in practice. MINA [6] can handle near-isometric shapes and estimate accurate correspondences. However, it fails to close the relative optimality gap and thus cannot certify global optimality. The same holds for PMSDP [25].

#### 4.5. Non-Isometric Shape Matching

**SHREC20** [11] contains 14 non-isometric, high resolution animal shapes, ranging from dog and cow to giraffe and elephant. It provides consensus-based sparse keypoint correspondences which we use in our evaluation. Note that these keypoints are manually selected by a human and an




Figure 6: **Qualitative results** on instances of DT4D-M [24] (①–③), SHREC20 [11] (④–⑥), SMAL [51] (⑦,⑧) and TOSCA [9] (⑨,⑩). While our approach consistently produces good results across all datasets, the other approaches have difficulties: SM-comb [38] (which is continuous and thus may match keypoints of shape  $\mathcal{X}$  to any point of shape  $\mathcal{Y}$ ) cannot always find a feasible solution, which results in incomplete solutions (③, ⑥, ⑨); PMSDP [25] especially struggles with non-isometric shape pairs (①, ④, ⑥), similarly as MINA [6] (④, ⑧).

absolute ground truth does not exist. In total, 25 shape pairs are randomly selected from the dataset.

Our proposed method outperforms all competitors on this dataset and finds the global optima for all 25 matching pairs (cf. Tab. 3). Furthermore, all competing methods deteriorate heavily in this non-isometric setting. While it is much harder for MINA [6] to deform the shapes, the high-dimensional embedding of LBO eigenfunctions used in PMSDP [25] breaks for non-isometric shapes, hence their unsatisfactory performance. SM-comb [38] suffers from the large scale variation and the inconsistent triangulation present in the dataset.

**DeformingThings4D-Matching** (DT4D-M) [24] is a recent dataset with humanoid shapes in various poses and inconsistent meshing. Dense but incomplete ground truth correspondences are provided for evaluation. We randomly created 60 shape pairs, consistently sample 30 keypoints and then (independently) subsample shapes to about 10k faces. In this way, the selected keypoints are not in exact

correspondence. This dataset often exhibits non-isometric and non-rigid deformations at the same time and poses a challenging task for all methods.

As shown in Fig. 4, ours yields the best accuracy and obtains the most globally optimal solutions among all baselines despite the present difficulties (Tab. 3). Note that  $\mathcal{E}_{\text{ori}}$  is less helpful here than in TOSCA since it is designed to disambiguate symmetry based on the orientation-aware  $\mathbf{h}^{(\cdot)}$  which is more accurate for isometric shapes. DT4D-M (cf. TOSCA) has less isometric shapes (some without any symmetry, e.g. *prisoner*), so  $\mathcal{E}_{\text{ori}}$  has less impact (Fig. 6). Besides the challenges mentioned above, many of its shapes contain non-manifold structures which empirically causes more difficulties.

#### 4.6. Global Scaling

We showed the scale invariance of our method in Lemma 2 in theory, but also validate our claims with experiments. Shapes with different scales exist naturally, such as the dog-camel and bison-elephant shown in Figs. 1a & 6,



	SM-comb	MINA	PMSDP	Ours
TOSCA (71)	13	8	4	<b>56</b>
SMAL (44)	31	9	2	<b>44</b>
SHREC20 (25)	0	5	0	<b>25</b>
DT4D-M (60)	3	1	4	<b>21</b>

Table 3: The **number of matching instances with certified global optimality** within 1h time budget. Numbers in parentheses in the 1st column are the total number of matching pairs. SM-comb [38] does not allow to set a time budget, so we let it run until it terminates.

but many methods struggle with such large scale changes. We randomly select five SHREC20 [11] pairs to further study the effect of shape scale by rescaling one shape while fixing the other. More specifically, we fix the shape  $\mathcal{X}$  and rescale shape  $\mathcal{Y}$  by a factor of  $\{0.1, 0.5, 1, 5, 10\}$  respectively, to create 5 instances with different scales of the same matching pair. As shown in Lemma 2 and Fig. 5 (right), our proposed method is scale-invariant and solves all matching instances to global optimality independent of the shape scale, hence consistently achieving the lowest *mean geodesic error*. While PMSDP [25] is also scale-invariant, it fails under non-isometric deformation. All other baselines depend on shape scale and thus manifest an accuracy reduction upon scale changes.

#### 4.7. Mesh Resolution

The mesh resolution of shapes often plays an vital role in the scalability of matching methods, especially for global methods. The elastic matching model [48] targeted by SM-comb [38] is constructed in the space of product-surfaces, and therefore the number of binary variables to be optimised increases quadratically with the number of mesh triangles. MINA [6] deforms each mesh triangle using an affine transformation, and in addition requires a global rotation matrix that drastically increases its runtime due to the involved discretisation. In our approach the continuous variables  $\hat{\mathbf{X}}, \hat{\mathbf{Y}}$  increase linearly with the number of mesh vertices while the number of binary variables  $\mathbf{P}$  remains unchanged. As we demonstrate in Fig. 1d, our approach scales much better compared to SM-comb and MINA, which stems from our more efficient matching formalism (e.g. without rotation matrices opposed to MINA, or without a quadratic number of binary variables opposed to SM-comb).

### 5. Discussion & Limitations

Despite the state-of-the-art performance, SIGMA has also limitations. We show some failure cases due to inconsistent orientation maps in Fig. 7. As proof-of-concept, we showed in Fig. 1b that SIGMA can also work for partial shapes. However, its performance is far from perfect under




Figure 7: (Left) **Mesh to point cloud matching**. Shapes in the top row are matched to respective shapes in the bottom row. Our method is able to tackle this challenging scenario. (Right) **failure modes**. For few shapes our orientation term does not help to disentangle intrinsic symmetry.

this challenging scenario and often cannot find the global optimum within an adequate time limit. This is due to the fact that the equality constraints on  $\mathbf{P}$  become the inequalities  $\mathbf{P}^T \mathbf{1} \leq \mathbf{1}, \mathbf{P} \mathbf{1} \leq \mathbf{1}$  (cf. Eqn. (4)) under partiality, and this increases the size of the solution search space, resulting in longer runtime. We can handle mesh to point cloud matching quite well, as shown in Fig. 7. Here, we replace the geodesic distance with Euclidean distances for the pruning strategy. Another challenge are matchings between pairs with topological changes. Here, SIGMA struggles since the mesh of one shape could not well-explain the deformation into the other shape anymore. We leave these challenging problem settings for future exploration.

### 6. Conclusion

We presented SIGMA, an initialisation-free MIP formulation for sparse shape matching problems, for which we demonstrate that it can be solved to global optimality for many instances. Our method is provably invariant to global scaling and rigid transformations of the input shapes, eliminating the need for an extrinsic shape (pre-)alignment required by many shape matching methods. Furthermore, we introduced the projected Laplace-Beltrami operator PLBO, which combines both intrinsic and extrinsic geometric information and remains invariant under  $E(3)$  group actions – we believe it may also be useful for a broad range of other geometry processing tasks. In our experiments, we demonstrated that our method outperforms competitive baselines in terms of matching quality, optimality and scalability, across a spectrum of challenging 3D shape benchmarks.

**Acknowledgement.** MG, ME, DC are supported by the ERC Advanced Grant SIMULACRON, by the Munich Center for Machine Learning and by the EPSRC Programme Grant VisualAI EP/T028572/1. ZL is supported by the Ministry of Culture and Science NRW.



## References

- [1] Yonathan Aflalo, Ron Kimmel, and Dan Raviv. Scale Invariant Geometry for Nonrigid Shapes. *SIAM Journal on Imaging Sciences*, 6(3):1579–1597, 2013. 3
- [2] MOSEK ApS. *The MOSEK Optimization Toolbox for MATLAB Manual.*, 2019. 5
- [3] Souhaib Attaiki, Gautam Pai, and Maks Ovsjanikov. DPFM: Deep Partial Functional Maps. In *3DV*, 2021. 1
- [4] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The Wave Kernel Signature: A Quantum Mechanical Approach to Shape Analysis. In *ICCV Workshop*, 2011. 3, 5
- [5] Florian Bernard, Frank Schmidt, Johan Thunberg, and Daniel Cremers. A Combinatorial Solution to Non-Rigid 3D Shape-to-Image Matching. In *CVPR*, 2017. 1
- [6] Florian Bernard, Zeeshan Khan Suri, and Christian Theobalt. MINA: Convex Mixed-Integer Programming for Non-Rigid Shape Alignment. In *CVPR*, 2020. 1, 2, 3, 5, 6, 7, 8
- [7] Florian Bernard, Christian Theobalt, and Michael Moeller. DS\*: Tighter Lifting-Free Convex Relaxations for Quadratic Matching Problems. In *CVPR*, 2018. 2
- [8] Mario Botsch, Mark Pauly, Markus Gross, and Leif Kobbelt. PriMo: Coupled Prisms for Intuitive Surface Modeling. In *SGP*, pages 11–20, 2006. 4
- [9] Alexander Bronstein, Michael Bronstein, and Ron Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer Publishing Company, Incorporated, 1 edition, 2008. 5, 6, 7
- [10] Michael Bronstein and Iasonas Kokkinos. Scale-Invariant Heat Kernel Signatures for Non-Rigid Shape Recognition. In *CVPR*, pages 1704–1711, 2010. 3
- [11] Roberto Dyke, Yu-Kun Lai, Paul Rosin, Stefano Zappalà, Seana Dykes, Daoliang Guo, Kun Li, Riccardo Marin, Simone Melzi, and Jingyu Yang. SHREC’20: Shape Correspondence with Non-Isometric Deformations. *Computers & Graphics*, 92:28–43, 2020. 5, 6, 7, 8
- [12] 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
- [13] Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep Shells: Unsupervised Shape Correspondence with Optimal Transport. *Neurips*, 2020. 1
- [14] Danielle Ezuz, Behrend Heeren, Omri Azencot, Martin Rumpf, and Mirela Ben-Chen. Elastic Correspondence between Triangle Meshes. *CGF*, 2019. 3
- [15] Pedro Felzenszwalb. Representation and Detection of Deformable Shapes. *TPAMI*, 27(2):208–220, 2005. 2
- [16] Eitan Grinspun, Anil Hirani, Mathieu Desbrun, and Peter Schröder. Discrete Shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 62–67, 2003. 4
- [17] Itay Kezurer, Shahar Kovalsky, Ronen Basri, and Yaron Lipman. Tight Relaxation of Quadratic Matching. *Comput. Graph. Forum*, 2015. 2
- [18] Vladimir G. Kim, Yaron Lipman, and Thomas Funkhouser. Blended Intrinsic Maps. *ACM Transactions on Graphics (TOG)*, 30(4), 2011. 5, 6
- [19] Yam Kushinsky, Haggai Maron, Nadav Dym, and Yaron Lipman. Sinkhorn Algorithm for Lifted Assignment Problems. *SIAM Journal on Imaging Sciences*, 12(2):716–735, 2019. 2
- [20] Zorah Löhner, Emanuele Rodola, Frank Schmidt, Michael Bronstein, and Daniel Cremers. Efficient Globally Optimal 2D-to-3D Deformable Shape Matching. In *CVPR*, 2016. 2
- [21] Eugene Lawler. The Quadratic Assignment Problem. *Management science*, 9(4):586–599, 1963. 2
- [22] D. Khuê Lê-Huu and Nikos Paragios. Alternating Direction Graph Matching. In *CVPR*, 2017. 2
- [23] 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
- [24] Robin Magnet, Jing Ren, Olga Sorkine-Hornung, and Maks Ovsjanikov. Smooth Non-Rigid Shape Matching via Effective Dirichlet Energy Optimization. In *3DV*, 2022. 5, 6, 7
- [25] Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, and Yaron Lipman. Point Registration via Efficient Convex Relaxation. *ACM Transactions on Graphics (TOG)*, 35(4):73, 2016. 1, 2, 3, 5, 6, 7, 8
- [26] Ahmad Nasikun, Christopher Brandt, and Klaus Hildebrandt. Fast Approximation of Laplace-Beltrami Eigenproblems. *Computer Graphics Forum*, 37(5), 2018.
- [27] 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):30, 2012. 2, 3
- [28] Gautam Pai, Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Fast Sinkhorn Filters: Using Matrix Scaling for Non-Rigid Shape Correspondence With Functional Maps. In *CVPR*, pages 384–393, 2021. 3
- [29] Mikhail Panine, Maxime Kirgo, and Maks Ovsjanikov. Non-Isometric Shape Matching via Functional Maps on Landmark-Adapted Bases. *Computer Graphics Forum (CGF)*, 41(6), 2022. 3
- [30] Panos Pardalos, Franz Rendl, and Henry Wolkowicz. The Quadratic Assignment Problem - A Survey and Recent Developments. *DIMACS Series in Discrete Mathematics*, 1993. 2
- [31] Ulrich Pinkall and Konrad Polthier. Computing Discrete Minimal Surfaces and their Conjugates. *Experimental Mathematics*, 1(2), 1993. 4
- [32] Jing Ren, Simone Melzi, Maks Ovsjanikov, and Peter Wonka. MapTree: Recovering Multiple Solutions in the Space of Maps. *ACM Transactions on Graphics (TOG)*, 2020. 3
- [33] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and Orientation-Preserving Correspondences via Functional Maps. In *SIGGRAPH Asia*, 2018. 5
- [34] Jing Ren, Peter Wonka, Gowtham Harihara, and Maks Ovsjanikov. Geometric Analysis of Shape Variability of Lower Jaws of Prehistoric Humans. *L’Anthropologie*, 2020. 1
- [35] F Rendl, P Pardalos, and H Wolkowicz. The Quadratic Assignment Problem: A Survey and Recent Developments. In *DIMACS workshop*, 1994. 2

- [36] Emanuele Rodola, Alex Bronstein, Andrea Albarelli, Filippo Bergamasco, and Andrea Torsello. A Game-Theoretic Approach to Deformable Shape Matching. In *CVPR*, 2012. [2](#), [3](#)
- [37] Paul Roetzer, Zorah Löhner, and Florian Bernard. Conjugate Product Graphs for Globally Optimal 2D-3D Shape Matching. In *CVPR*, 2023. [2](#)
- [38] Paul Roetzer, Paul Swoboda, Daniel Cremers, and Florian Bernard. A Scalable Combinatorial Solver for Elastic Geometrically Consistent 3D Shape Matching. In *CVPR*, 2022. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [39] Yusuf Sahillioğlu. A Genetic Isometric Shape Correspondence Algorithm with Adaptive Sampling. *ACM Transactions on Graphics (TOG)*, 36, 2018. [3](#)
- [40] Yusuf Sahillioğlu. Recent Advances in Shape Correspondence. *The Visual Computer*, 36(1), 2020. [2](#)
- [41] Christian Schellewald and Christoph Schnörr. Probabilistic Subgraph Matching Based on Convex Relaxation. In *EMM-CVPR*, 2005. [2](#)
- [42] Frank Schmidt, Eno Töppe, and Daniel Cremers. Efficient Planar Graph Cuts with Applications in Computer Vision. In *CVPR*, 2009. [2](#)
- [43] Thomas Schoenemann and Daniel Cremers. A Combinatorial Solution for Model-Based Image Segmentation and Real-time Tracking. *TPAMI*, 32(7):1153–1164, 2009. [2](#)
- [44] Justin Solomon, Mirela Ben-Chen, Adrian Butscher, and Leonidas Guibas. As-Killing-As-Possible Vector Fields for Planar Deformation. In *Computer Graphics Forum*, 2011. [4](#)
- [45] Olga Sorkine and Marc Alexa. As-Rigid-As-Possible Surface Modeling. In *SGP*, 2007. [4](#)
- [46] Jonathan Taylor, Jamie Shotton, Toby Sharp, and Andrew Fitzgibbon. The Vitruvian Manifold: Inferring Dense Correspondences for One-Shot Human Pose Estimation. In *CVPR*, pages 103–110, 2012. [1](#)
- [47] Matthias Vestner, Zorah Löhner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodola, Alex Bronstein, Michael Bronstein, Ron Kimmel, and Daniel Cremers. Efficient Deformable Shape Correspondence via Kernel Matching. In *3DV*, 2017. [2](#), [3](#)
- [48] Thomas Windheuser, Ulrich Schlickewei, Frank Schmidt, and Daniel Cremers. Geometrically Consistent Elastic Matching of 3D Shapes: A Linear Programming Solution. In *ICCV*, 2011. [1](#), [2](#), [5](#), [8](#)
- [49] Zhijie Wu, Xiang Wang, Di Lin, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. SAGNet: Structure-Aware Generative Network for 3D-Shape Modeling. *ACM Transactions on Graphics (TOG)*, 2019. [1](#)
- [50] Jingyu Yang, Ke Li, Kun Li, and Yu-Kun Lai. Sparse Non-rigid Registration of 3D Shapes. In *SGP*, 2015. [3](#)
- [51] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D Menagerie: Modeling the 3D Shape and Pose of Animals. In *CVPR*, 2017. [5](#), [6](#), [7](#)