


Conjugate Product Graphs for Globally Optimal 2D-3D Shape Matching

Paul Roetzer¹ Zorah Löhner² Florian Bernard¹
 University of Bonn¹ University of Siegen²



Results of Löhner *et al.* [28] (top) and ours (bottom) on the TOSCA dataset. (i) Matching with our approach (ii) 2D to 3D deformation transfer

Figure 1. We propose a novel formalism for **globally optimal 2D contour to 3D shape matching** based on shortest paths in the *conjugate product graph*. For the first time we make it possible to incorporate higher-order costs within a shortest path-based matching formalism, which in turn enables to integrate powerful priors, *e.g.* favouring locally rigid deformations. **Left:** Our method produces compelling 2D-3D matchings that significantly outperform the previous state of the art [28]. **Right:** Sketch-based 2D to 3D deformation transfer by (i) computing a 2D-3D matching using our approach, (ii) manipulating the 2D sketch, and then transferring 2D deformations to the 3D shape.

Abstract

We consider the problem of finding a continuous and non-rigid matching between a 2D contour and a 3D mesh. While such problems can be solved to global optimality by finding a shortest path in the product graph between both shapes, existing solutions heavily rely on unrealistic prior assumptions to avoid degenerate solutions (*e.g.* knowledge to which region of the 3D shape each point of the 2D contour is matched). To address this, we propose a novel 2D-3D shape matching formalism based on the conjugate product graph between the 2D contour and the 3D shape. Doing so allows us for the first time to consider higher-order costs, *i.e.* defined for edge chains, as opposed to costs defined for single edges. This offers substantially more flexibility, which we utilise to incorporate a local rigidity prior. By doing so, we effectively circumvent degenerate solutions and thereby obtain smoother and more realistic matchings, even when using only a one-dimensional feature descriptor. Overall, our method finds globally optimal and continuous 2D-3D matchings, has the same asymptotic complexity as previous solutions, produces state-of-the-art results for shape matching and is even capable of matching partial shapes. Our code is publicly available.¹

¹<https://github.com/paul0noah/sm-2D3D>

1. Introduction

In recent years the computer vision community has put great effort into the matching of either two 2D or two 3D shapes. However, the task of matching a 2D shape to a 3D shape is a problem that has received less attention, even though it has a high practical relevance due to its wide variety of applications. For example, 2D-3D shape matching has the potential to bridge the gap between the 2D and 3D domain by making interaction with 3D objects more accessible to non-experts, who typically find the manipulation of 2D shapes more intuitive than operating 3D modelling tools. In addition to the modelling and manipulation of 3D shapes using 2D sketches (see Fig. 1), 2D-3D shape matching is relevant for 3D shape retrieval from 2D queries (*e.g.* searching a hand-drawn sketch in a 3D shape database), for augmented reality interactions (*e.g.* using hand-gestures to select 3D objects), for 3D image analysis based on 2D images (*e.g.* matching 2D X-ray image segmentations to 3D CT image segmentations), or for multi-modal 2D-3D shape analysis.

2D-3D shape matching can be framed as finding a continuous mapping of a 2D contour (*e.g.* a sketch of an animal outline) to a 3D shape (*e.g.* a 3D model of this animal), see Fig. 1 (left). Here, the matched 2D contour that is deformed to the 3D shape should resemble the original

2D shape as much as possible, *i.e.* spatial shape deformations should be small. While humans have an intuitive and implicit understanding of *good* 2D-3D matchings, unfortunately, it is non-trivial to transfer this understanding into a rigorous mathematical framework: left-right flips are not distinguishable; the 2D shape does not contain all parts of the 3D shape (*e.g.* 2D shape of the wolf in Fig. 1 (left) contains only two legs); usually there is more than one *good* solution; and even slight deviations from a *good* matching can either be another good matching or can be a bad matching (*e.g.* zig-zagging on the 3D shape). In addition, phrasing 2D-3D shape matching as an optimisation problem requires to compute features on both shapes that allow to distinguish corresponding points from non-corresponding points – this is particularly difficult as many of the widely-used features for 2D or 3D shapes do not have a natural counterpart in the other domain, and are thus not directly comparable.

Nevertheless, previous work shows that matching a 2D contour to a 3D shape can be efficiently and optimally solved based on shortest paths in product graphs [28]. However, existing solutions require strong, unnatural assumptions (*e.g.* a coarse pre-matching, see Sec. 3.2) in order to resolve (some of) the above-mentioned difficulties. In this paper we present a novel graph-based formalism that relaxes previous unnatural assumptions, which in turn allows to solve a substantially more difficult setting of 2D-3D shape matching. Our main contributions are:

- We present a novel matching formalism based on conjugate product graphs that allow to encode more expressive higher-order information.
- For the first time this makes it possible to impose a local rigidity prior that penalises deformations, which in turn leads to previously unseen matching quality.
- Opposed to involved high-dimensional feature descriptors that were previously used (*e.g.* spectral features), our method requires only a simple one-dimensional feature that encodes the local object thickness – a feature that can naturally be defined for 2D and 3D shapes.
- Overall, our technical contributions allow us to solve 2D-3D matching for the first time without the requirement of a coarse pre-matching.

2. Related Work

In the following we summarise existing works that we consider most relevant in the context of this paper.

Geometric Feature Descriptors. Most matching approaches rely on point-wise features to decide what are good potential matches. For 2D contours, popular features are cumulative angles [50], curvature [27, 49] and various distance metrics [29, 33, 42, 49, 50]. One (for our work) notable example from the class of distance-based metrics is to consider the distance from each point to other parts of the contour along several fixed rays [33].

On 3D shapes, other feature types are predominant because the geometry is more complicated and 2D features often do not have a direct equivalent in 3D. While curvature does exist in 2D and in 3D, in 3D there are multiple notions of curvature, like mean, Gaussian and directional curvature. More popular are higher-dimensional features like the heat kernel signature [46] or wave kernel signature [1], which are based on spectral properties of the 3D surface, or the SHOT descriptor [48] based on the distribution of normals in the neighbourhood of a vertex. Recent approaches aim to learn suitable features for a specific matching pipeline [30, 31]. Overall, there is a discrepancy between 2D and 3D features, and even for features that can conceptually be calculated for both domains, they are typically not directly comparable. While [28] successfully addresses 2D-3D *shape retrieval* based on spectral 2D and 3D features, our experiments confirm that these features are insufficient to achieve precise *correspondences*. Similarly, many approaches learn multi-modal or multi-dimensional descriptors for entire shapes [17, 37, 51], but these are only useful in retrieval settings and not capable of point-to-point comparisons needed for finding reliable correspondences. In this paper we instead shift focus to incorporating a powerful deformation prior, so that in turn substantially simpler feature descriptors are sufficient. We demonstrate that this allows to consider simple distance-based features which can be consistently computed both in 2D and 3D.

2D-3D Matching. Matching pairs of 2D objects is well-researched and it is widely known that respective solutions can be represented as paths in a graph. With that, shortest path algorithms can be used to efficiently find globally optimal solutions. This has for example been done for open contours, known as dynamic time warping [41], and closed contours [42], including invariance to scale and partiality [33]. Similarly, it was shown that matching 2D contours to 2D images (*e.g.* for template-based image segmentation) can be addressed using a similar framework [43]. Matching two 3D shapes is considerably harder as the solution is not a shortest path (in the product manifold) anymore but rather a minimal surface embedded in four-dimensional space. Thus, imposing constraints on the continuity of the solution is not possible in an efficient way [38, 53, 54]. We note that there are works which consider (geodesic) shortest paths in the space of 3D shapes, but they are not solvable to global optimality [14, 21]. While from an algorithmic perspective finding a 2D-3D matching is easier than the 3D-3D case (as the former also amounts to a (cyclic) shortest path problem [28]), quantifying matching costs is significantly more difficult for the 2D-3D case (*cf.* previous paragraph on feature descriptors). In this work we build upon the path-based 2D-3D matching formalism of [28] and propose a novel formalism that enables the use of higher-order costs (defined for chains of edges, opposed to costs of single

edges). In turn, our formalism allows for the first time to incorporate a spatial deformation prior, so that our framework requires substantially less descriptive features – in fact, we demonstrate that a one-dimensional distance-based feature descriptor is sufficient to successfully solve 2D-3D shape matching.

Extensions of Graph Representations. Graphs are not only relevant for diverse subfields of visual computing, such as e.g. image analysis [5, 7, 13, 24, 25, 36, 39, 40, 44], recognition [9, 12], tracking [19, 56], or mesh processing [4, 23, 28, 34], but also for a wide variety of other application domains, for example in DNA research [11], language processing [52], or social sciences [22]. In graph theory, many graph extensions have been proposed, including multilayer networks [26], dual graphs [16], hypergraphs [3], and product graphs [15], to name just a few. The product graph extends the concept of the Cartesian product (and other types of products) to graphs by additionally encoding neighbourhood information. This has been used in the context of different matching problems, including 2D-2D [42], 3D-3D [53], and 2D-3D [28] settings. Another graph extension relevant for this paper is the *conjugate graph* (also known as *line graph*), which encodes connectivity information into the vertices instead of edges [16]. This has for example been used for route planning [55] or graph link prediction [8]. In this work, we propose to combine product graphs with conjugate graphs (in fact we consider the conjugate graph of a product graph between two shapes) and showcase that this substantially increases modelling expressiveness and flexibility, and therefore allows for globally optimal 2D-3D shape matching.

3. Background & Notation

In this section we introduce our notation (also see Tab. 1), conjugate graphs, and the formalism for the matching of shapes as shortest path problem on a product graph.

Definition 1 (Directed Graph). A directed graph \mathcal{G} is defined as a tuple $(\mathcal{V}_{\mathcal{G}}, \mathcal{E}_{\mathcal{G}})$ of vertices $\mathcal{V}_{\mathcal{G}}$ and oriented edges $\mathcal{E}_{\mathcal{G}} \subset \mathcal{V}_{\mathcal{G}} \times \mathcal{V}_{\mathcal{G}}$ (i.e. oriented edge $(v_1, v_2) \in \mathcal{E}_{\mathcal{G}}$ does not imply $(v_2, v_1) \in \mathcal{E}_{\mathcal{G}}$).

We directly work with discrete graph-based representations for shapes, i.e. 2D shapes are represented as contours sampled at m many points, and 3D shapes are represented as (manifold) triangular surface meshes:

Definition 2 (2D Shape). We define a 2D shape (or contour) \mathcal{M} as a tuple $(\mathcal{V}_{\mathcal{M}}, \mathcal{E}_{\mathcal{M}})$ of m vertices $\mathcal{V}_{\mathcal{M}}$ and m oriented edges $\mathcal{E}_{\mathcal{M}} \subset \mathcal{V}_{\mathcal{M}} \times \mathcal{V}_{\mathcal{M}}$ s.t. \mathcal{M} is a directed cycle graph.

Definition 3 (3D Shape). We define a 3D shape \mathcal{N} as a tuple $(\mathcal{V}_{\mathcal{N}}, \mathcal{E}_{\mathcal{N}})$ of vertices $\mathcal{V}_{\mathcal{N}}$ and oriented edges $\mathcal{E}_{\mathcal{N}} \subset \mathcal{V}_{\mathcal{N}} \times \mathcal{V}_{\mathcal{N}}$ such that \mathcal{N} forms a 2D manifold in 3D space (triangular surface mesh, possibly with multiple boundaries).

We also consider an extended edge set, which contains all vertices represented as self-edges:

Definition 4 (Extended Set of Edges). For a (2D or 3D) shape $\mathcal{X} = (\mathcal{V}_{\mathcal{X}}, \mathcal{E}_{\mathcal{X}})$, we define the extended set of edges $\mathcal{E}_{\mathcal{X}}^+ = \mathcal{E}_{\mathcal{X}} \cup \{(a, a) \mid a \in \mathcal{V}_{\mathcal{X}}\} \subset \mathcal{V}_{\mathcal{X}} \times \mathcal{V}_{\mathcal{X}}$. We call the additional edges degenerate edges.

Symbol	Description
$\mathcal{M} = (\mathcal{V}_{\mathcal{M}}, \mathcal{E}_{\mathcal{M}})$	2D shape (closed contour)
$\mathcal{N} = (\mathcal{V}_{\mathcal{N}}, \mathcal{E}_{\mathcal{N}})$	3D shape (manifold triangular surface mesh)
$e^{\mathcal{M}}, e^{\mathcal{N}}$	edge $e^{\mathcal{M}}$ of contour, edge $e^{\mathcal{N}}$ of mesh
$\mathcal{E}_{\mathcal{X}}^+$	extended set of edges of shape $\mathcal{X} \in \{\mathcal{M}, \mathcal{N}\}$
$\mathcal{P} = (\mathcal{V}, \mathcal{E})$	product graph of $\mathcal{M} \times \mathcal{N}$ with product vertices \mathcal{V} and product edges \mathcal{E}
e, v	edge e , vertex v of \mathcal{P}
$\mathcal{P}^* = (\mathcal{V}^*, \mathcal{E}^*)$	conjugate product graph of \mathcal{P} with vertices \mathcal{V}^* and edges \mathcal{E}^*
e^*, v^*	edge e^* , vertex v^* of \mathcal{P}^*

Table 1. Summary of the notation used in this paper.

3.1. Conjugate Graphs

Definition 5 (Conjugate Graph [16]). The conjugate graph \mathcal{G}^* of a directed graph \mathcal{G} is defined as a tuple $(\mathcal{V}_{\mathcal{G}}^*, \mathcal{E}_{\mathcal{G}}^*)$ with

$$\mathcal{V}_{\mathcal{G}}^* = \mathcal{E}_{\mathcal{G}}, \quad \mathcal{E}_{\mathcal{G}}^* = \{((v_1, v_2), (v_2, v_3)) \in \mathcal{V}_{\mathcal{G}}^* \times \mathcal{V}_{\mathcal{G}}^*\}.$$

Intuitively, the edges of \mathcal{G} become the vertices of \mathcal{G}^* and the conjugate edges connect pairs of adjacent edges from \mathcal{G} . Fig. 2 illustrates the construction of the conjugate graph.




Figure 2. Illustration of the **conjugate graph**. (a) Input graph \mathcal{G} . (b) Each edge of \mathcal{G} becomes a conjugate vertex (●). (c) The conjugate graph \mathcal{G}^* is now formed by connecting the newly introduced vertices (●) by edges (→) according to Definition 5 (e.g. conjugate vertex $(1, 3)$ and $(3, 4)$ are connected by a directed edge since they are both adjacent to vertex 3 in \mathcal{G}).

3.2. Matching Formalism

Next, we introduce the product graph \mathcal{G} between a 2D contour and a 3D shape, and we summarise Löhner *et al.*'s [28] representation of a 2D-3D shape matching as shortest (cyclic) path in the product graph.

Definition 6 (Product Graph). *The product graph \mathcal{P} of the 2D contour \mathcal{M} and the 3D shape \mathcal{N} is a tuple $(\mathcal{V}_{\mathcal{P}}, \mathcal{E}_{\mathcal{P}})$ of product vertices $\mathcal{V}_{\mathcal{P}}$ and product edges $\mathcal{E}_{\mathcal{P}}$, where*

$$\begin{aligned} \mathcal{V}_{\mathcal{P}} &= \mathcal{V}_{\mathcal{M}} \times \mathcal{V}_{\mathcal{N}}, \\ \mathcal{E}_{\mathcal{P}} &= \{(e_1^{\mathcal{M}}, e_2^{\mathcal{N}}) \in \mathcal{V}_{\mathcal{P}} \times \mathcal{V}_{\mathcal{P}} \mid e_1^{\mathcal{M}} \in \mathcal{E}_{\mathcal{M}}^+, e_2^{\mathcal{N}} \in \mathcal{E}_{\mathcal{N}}^+, \\ &\quad e_1^{\mathcal{M}} \text{ or } e_2^{\mathcal{N}} \text{ non-deg.}\}. \end{aligned}$$

The product graph \mathcal{P} is visualised in Fig. 3 (left). To simplify the notation, we will refer to the vertices and edges of the product graph as \mathcal{V} and \mathcal{E} for the remainder of the paper. A matching between the 2D contour and the 3D shape can be represented as the subset $\mathcal{C} \subset \mathcal{V}_{\mathcal{M}} \times \mathcal{V}_{\mathcal{N}}$ where tuples in \mathcal{C} indicate which vertices of the 2D contour and 3D shape are in correspondence. Desirable properties of such matchings are that a) each vertex on \mathcal{M} is matched to at least one vertex on \mathcal{N} , and b) the matching is continuous, *i.e.* if two vertices on \mathcal{M} are adjacent, their matches on \mathcal{N} should also be adjacent. These properties can be guaranteed if the solution is a (cyclic) path that goes through all layers of the product graph (cf. Fig. 3). A path that minimises costs defined on the (product graph) edges can efficiently be computed based on Dijkstra’s algorithm [10]. To ensure the path is cyclic, Dijkstra’s algorithm needs to be run multiple times (once for each vertex of the 3D mesh); however, the number of Dijkstra runs can be drastically reduced based on a simple branch and bound strategy (see [28] and Appendix).

Despite the theoretical elegance of Lähler *et al.*’s formalism, a major limitation is that shortest paths only take into account costs of individual edges. With that, the approach is not capable of penalising local deformations induced by a matching (which is only possible with *pairs* of product graph edges). Instead, the authors use high-dimensional features in combination with knowledge about pre-matched segmentations between 2D and 3D shapes. While such a pre-matching drastically reduces the search space and avoids many degenerate solutions, the knowledge of such a pre-matching is typically not available in practice.

4. Our 2D-3D Shape Matching Approach

In the following, we present our solution for 2D-3D shape matching which allows to incorporate higher-order edge costs which in turn is enabled through conjugate product graphs. With this, spatial deformations can be naturally penalised within a shortest path matching framework while still finding a globally optimal matching in polynomial time (with the same asymptotic complexity as [28]), see Section 3.2. We emphasise that we address a significantly more difficult problem setting than [28] since we do not rely on the unrealistic assumption that a coarse 2D-3D pre-matching is available.

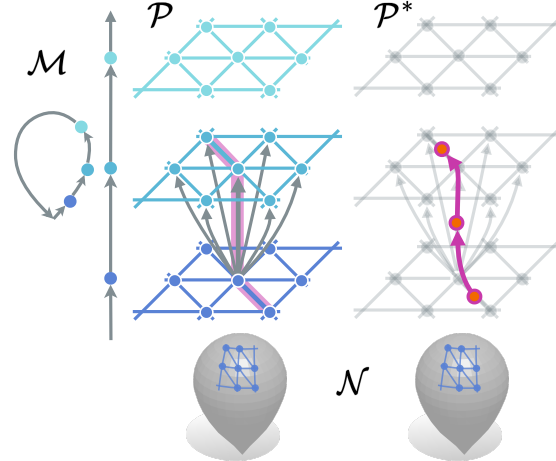


Figure 3. Illustration of the product graph \mathcal{P} (left) and the conjugate product graph \mathcal{P}^* (right) for a water drop shape. **Left:** the product graph \mathcal{P} is structured into layers ($\bullet, \bullet, \bullet$), where each layer represents a single vertex on \mathcal{M} and the entire 3D shape \mathcal{N} . **Right:** we illustrate (part of) the conjugate product graph \mathcal{P}^* for the three-edge path $\bullet-\bullet-\bullet$ in \mathcal{P} (highlighted in pink), which becomes $\bullet-\bullet-\bullet$ in \mathcal{P}^* . Conjugate product vertices \bullet are shown in orange and conjugate product edges --- are shown in pink.




Figure 4. Hierarchical relationship between edges of the involved graphs. An edge $e^* \in \mathcal{E}^*$ of the conjugate product graph \mathcal{P}^* is formed by two edges $e_1, e_2 \in \mathcal{E}$ of the product graph \mathcal{P} (which correspond to vertices $v_1^*, v_2^* \in \mathcal{V}^*$ in the conjugate product graph \mathcal{P}^* , respectively). Each edge e_\bullet of \mathcal{P} is formed by one edge $e_\bullet^{\mathcal{M}} \in \mathcal{E}_{\mathcal{M}}^+$ and one edge $e_\bullet^{\mathcal{N}} \in \mathcal{E}_{\mathcal{N}}^+$ of the shapes \mathcal{M} and \mathcal{N} , respectively.

4.1. Conjugate Product Graphs

Our formalism builds upon conjugate product graphs, *i.e.* the conjugate graph (Definition 5) of a product graph (Definition 6). We refer to the conjugate product graph as $\mathcal{P}^* = (\mathcal{V}^*, \mathcal{E}^*)$. Here, edges in the product graph become vertices in the conjugate product graph and are connected based on the adjacency of vertices in the product graph, see Definition 5. Thus, an edge $e^* \in \mathcal{E}^*$ in \mathcal{P}^* has the scope of two edges in the product graph \mathcal{P} , *i.e.* $e^* = (e_1, e_2)$, $e_1, e_2 \in \mathcal{E}$, see Fig. 4. In turn, this enables the definition of cost functions that consider two product graph edges simultaneously which allows to integrate powerful priors, *i.e.* regularise the problem so that locally rigid deformations are favoured and thus resulting matchings are smoother and more realistic (cf. Fig. 1 and Fig. 8). We note that higher-order costs can also be defined by repeating the conjugation process, *e.g.* an edge in the conjugate

of the conjugate product graph is formed by three edges of the product graph so that costs can be defined for triplets of product graph edges (and so on). For brevity and a simpler exposure, in the following we restrict ourselves w.l.o.g. to second-order costs.

4.2. Cost Function

We define our cost function $d : \mathcal{E}^* \rightarrow \mathbb{R}$ for every edge $e^* = (v_1^*, v_2^*) = (e_1, e_2) \in \mathcal{E}^*$ in the conjugate product graph as

$$d(e^*) = d_{\text{data}}(e^*) + d_{\text{reg}}(e^*). \quad (1)$$

d_{data} is the data term which measures the similarity between the product edges based on feature descriptors. d_{reg} is a local rigidity regulariser which ensures that adjacent edges on the 2D contour are deformed similarly to adjacent elements on the 3D shape. We first describe the data term followed by the local rigidity regulariser.

Data Term. A major difficulty when comparing 2D and 3D shapes is that many of the existing geometric feature descriptors cannot be consistently defined for 2D and 3D shapes (e.g. although the notion of curvature exists for both shapes, in 3D the curvature is direction-dependent, which makes it difficult to compare 2D and 3D curvature). We build a simple one-dimensional descriptor based on the observation that corresponding points $i \in \mathcal{M}$ and $j \in \mathcal{N}$ of the same shape class should have a similar distance to the other side going through the interior of the respective shape, see Fig. 5a. As such, we consider *local thickness* as feature descriptor. It is computed as follows:

- For a vertex $i \in \mathcal{M}$, we compute the *2D local thickness* ℓ_i^{2D} by following the inverted vertex-normal (through the interior of the shape) until the (first) intersection with the contour [33].
- For a vertex $j \in \mathcal{N}$, we compute the *3D local thickness* ℓ_j^{3D} by following the inverted vertex-normal (through the interior of the shape) until the (first) intersection with the mesh. We employ a triangle-ray-intersection algorithm for this [35].

With that, we define the local thickness difference for each conjugate product edge $e^* = (e_1, e_2) \in \mathcal{E}^*$, so that our data term $d_{\text{data}}(e^*)$ reads

$$d_{\text{data}}(e^*) = \psi_1(|\ell_i^{2D} - \ell_j^{3D}|), \quad (2)$$

where ℓ_i^{2D} and ℓ_j^{3D} are local thickness values at vertices $i \in \mathcal{V}_{\mathcal{M}}$ and $j \in \mathcal{V}_{\mathcal{N}}$ on the 2D and 3D shape, respectively. Large deformations of shapes potentially lead to large, local outliers of respective local thickness. To reduce the influence of such outliers, we additionally apply the function $\psi_1(\cdot)$ to the absolute value of the local thickness difference. For example, $\psi_1(\cdot)$ can be chosen to be a robust loss function. We avoid taking into account the same local thickness value multiple times by computing the local thickness




Figure 5. (a) The **local thickness** ℓ_p for the point p is found by intersecting the ray from p in the opposite normal direction (light red) with the shape. (b) Illustration of finding the **rotation** R_e that aligns the 3D coordinate frame defined for a 2D contour edge $e^{\mathcal{M}}$ and the 3D coordinate frame defined for a 3D shape edge $e^{\mathcal{N}}$. The black vector shows the shape edge, the red vector the normal and the green vector their cross product.

difference $d_{\text{data}}(e^*)$ at a conjugate product vertex e^* solely with the local thickness at vertex i shared by $e_1^{\mathcal{M}}$ and $e_2^{\mathcal{M}}$, and respectively the local thickness at vertex j shared by $e_1^{\mathcal{N}}$ and $e_2^{\mathcal{N}}$. We have found that despite its simplicity, local thickness is an effective one-dimensional feature descriptor that, in combination with our local rigidity regulariser, enables faithful 2D-3D shape matchings, see Sec. 5.

Regularisation. Inspired by [4, 45], we employ a regularisation term which favours deformations that are locally rigid. To compute the regularisation of the conjugate product edge $e^* = (e_1, e_2) = ((e_1^{\mathcal{M}}, e_1^{\mathcal{N}}), (e_2^{\mathcal{M}}, e_2^{\mathcal{N}}))$ (cf. Fig. 4), we define a local 3D coordinate frame for each of the four (shape) edges $e_1^{\mathcal{M}}, e_1^{\mathcal{N}}, e_2^{\mathcal{M}}, e_2^{\mathcal{N}}$. To this end, we embed the 2D contour into 3D space by adding a third constant coordinate. With that, for both 2D contour and 3D shape, we can define a local 3D coordinate frame based on the normalised edge direction, outward-pointing unit normal, and their cross product. Subsequently, we solve the orthogonal Procrustes problem [47] in order to compute the rotation R_{e_1} that aligns the 3D coordinate frame of $e_1^{\mathcal{M}}$ to the 3D coordinate frame of $e_1^{\mathcal{N}}$, and the rotation R_{e_2} that aligns the 3D coordinate frame of $e_2^{\mathcal{M}}$ to the coordinate of $e_2^{\mathcal{N}}$, see Fig. 5b. In presence of degenerate edges we simply use the previous edge, see also Sec. 4.3.

By computing the geodesic distance between R_{e_1} and R_{e_2} on the Lie group $\text{SO}(3)$, we can quantify the amount of non-rigidity of the matching that is induced by e^* . For computational efficiency, we consider unit quaternion representations q_{e_\bullet} of R_{e_\bullet} , so that our local rigidity regularisation term d_{reg} for the conjugate product edge e^* reads

$$d_{\text{reg}}(e^*) = \psi_2(\arccos(\langle q_{e_1}, q_{e_2} \rangle)), \quad (3)$$

where $\langle \cdot, \cdot \rangle$ is the inner product for quaternions [18]. $\psi_2(\cdot)$ is again a robust loss function, see Sec. 4.3.

4.3. Theoretical Analysis and Implementation

In the following we provide a theoretical analysis and additional implementation details.

Cyclic Shortest Paths. To find a *cyclic* shortest path,

we can run an ordinary (non-cyclic) shortest path algorithm (e.g. Dijkstra’s algorithm [10]) $|\mathcal{V}_{\mathcal{N}}|$ many times. To this end, we duplicate the last layers in the conjugate product graph (see Fig. 3), disconnect the duplicate layers from each other, and for each vertex from the ‘upper duplicate’ find the shortest path to the corresponding vertex in the ‘lower duplicate’. The globally optimal shortest *cyclic* path is now formed by the minimum among the $|\mathcal{V}_{\mathcal{N}}|$ individual paths. To reduce the number of shortest paths that need to be computed, we can instead resort to a more efficient branch-and-bound strategy, we refer to the Appendix for details.

Degenerate Cases. Conjugate product vertices containing degenerate edges of the 3D shape do not contain directional information on the 3D shape which we need to compute d_{reg} . We inject the relevant directional information into the conjugate product graph by introducing new conjugate product vertices that reflect (non-degenerate) edges on the 3D shape adjacent to respective degenerate 3D edge.

Pruning. To decrease the size of the conjugate product graph \mathcal{P}^* , we apply a pruning strategy. To this end, we prune conjugate product edges that reflect local turning points on the 3D shape since such paths represent undesirable matchings. In addition, we prune edges that first represent a degenerate edge of \mathcal{M} , followed by a degenerate edge of \mathcal{N} (or vice-versa). Such combinations are equivalent to a matching with two non-degenerate edges. Overall, our pruning reduces the graph size (and thus runtime) and excludes obvious non-desirable solutions.

Runtime Analysis. The runtime of our algorithm depends on the size of the conjugate product graph and the number of shortest path runs. The number of vertices in \mathcal{P}^* corresponds to the number of edges in \mathcal{P} . The number of edges in \mathcal{P}^* can be approximated by $c \cdot |\mathcal{V}_{\mathcal{M}}| \cdot (|\mathcal{E}_{\mathcal{N}}| + |\mathcal{V}_{\mathcal{N}}|)$ where c is a constant that is related to the maximum number of neighbours of the vertices in \mathcal{N} . Tab. 2 sums up the sizes of the product graph \mathcal{P} and the conjugate product graph \mathcal{P}^* .

	# vertices	# edges
\mathcal{P}	$ \mathcal{V}_{\mathcal{M}} \cdot \mathcal{V}_{\mathcal{N}} $	$ \mathcal{V}_{\mathcal{M}} \cdot (2 \mathcal{E}_{\mathcal{N}} + \mathcal{V}_{\mathcal{N}})$
\mathcal{P}^*	$ \mathcal{V}_{\mathcal{M}} \cdot (2 \mathcal{E}_{\mathcal{N}} + \mathcal{V}_{\mathcal{N}})$	$c \cdot \mathcal{V}_{\mathcal{M}} \cdot (2 \mathcal{E}_{\mathcal{N}} + \mathcal{V}_{\mathcal{N}})$

Table 2. Comparison of sizes of the product graph \mathcal{P} and the conjugate product graph \mathcal{P}^* .

Using $|\mathcal{E}_{\mathcal{N}}| \approx 3|\mathcal{V}_{\mathcal{N}}|$ [6] shows that the conjugate product graph \mathcal{P}^* has 7 times more vertices and $c \approx 11$ (see Appendix) times more edges than the product graph \mathcal{P} , which shows that asymptotically both graphs have the same size. In the worst case $\mathcal{O}(|\mathcal{V}_{\mathcal{N}}|)$ shortest path problems – one for each vertex in \mathcal{N} – have to be solved. Together with the complexity of each Dijkstra run the final time complexity can be estimated as $\mathcal{O}(|\mathcal{V}_{\mathcal{M}}| \cdot |\mathcal{V}_{\mathcal{N}}|^2 \cdot \log(|\mathcal{V}_{\mathcal{N}}|))$, which is the same as in [28]. We provide more details as well as runtime plots in the Appendix.

Implementation Details. We implement the shortest path algorithm in C++ wrapped in a MATLAB [32] mex-function. Computation of quantities on meshes, mesh simplification as well as local thickness computations are done using [20, 32]. For all experiments we choose $\psi_1(x)$ to be the robust loss function of [2], for which we choose $\alpha_1 = -2$ and $c_1 = 0.15$. For $\psi_2(x)$ we also choose the same loss function with $\alpha_2 = 0.7$ and $c_2 = 0.6$, but with a cubic bowl instead of a quadratic bowl as we want to ensure that small errors due to discretisation artefacts are not penalised. The choice of different $\psi_1(x)$ and $\psi_2(x)$ is required since d_{data} and d_{reg} have different ranges.

5. Experiments

In this section we compare our method on two datasets, conduct an ablation study, showcase results on partial shapes and for sketch-based shape manipulation. We emphasise that the matching of contours to 3D meshes is ill-posed: the same contour can arise from different configurations, *i.e.* the ground-truth is not necessarily unique, the space of solutions that seem natural is even bigger, and evaluation criteria that capture this non-uniqueness do not exist.

Datasets. We evaluate on the following two datasets:

- **TOSCA 2D-3D** [28]: 80 shapes of 9 different classes (humans, animals, etc.) in different poses. For each class exists at least one 2D shape.
- **FAUST 2D-3D** [28]: 100 human shapes in different poses subdivided into 10 classes. Each class has one 2D shape. Ground-truth correspondences between 2D and 3D are available for all instances.

Both datasets contain segmentation information across all shapes which form consistent 2D part to 3D part mappings.

Competing Approach. The only other method able to produce continuous matchings between 2D contours and 3D shapes is [28]. Due to their weaker model expressiveness that prevents the incorporation of a deformation prior, they use global spectral features and a pre-matched segmentation as additional feature in order to prevent degenerate solutions (e.g. collapsing). To enable a fair comparison, for both methods we provide results with and without this pre-matching. However, we consider the pre-matching as unrealistic prior knowledge, and therefore regard the cases without pre-matching as main results. As we show in Fig. 1, our results are superior without the segmentation term even in comparison to [28] using the segmentation term.

5.1. Matching

Next we evaluate our approach on the task of 2D-3D shape matching. First, we introduce a new error metric designed for the ambiguous setting of matching a contour onto a mesh. Subsequently, we compare quantitatively and qualitatively to the approach by Löhner *et al.* [28].




Figure 6. **Quantitative comparison** for the FAUST and TOSCA datasets. **Left:** Cumulative segmentation errors. The y-axis shows the percentage of points in the correct segment, and the x-axis the geodesic error threshold. Vacuously, when integrating the segmentation information into the optimisation (methods with suffix ‘-Seg’, dashed lines), the results are ‘perfect’ for both methods. **Right:** Cumulative geodesic errors on FAUST with and without left-right flips (manually removed for all approaches), which confirms that in many cases our method finds plausible solutions but does not resolve the intrinsic symmetry ambiguity. The y-axis shows the percentage of points below the x-axis threshold. We can see that our method consistently outperforms Lähler *et al.* [28]. Scores shown in the legends are respective areas under the curves.

Error Metric. We use two different error metrics: a) geodesic error and b) segmentation error. We only evaluate the geodesic error on FAUST due to the lack of 2D-3D ground truth correspondences in the TOSCA dataset. Additionally, there exist many valid matchings that may not correspond to the ground truth because the problem is ill-posed as explained above. Hence, we aim to derive a more robust quantitative evaluation for 2D-3D matchings. For that, we utilise part-based shape segmentations, which are available for all classes in the FAUST and TOSCA datasets and are generally consistent between 2D and 3D shapes. We argue that a good solution must have the same segmentation in the target domain, *i.e.* on the 3D shape, as in the source domain, *i.e.* on the 2D shape. For both we plot the cumulative curves measuring for each geodesic error value the percentage of matches with an error lower than this.

Geodesic Error. Let $(x, y) \in \mathcal{C} \subset \mathcal{V}_M \times \mathcal{V}_N$ be a computed match and \hat{y} be the ground-truth match of x . The normalised geodesic error of this matching is defined as

$$\varepsilon_{\text{geo}}(x, y) = \frac{\text{dist}_{\mathcal{N}}(y, \hat{y})}{\text{diam}(\mathcal{N})}. \quad (4)$$

Here $\text{dist}_{\mathcal{N}} : \mathcal{N} \times \mathcal{N} \rightarrow \mathbb{R}_0^+$ is the geodesic distance on \mathcal{N} and $\text{diam}(\mathcal{N}) = \max_{x, y \in \mathcal{N}} \text{dist}_{\mathcal{N}}(x, y)$.

Segmentation Error. Let $\sigma_{\mathcal{M}}(x)$ be the source segment of its matched point $y \in \mathcal{N}$ and let $\sigma_{\mathcal{N}}(y)$ be its target segment. We define the segmentation error as

$$\varepsilon_{\text{seg}}(x, y) = \min_{y' \in \mathcal{N}} \frac{\text{dist}_{\mathcal{N}}(y, y')}{\text{diam}(\mathcal{N})} \text{ s.t. } \sigma_{\mathcal{N}}(y') = \sigma_{\mathcal{M}}(x). \quad (5)$$

For shapes with symmetries or other ambiguities, we choose the best of all plausible segmentation combinations.

Quantitative Matching Results. In Fig. 6 (left) we show that our method outperforms the competing method by Lähler *et al.* [28] by a great margin in terms of the



Figure 7. **Qualitative results** of our method on FAUST. We can see the occurrence of left-right-flips (indicated by \swarrow) which nevertheless can be considered as plausible matchings.

segmentation error, both on FAUST on TOSCA. Since for FAUST ground truth is available, in Fig. 6 (right) we show the percentage of correct matchings, for which our method is superior. In addition, when left-right-flips (which form plausible solutions that stem from shape symmetries) are removed, our method (without pre-matching) is on par with the approach by Lähler *et al.* that uses pre-matching.

Qualitative Matching Results. We also compare our method qualitatively to Lähler *et al.* [28]. Even though our method is not using segmentation information, matchings computed with our approach are consistently of better quality and reflect a more plausible path on the 3D shape, *i.e.* are locally straight, see Fig. 1, Fig. 7 and Fig. 8.

5.2. Ablation Studies

In the Appendix we provide ablation studies considering different parts of our cost function, discretisation, shape discrepancies and noise.

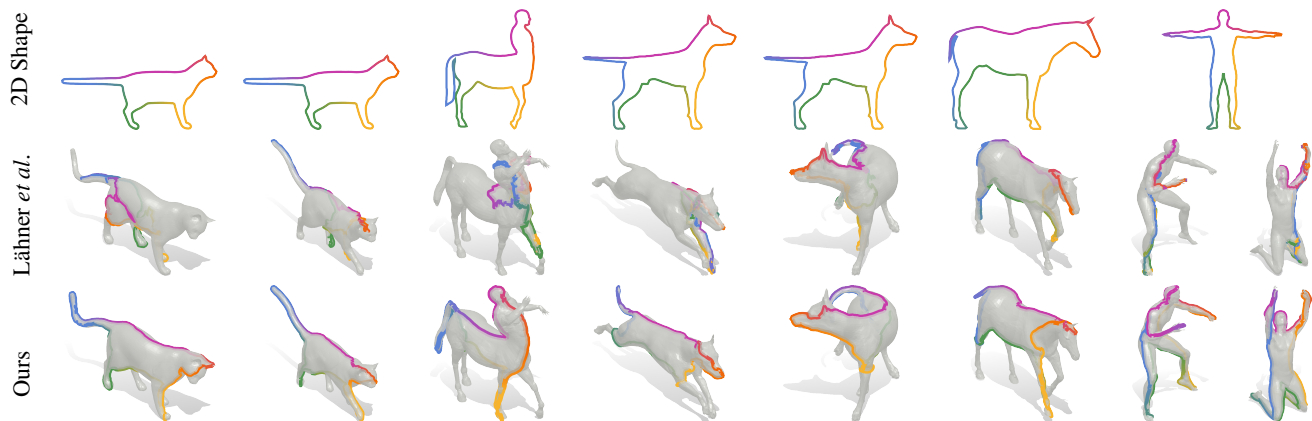


Figure 8. **Qualitative comparison** of the method by Lähler *et al.* [28] (second row) and our approach (third row) on TOSCA. Our approach results in more plausible matchings despite that Lähler *et al.* use a coarse segmentation-based pre-matching. Our local rigidity regulariser, which is enabled by our novel conjugate product graph formalism, ensures that resulting paths on 3D target shapes are much smoother.




Figure 9. **Qualitative comparison** of Lähler *et al.* [28] and ours on partial FAUST shapes. The global features of [28] result in poor matchings in scenarios without full shape, whereas we use local features and thus obtain valid partial matchings.

5.3. Partial Shapes

We show experiments on partial shapes, for which we removed parts of either the 2D or 3D shape in FAUST, see Fig. 9. Our approach is substantially more robust in the partial setting compared to Lähler *et al.* [28], likely due to the locality of our features and strong spatial regularisation, in contrast to the global spectral features of [28].

5.4. Sketch-Based 3D Shape Manipulation

We show the high quality of our matchings by performing 2D sketch-based 3D shape manipulation. After deforming the contour, the 3D shape is brought into a corresponding pose through as-rigid-as-possible shape deformation [45], see Fig. 1 (right). Details can be found in the Appendix.

6. Discussion & Limitations

Our experimental results confirm that conjugate product graphs enable 2D-3D shape matching without the need of a coarse pre-matching. Even though we compute results to global optimality, scenarios like symmetries (*e.g.* for human shapes) lead to ambiguities that are challenging to reflect in the cost function, which may result in matchings that collapse to one side of the 3D shape, see Fig. 7 (bottom-right). Although our method has the same asymptotic complexity as [28], in practice the computation is slower due to the conjugate product graph being larger (by a constant factor) than the product graph (cf. Tab. 2, also see Appendix).

7. Conclusion

We presented conjugate product graphs for 2D-3D shape matching, which for the first time allows for the incorporation of higher-order costs within path-based matching formalisms. Our novel concept significantly increases model expressiveness and flexibility, allowing to inject desirable properties, like local rigidity regularisation, into respective optimisation problems. Our results show significant improvements in challenging matching settings, even allowing for 2D sketch-based 3D shape manipulation. Since our powerful higher-order regularisation allows to get rid of the need for global features, our method is the first that solves *partial* 2D-3D shape matching. We believe that our work is of high relevance to the field of shape analysis, and hope to inspire more work on inter-dimensional applications.

Acknowledgements. PR is funded by the TRA Modelling (University of Bonn) as part of the Excellence Strategy of the federal and state governments. ZL is funded by a KI-Starter grant of the Ministry of Culture and Science NRW.

References

- [1] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. Pose-Consistent 3D Shape Segmentation Based on a Quantum Mechanical Feature Descriptor. In *Pattern Recognition*. Springer Berlin Heidelberg, 2011. 2
- [2] Jonathan T Barron. A general and adaptive robust loss function. In *CVPR*, 2019. 6
- [3] Claude Berge. *Hypergraphs: combinatorics of finite sets*. Elsevier, 1984. 3
- [4] Florian Bernard, Frank R Schmidt, Johan Thunberg, and Daniel Cremers. A combinatorial solution to non-rigid 3d shape-to-image matching. In *CVPR*, 2017. 3, 5
- [5] Florian Bernard, Christian Theobalt, and Michael Moeller. Ds*: Tighter lifting-free convex relaxations for quadratic matching problems. In *CVPR*, 2018. 3
- [6] Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Lévy. *Polygon Mesh Processing*. CRC press, 2010. 6
- [7] Yuri Boykov and Gareth Funka-Lea. Graph cuts and efficient nd image segmentation. *International journal of computer vision*, 70(2):109–131, 2006. 3
- [8] Lei Cai, Jundong Li, Jie Wang, and Shuiwang Ji. Line graph neural networks for link prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 3
- [9] Donatello Conte, Pasquale Foggia, Carlo Sansone, and Mario Vento. Thirty years of graph matching in pattern recognition. *International journal of pattern recognition and artificial intelligence*, 2004. 3
- [10] Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1, 1959. 4, 6
- [11] Hossein Eghdami and Majid Darehmiraki. Application of dna computing in graph theory. *Artificial Intelligence Review*, 2012. 3
- [12] Pedro F Felzenszwalb and Daniel P Huttenlocher. Pictorial structures for object recognition. *International journal of computer vision*, 2005. 3
- [13] Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, 1984. 3
- [14] Joan Glaunès, Anqi Qiu, Michael I Miller, and Laurent Younes. Large deformation diffeomorphic metric curve mapping. *International journal of computer vision*, 80:317–336, 2008. 2
- [15] Richard H Hammack, Wilfried Imrich, Sandi Klavžar, Wilfried Imrich, and Sandi Klavžar. *Handbook of product graphs*. CRC press Boca Raton, 2011. 3
- [16] Frank Harary. *Graph Theory*. Addison-Wesley, 1969. 3
- [17] Robert Herzog, Daniel Mewes, Michael Wand, Leonidas Guibas, and Hans-Peter Seidel. Lesss: Learned shared semantic spaces for relating multi-modal representations of 3d shapes. *Computer Graphics Forum*, 2015. 2
- [18] Du Q Huynh. Metrics for 3d rotations: Comparison and analysis. *Journal of Mathematical Imaging and Vision*, 2009. 5
- [19] Umar Iqbal, Anton Milan, and Juergen Gall. Posetrack: Joint multi-person pose estimation and tracking. In *CVPR*, 2017. 3
- [20] Alec Jacobson et al. gptoolbox: Geometry processing toolbox, 2021. <http://github.com/alecjacobson/gptoolbox>. 6
- [21] Ian H Jermyn, Sebastian Kurtek, Eric Klassen, and Anuj Srivastava. Elastic shape matching of parameterized surfaces using square root normal fields. In *ECCV*. Springer, 2012. 2
- [22] Long Jin, Yang Chen, Tianyi Wang, Pan Hui, and Athanasios V Vasilakos. Understanding user behavior in online social networks: A survey. *IEEE communications magazine*, 2013. 3
- [23] Evangelos Kalogerakis, Aaron Hertzmann, and Karan Singh. Learning 3d mesh segmentation and labeling. In *ACM SIGGRAPH*, 2010. 3
- [24] Jörg Hendrik Kappes, Markus Speth, Bjoern Andres, Gerhard Reinelt, and Christoph Schnörr. Globally optimal image partitioning by multicuts. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*. Springer, 2011. 3
- [25] Margret Keuper, Evgeny Levinkov, Nicolas Bonneel, Guillaume Lavoue, Thomas Brox, and Bjoern Andres. Efficient decomposition of image and mesh graphs by lifted multicuts. In *CVPR*, 2015. 3
- [26] Mikko Kivelä, Alex Arenas, Marc Barthélemy, James P Gleeson, Yamir Moreno, and Mason A Porter. Multilayer networks. *Journal of complex networks*, 2014. 3
- [27] Zhang Kun, Ma Xiao, and Li Xinguo. Shape matching based on multi-scale invariant features. *IEEE Access*, 2019. 2
- [28] Zorah Lähner, Emanuele Rodolà, Frank R Schmidt, Michael M Bronstein, and Daniel Cremers. Efficient globally optimal 2d-to-3d deformable shape matching. In *CVPR*, 2016. 1, 2, 3, 4, 6, 7, 8
- [29] Haibin Ling and David W. Jacobs. Shape classification using the inner-distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2007. 2
- [30] O. Litany, T. Remez, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *ICCV*, 2017. 2
- [31] Riccardo Marin, Marie-Julie Rakotosaona, Simone Melzi, and Maks Ovsjanikov. Correspondence learning via linearly-invariant embedding. In *NeurIPS*, 2020. 2
- [32] The Mathworks, Inc., Natick, Massachusetts. *MATLAB version 9.13.0.2049777 (R2022b)*, 2022. 6
- [33] Damien Michel, Iasonas Oikonomidis, and Antonis Argyros. Scale invariant and deformation tolerant partial shape matching. *Image and Vision Computing*, 2011. 2, 5
- [34] Joseph SB Mitchell, David M Mount, and Christos H Papadimitriou. The discrete geodesic problem. *SIAM Journal on Computing*, 1987. 3
- [35] Tomas Möller and Ben Trumbore. Fast, minimum storage ray/triangle intersection. In *ACM SIGGRAPH 2005 Courses*. 2005. 5
- [36] Eric N Mortensen and William A Barrett. Intelligent scissors for image composition. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, 1995. 3
- [37] Jie Qin, Shuaihang Yuan, Jiabin Chen, Boulbaba Ben Amor, Yi Fang, Nhat Hoang-Xuan, Chi-Bien Chu, Khoi-Nguyen Nguyen-Ngoc, Thien-Tri Cao, Nhat-Khang Ngo, Tuan-Luc

- Huynh, Hai-Dang Nguyen, Minh-Triet Tran, Haoyang Luo, Jianning Wang, Zheng Zhang, Zihao Xin, Yang Wang, Feng Wang, Ying Tang, Haiqin Chen, Yan Wang, Qunying Zhou, Ji Zhang, and Hongyuan Wang. Shrec'22 track: Sketch-based 3d shape retrieval in the wild. *Computers & Graphics*, 2022. 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. 2
- [39] Stefan Roth and Michael J Black. Fields of experts: A framework for learning image priors. In *CVPR*. IEEE, 2005. 3
- [40] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. “grabcut” interactive foreground extraction using iterated graph cuts. *ACM transactions on graphics (TOG)*, 2004. 3
- [41] Hiroaki Sakoe and Seibi Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing*, 1978. 2
- [42] Frank R Schmidt, Dirk Farin, and Daniel Cremers. Fast matching of planar shapes in sub-cubic runtime. In *IEEE International Conference on Computer Vision*, 2007. 2, 3
- [43] Thomas Schoenemann, Frank R. Schmidt, and Daniel Cremers. Image segmentation with elastic shape priors via global geodesics in product spaces. In *BMVC*, 2008. 2
- [44] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 2000. 3
- [45] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *SGP*, 2007. 5, 8
- [46] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *SGP*, 2009. 2
- [47] Jos MF Ten Berge. Orthogonal procrustes rotation for two or more matrices. *Psychometrika*, 1977. 5
- [48] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *ECCV*, 2010. 2
- [49] Remco C Veltkamp. Shape matching: Similarity measures and algorithms. In *Proceedings International Conference on Shape Modeling and Applications*, 2001. 2
- [50] Remco C Veltkamp and Michiel Hagedoorn. State of the art in shape matching. *Principles of visual information retrieval*, 2001. 2
- [51] Fang Wang, Le Kang, and Yi Li. Sketch-based 3d shape retrieval using convolutional neural networks. In *CVPR*, 2015. 2
- [52] Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT and EMNLP*, 2005. 3
- [53] Thomas Windheuser, Ulrich Schlickewei, Frank R. Schmidt, and Daniel Cremers. Geometrically consistent elastic matching of 3D shapes: A linear programming solution. In *ICCV*, 2011. 2, 3
- [54] Thomas Windheuser, Ulrich Schlickewei, Frank R. Schimdt, and Daniel Cremers. Large-Scale Integer Linear Programming for Orientation Preserving 3D Shape Matching. *Computer Graphics Forum*, 2011. 2
- [55] Stephan Winter. Modeling costs of turns in route planning. *GeoInformatica*, 2002. 3
- [56] Li Zhang, Yuan Li, and Ramakant Nevatia. Global data association for multi-object tracking using network flows. In *CVPR*, 2008. 3