

A Scalable Combinatorial Solver for Elastic Geometrically Consistent 3D Shape Matching

Paul Roetzer^{1,2}
 TU Munich¹

Paul Swoboda³
 University of Bonn²

Daniel Cremers¹
 MPI Informatics³

Florian Bernard²




Figure 1. We propose a **novel combinatorial solver for the non-rigid matching of 3D shapes** based on discrete orientation-preserving diffeomorphisms [66] (left). For the first time we utilize an orientation-preserving diffeomorphism to constrain the challenging problem of non-rigidly matching a pair of partial shapes without availability of complete shapes (center). Our solver scales significantly better compared to existing solvers and can thus handle shapes with practically relevant resolutions (right).

Abstract

We present a scalable combinatorial algorithm for globally optimizing over the space of geometrically consistent mappings between 3D shapes. We use the mathematically elegant formalism proposed by Windheuser et al. [66] where 3D shape matching was formulated as an integer linear program over the space of orientation-preserving diffeomorphisms. Until now, the resulting formulation had limited practical applicability due to its complicated constraint structure and its large size. We propose a novel primal heuristic coupled with a Lagrange dual problem that is several orders of magnitudes faster compared to previous solvers. This allows us to handle shapes with substantially more triangles than previously solvable. We demonstrate compelling results on diverse datasets, and, even showcase that we can address the challenging setting of matching two partial shapes without availability of complete shapes. Our code is publicly available at <http://github.com/paul0noah/sm-comb>.

1. Introduction

The shape matching problem is widely studied in graphics and vision due to its high relevance in numerous appli-

cations, including 3D reconstruction, tracking, shape modeling, shape retrieval, interpolation, texture transfer, or the canonicalization of geometric data for deep learning. Shape matching refers to finding correspondences between parts of shapes – for example, the shapes may be represented as triangular surface meshes, and correspondences may be obtained between vertices or triangles of individual shapes. While identifying such correspondences is a relatively easy task for humans, from a computational perspective shape matching is much more challenging. This is because many formulations lead to high-dimensional combinatorial optimization problems. While some of them are efficiently solvable, e.g. based on the linear assignment problem [47], such simple approaches do not take geometric relations between the shape parts into account and thus typically lead to poor matchings. In contrast, most practically relevant formulations account for geometric consistency in some form, which in turn lead to significantly more difficult optimization problems (e.g. the NP-hard quadratic assignment problem [51], or mixed-integer programming formulations [3]).

A decade ago Windheuser et al. [66, 67] proposed an elegant formalism for the deformable matching of 3D shapes. Most notably, their approach is able to account for geometric consistency based on an orientation-preserving discrete diffeomorphism. While this formulation is conceptually ap-

pealing, it requires to solve a difficult integer linear programming (ILP) problem. The authors proposed to relax the binary variables to continuous variables, so that a linear programming (LP) problem is obtained. Although this LP formulation is convex and can thus be solved to global optimality, the resulting optimization problem is prohibitively large, and in turn only shapes with low resolution (at most 250 triangles) can be matched with existing solvers.

In this work we close this gap and propose a scalable combinatorial solver for geometrically consistent deformable 3D shape matching. Our main contributions are:

- We propose a novel primal heuristic, which, together with a Lagrange dual problem, gives rise to a combinatorial solver for orientation-preserving deformable 3D shape matching based on the formalism by Windheuser et al. [66,67].
- We show that in the special case of consistent triangulations between both shapes our primal heuristic can readily be used to solve the ILP formulation to global optimality in polynomial time.
- Our solver is orders of magnitude faster than previous solvers, which allows us to handle shapes with a substantially higher resolution.
- Experimentally we demonstrate state-of-the-art results on a range of different shape matching problems, and we showcase the applicability to the difficult case of partial-to-partial shape matching.

2. Related Work

In the following we give an overview of existing shape matching approaches. For a more extensive overview we refer interested readers to the survey papers [59,64].

Rigid shape matching methods consider the matching of shapes under the assumption that one shape undergoes a rigid-body transformation. For known correspondences, the resulting orthogonal Procrustes problem has a closed-form solution [61]. In general, the assumption that correspondences are known is a severe limitation and typically one also wants to find the correspondences. Such tasks are commonly tackled via local optimization, *e.g.* via the Iterative Closed Point (ICP) algorithm [6] or variants thereof [7, 21, 48]. The main drawback of local methods is that they heavily depend on the initialization and generally do not obtain global optima. There are also global approaches, for example based on semidefinite programming [43], Branch and Bound (BnB) [49], or quasi-BnB [15]. A major limitation of *rigid* matching approaches is that in practice shapes are often non-rigidly deformed, for which rigid transformation models are too restricted.

For shapes that undergo an elastic transformation, **non-rigid shape matching** is more suitable. One example is the functional map framework [50], which have led to remarkable results for isometric shapes [50,55]. Yet, their

main drawback is that they are generally sensitive to noise and less reliable in non-isometric settings. A wide variety of extensions of functional maps was developed to overcome specific drawbacks, *e.g.* related to computation time [32], missing parts [56], orientation preservation [55], deblurring [22], denoising [54] non-isometric shapes [18], or multi-matching [26,33].

An alternative non-rigid shape matching formulation is via the **Quadratic Assignment Problem (QAP)**, also known as graph matching. Due to the NP-hardness of the QAP [51], numerous heuristic methods have been proposed [38,39]. There also exist various convex relaxation methods for the QAP [4, 16, 25, 35, 63]. Although these methods are appealing as they can produce optimality bounds, they cannot guarantee to find global optima in all cases. In [2], a globally optimal but exponential-time BnB method was presented. In [65], the authors consider an iterative linearization of the QAP for addressing shape matching. Recently, simulated annealing was used for QAP-based shape matching [31]. Similar to the QAP, the ILP problem we address is NP-hard in general and theoretically solvable with (exponential-time) BnB methods.

There are various alternative **local methods** for non-rigid shape matching, *e.g.* based on Gromov-Hausdorff distances [9,46], as-conformal-as-possible formulations [42], iterative spectral upsampling [45], higher-order projected power iterations, [5], a discrete solver for functional map energies [53], a hybrid spatial-spectral approach [68], triangle-based deformation models [62], elastic membrane energy optimization [23], and many more. A major downside of local methods is that they heavily depend on a good initialization.

Opposed to local optimization approaches are **global methods**, which have the strong advantage that they are independent of the initialization. Globally optimal non-rigid matchings can be obtained using shortest path algorithms for 2D shape matching [13, 24] and for 2D-to-3D matching [36]. Unfortunately, shortest path algorithms are not applicable for 3D-to-3D shape matching, since matchings cannot be represented as a shortest path but rather form a minimal surface [66]. An alternative formulation based on a convex relaxation was proposed in [11], which, however relies on an extrinsic term that requires a good spatial alignment. In [3] a BnB strategy is used to tackle a mixed-integer programming formulation that utilizes a low-dimensional discrete matching model. Despite the exponential worst-case complexity, it was shown that global optimality can be certified in most of the considered shape matching instances. The framework by Windheuser et al. [60, 66, 67] considers an ILP formulation for the orientation-preserving diffeomorphic matching of 3D shapes. The main issue is that to date there is no efficient combinatorial solver that can solve large instances of respective problems, so that their

framework is currently impracticable. The framework will be discussed in more detail in Sec. 3. The main objective of our work is to propose the first efficient solver tailored specifically to this formalism.

Learning-based shape matching. A wide variety of learning-based approaches has been used to address shape matching. While unsupervised approaches do not need time-consuming labeling of data [14, 19, 20, 30, 58, 69], supervised methods rely on the availability of labeled data [10, 28, 40, 52, 57]. Overall, learning-based techniques are ideal to obtain task-specific features to solve shape matching problems, and it was demonstrated that respective approaches achieve remarkable results. However, a major difficulty is that they often lack the ability to generalize to other types of shapes, which in contrast is a major strength of learning-free approaches. In this work we focus on a specific formalism in the latter class of optimization-based learning-free methods, and we believe that our work may be a first step towards integrating such powerful methods into modern learning approaches in order to eventually achieve the best from both worlds.

3. Background of the Shape Matching ILP

In the following we recapitulate the shape matching integer linear program (ILP) of Windheuser et al. [66]. We provide the used notation in Tab. 1.

X, Y	shapes (triangular surface meshes)
(V_X, E_X, F_X)	vertices, edges and triangles of shape X
F	product space of vertex-triangle, edge-triangle and triangle-triangle pairs
E	edge product space
$\Gamma \in \{0, 1\}^{ F }$	indicator vector of matches
π_X, π_Y	product vector projection on X and Y
∂	geometric consistency constraints
\mathbb{E}	matching energy

Table 1. Notation used in this paper.

Definition 1 (Shape). We define a shape X as a triplet (V_X, E_X, F_X) of vertices V_X , edges $E_X \subset V_X \times V_X$ and triangles $F_X \subset V_X \times V_X \times V_X$, such that the manifold induced by the triangles is oriented and has no boundaries¹.

In the remainder we will consider two shapes $X = (V_X, E_X, F_X)$ and $Y = (V_Y, E_Y, F_Y)$ that shall be elastically matched to each other. A matching between X and Y is defined by triangle-triangle, triangle-edge or triangle-vertex correspondences between pairs of elements in X and Y . We allow for triangle-edge and triangle-vertex matches

¹We note that partial shapes can be handled by closing holes and defining suitable costs for matching holes, see Appendix for more details.

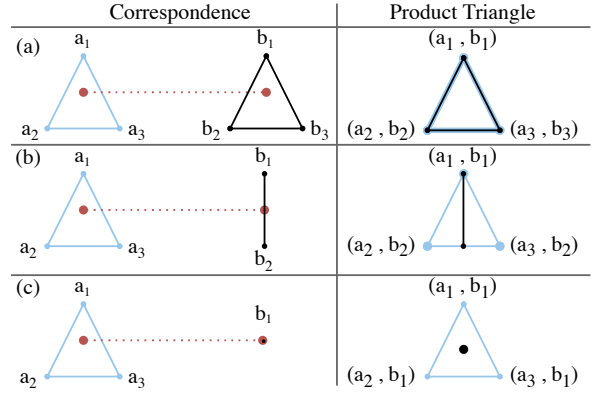


Figure 2. **Correspondences represented as product triangles of F :** If a product triangle (right) is part of the computed solution it implies that the respective triangle is matched to a triangle (top row), an edge (middle) or a vertex (bottom). Each product triangle is associated with a local matching cost that encodes feature similarity and costs for stretching/compressing and bending.

to account for compressing and stretching shapes. For convenience, by \bar{F}_\bullet we denote the set of *degenerate* triangles, that in addition to the triangles F_\bullet also contains triangles formed by edges (a triangle with two vertices at the same position) and triangles formed by vertices (a triangle with three vertices at the same position). Similarly, we consider edge products and the set of degenerate edges \bar{E}_\bullet .

Definition 2 (Product Spaces). Let two shapes X and Y be given. The triangle product space is defined as

$$F := \left\{ \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \\ a_3, b_3 \end{pmatrix} \mid \begin{array}{l} (a_1 a_2 a_3 \in F_X \wedge b_1 b_2 b_3 \in \bar{F}_Y) \vee \\ (a_1 a_2 a_3 \in \bar{F}_X \wedge b_1 b_2 b_3 \in F_Y) \end{array} \right\}.$$

The edge product space is defined as

$$E := \left\{ \begin{pmatrix} a_1, b_1 \\ a_2, b_2 \end{pmatrix} \mid \begin{array}{l} (a_1 a_2 \in E_X \wedge b_1 b_2 \in \bar{E}_Y) \vee \\ (a_1 a_2 \in \bar{E}_X \wedge b_1 b_2 \in E_Y) \end{array} \right\}.$$

An illustration of possible correspondences represented as product triangles, *i.e.* elements of the product triangle space, can be seen in Fig. 2. In order to guarantee a geometrically consistent matching, we impose two types of constraints:

(i) **Projection constraints π .** We require that all triangles from X must be matched to a vertex, edge or triangle from Y , and vice versa.




Figure 3. Two **product triangles** are neighboring if they share the same product edge with opposite orientation.

(ii) **Geometric consistency constraints** ∂ . Neighboring elements of X must be matched to neighboring elements of Y . Whenever a product edge is part of the matching, there must exist exactly two product triangles sharing the product edge oriented in opposite directions, see Fig. 3. This ensures geometric consistency by requiring that the matching is a two-manifold in the product space, which also implies that the respective 3D mapping is orientation preserving.

Together with the energy $\mathbb{E} \in \mathbb{R}^{|F|}$ that quantifies the matching cost of individual correspondences (elements of the product space), the ILP shape matching problem reads

$$\min_{\Gamma \in \{0,1\}^{|F|}} \mathbb{E}^\top \Gamma \quad \text{s.t.} \quad \begin{pmatrix} \pi_X \\ \pi_Y \\ \partial \end{pmatrix} \Gamma = \begin{pmatrix} \mathbf{1}_{|F_X|} \\ \mathbf{1}_{|F_Y|} \\ \mathbf{0}_{|E|} \end{pmatrix}, \quad (\text{ILP-SM})$$

where Γ is the matching vector represented as indicator vector of the triangle product space F . For more details about (ILP-SM) we refer to [60, 66, 67].

Proposition 3. *If only triangle-triangle correspondences are allowed, problem (ILP-SM) is solvable to global optimality in polynomial time.*

Proof. The two-manifold property of the matching and of both shapes X and Y yields that a single triangle-triangle match determines all other triangle-triangle matches due to geometric consistency. \square

Despite the conceptual elegance of the formulation, as soon as degenerate matchings are allowed for (e.g. triangle-edge or triangle-vertex), which is necessary for virtually any real-world shape matching instance, problem (ILP-SM) is significantly more difficult. Specifically, it belongs to the class of ILPs, which are NP-hard in general. Windheuser et al. [60, 66, 67] address this based on an LP relaxation that relaxes the binary variables to continuous ones. However, even their convex LP formulation has several drawbacks that impede a practical application:

(i) The relaxed LP formulation involves a *very large number of variables*. For example, non-rigidly matching two 3D shapes with 1000 faces each leads to a total of about $2 \cdot 10^7$ variables. The authors implement an efficient parallelized GPU-based primal-dual solver [17], which can handle problems with up to 250 faces (leading to about 10^6 binary variables), requiring a total time of about 2 h. The same applies to modern state-of-the-art LP solvers [29]. Currently no solver exists that can solve even moderately-sized instances of problem (ILP-SM).

(ii) Windheuser et al. attempt to address this based on a coarse-to-fine scheme. Initially, at the coarsest scale, they match severely downsampled shapes, and repeatedly apply their framework to refine the non-rigid matching only in the local neighborhood of the matching at the previous coarser scale. Overall, this has the downside that the final matching

substantially depends on the initial matching at the coarsest scale. With that, there is the risk that the initial low resolution shapes do not contain sufficient details for finding a reliable matching.

(iii) Furthermore, solutions of the continuous relaxation are generally not binary, so that a rounding step is necessary to obtain a discrete solution. To this end, the authors propose to repeatedly solve the expensive LP relaxations while gradually fixing more and more variables to be binary.

Overall, to make problem (ILP-SM) practicable, being able to solve larger shape matching instances is of great importance. To address this, we propose an efficient combinatorial solver based on the Lagrange decomposition to optimize problem (ILP-SM).

4. Lagrange Decomposition

Next, we introduce our Lagrange decomposition reformulation for problem (ILP-SM), which is amenable to dual optimization.

Decomposition into subproblems. We associate a small subproblem \mathcal{S} for each row of the constraint matrix in (ILP-SM):

Definition 4 (Subproblems). *For each individual projection and geometry consistency constraint we define a set of subproblems as*

$$\forall j : \mathcal{S}_j^{\pi_X} = \{\Gamma \in \{0,1\}^{|F|} : \sum_i (\pi_X)_{ji} \Gamma_i = 1\}, \quad (1)$$

$$\forall j : \mathcal{S}_j^{\pi_Y} = \{\Gamma \in \{0,1\}^{|F|} : \sum_i (\pi_Y)_{ji} \Gamma_i = 1\}, \quad (2)$$

$$\forall j : \mathcal{S}_j^\partial = \{\Gamma \in \{0,1\}^{|F|} : \sum_i (\partial)_{ji} \Gamma_i = 0\}. \quad (3)$$

We write the set of all subproblems as

$$\mathcal{S} = \{\mathcal{S}_{j=1,\dots,|F_X|}^{\pi_X}, \mathcal{S}_{j=1,\dots,|F_Y|}^{\pi_Y}, \mathcal{S}_{j=1,\dots,|E|}^\partial\}. \quad (4)$$

With the above decomposition we can write the Lagrange dual shape matching problem as

$$\begin{aligned} \max_{\lambda = \{\lambda^{\mathcal{S}}\}} \quad & \sum_{\mathcal{S} \in \mathcal{S}} \min_{\Gamma \in \mathcal{S}} \langle \lambda^{\mathcal{S}}, \Gamma \rangle \\ \text{s.t.} \quad & \sum_{\mathcal{S} \in \mathcal{S}} \lambda^{\mathcal{S}} = \mathbb{E}. \end{aligned} \quad (\text{LD-SM})$$

Min-marginals. While we cannot easily derive a matching from a solution λ of the dual problem (LD-SM), we can nevertheless obtain important dual costs that will guide our primal solution search. We use dual costs based on min-marginals, which are defined as follows:

Definition 5 (Min-marginals). *For any subproblem $\mathcal{S} \in \mathcal{S}$ we define the min-marginal for the i -th variable as the difference of the optima with the corresponding variable fixed to 1 vs. 0 as*

$$m_i^{\mathcal{S}} = \min_{\Gamma \in \mathcal{S}: \Gamma_i = 1} \langle \lambda^{\mathcal{S}}, \Gamma \rangle - \min_{\Gamma \in \mathcal{S}: \Gamma_i = 0} \langle \lambda^{\mathcal{S}}, \Gamma \rangle. \quad (5)$$

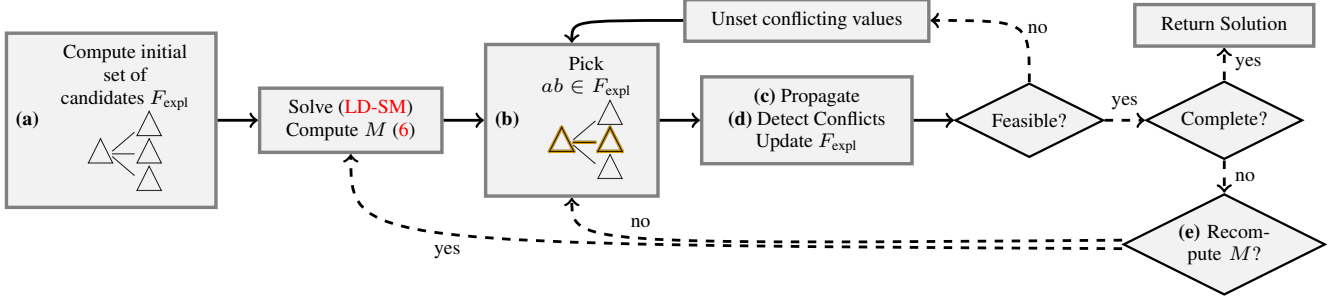


Figure 4. The **pipeline of our combinatorial solver** for the ILP shape matching problem (**ILP-SM**). The individual stages (a)-(e) are explained in Sec. 5.

In words, a min-marginal quantifies by how much a variable wants to attain the value 1 resp. 0 in the subproblem \mathcal{S} .

The total min-marginal is defined as the sum over all min-marginals

$$M_i = \sum_{\mathcal{S} \in \mathcal{S}} m_i^{\mathcal{S}}. \quad (6)$$

If for each variable all min-marginals have the same sign, and the total min-marginal is non-zero, we can directly reconstruct a primal solution by setting $\Gamma_i = 1$ if $M_i < 0$ $\Gamma_i = 0$ if $M_i > 0$. This case occurs when the relaxation defined by the Lagrange decomposition (**LD-SM**) is tight, which is not true in general. If not tight, the above reconstruction will result in infeasible Γ . However, in that case good solutions will mostly agree with the sign of the total min-marginal, which we exploit in our primal rounding strategy. We optimize the Lagrange decomposition and compute min-marginals with the approximate solver [37].

5. Primal Rounding

In the following we introduce our heuristic for primal rounding, i.e. obtaining primal solutions based on the dual costs that were computed through the Lagrange decomposition (**LD-SM**). Before we explain details, we provide a high-level summary of the main concept: First, we pick a suitable initial product triangle as first correspondence candidate. After we have solved the Lagrange decomposition and obtained the total min-marginals, we iteratively add product triangle matchings, i.e. we select elements of the matching vector Γ that are to be set to 1. This may induce additional product triangle correspondences (due to the constraints in problem (**ILP-SM**)), so we also force assignments of other variables. In case a conflict arises, we detect it and backtrack. After a given number of variable assignments, or if too many backtracking steps occur, we solve the Lagrange decomposition (**LD-SM**) again, while fixing the already found correspondences. With the updated total min-marginals we start the search again until we find a complete solution. Our overall pipeline is shown in Fig. 4, which we explain next in detail.

(a) Initialization. Given the empty matching, we choose the first triangle-triangle matching as follows:

- (i) Consider the shape $Z \in \{X, Y\}$ with fewer triangles.
- (ii) In Z , we choose the most regular triangle z (all angles between 20° and 90° , area is close to the mean area of all triangles, triangle lies in low curvature region).
- (iii) We select all elements in Γ that form non-degenerate triangle-triangle matchings of z .
- (iv) Among these, we choose the matching candidate with smallest total min-marginals.

(b) Exploration of candidates. We maintain a set of individual triangle-triangle matchings, i.e. elements of the matching vector Γ , to explore in subsequent iterations $F_{\text{expl}} \subset F$. After adding a product triangle $ab = (a_1b_1, a_2b_2, a_3b_3) \in F_{\text{expl}}$ to our matching Γ , we add all other product triangles $a'b' \in F$ that share a product edge oriented in the opposite direction with the currently selected ab to F_{expl} . In other words, in the next iteration we select one of the tentative product triangles such that the **Geometric Consistency I** constraint (Tab. 2) is fulfilled for one of the already selected product triangles. We explore product triangles in F_{expl} according to their total min-marginals in ascending order (see Appendix for a formal description).

Overall, the purpose of exploring only the elements of F_{expl} minimizes the possibility of obtaining several disjoint submatchings that do not fit together geometrically.

(c) Constraints & propagation. After each individual variable assignment we analyze all constraints that involve the modified variable, and then check whether any other variable assignments are forced. If so, we set the forced variable and recursively propagate. All used constraints and propagation rules are summarized shown in Tab. 2.

(d) Conflict detection. After each variable assignment we check for potential conflicts due to two potential cases:

- (i) *Infeasibility*: individual constraints are not satisfiable anymore by the variables that are not set yet.
- (ii) *Contradicting assignments*: variables which are already set to one would have to be set to zero by the propagation and vice versa.

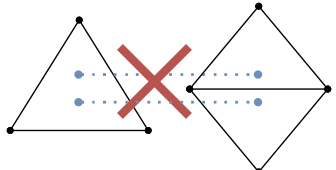
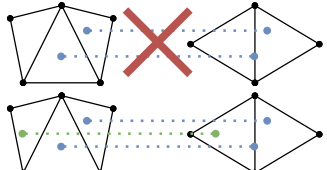
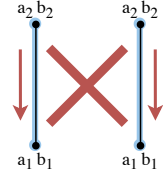
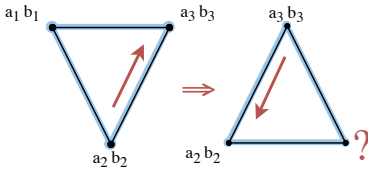
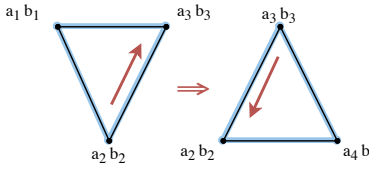
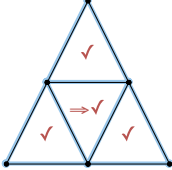
<p style="text-align: center;">Injectivity</p>  <p>Constraint: Each face can only have one matching. Propagation: Matching a triangle from one shape implies all other correspondences involving the same triangle have to be set to 0.</p>	<p style="text-align: center;">Surjectivity</p>  <p>Constraint: Each face of each shape needs to have a matching. Propagation: If all correspondences involving a given triangle are 0, then the remaining one is set to 1.</p>	<p style="text-align: center;">Two-manifold</p>  <p>Constraint: The same product edge orientation can only occur once in the solution Propagation: If a product edge is part of a matching, all other occurrences of the same product edge are set to 0.</p>
<p style="text-align: center;">Geometric Consistency I</p>  <p>Constraint: Each product edge needs to have a counterpart (opposite orientation). Propagation: A solution must include exactly one of the product triangles containing the opposite orientation of the product edge.</p>	<p style="text-align: center;">Geometric Consistency II</p>  <p>Constraint: Each product edge needs to have a counterpart (opposite orientation). Propagation: If there is only one counterpart left within the possible matches, it has to be part of the solution.</p>	<p style="text-align: center;">Closedness</p>  <p>Constraint: No holes are allowed in the matching. Propagation: Whenever all three neighbors of a product face are part of the solution, the product face itself is also part of the solution.</p>

Table 2. **Matching constraints and the derived propagation rules** used in the primal heuristic. Product triangles are visualized with blue-black edges, triangles of 3D shapes with black edges.

In cases of conflicts we perform backtracking by undoing the assignments and respective propagations that participated in the conflict.

(e) Recomputation of min-marginals. In order to better reflect the quality of product triangle candidates in F_{expl} w.r.t. already obtained partial matching, we regularly recompute (total) min-marginals. Whenever a certain amount of product triangles is set, we fix the respective variables in (ILP-SM), dualize the subproblem to obtain a new reduced Lagrange decomposition (LD-SM), and eventually optimize again to obtain the updated total min-marginals M_i . We refer to the Appendix for details.

6. Experiments

In the following we experimentally evaluate our approach on various datasets in a range of different settings.

Shape matching data. In our experiments we consider shape matching instances from several datasets: TOSCA [8], TOSCA partial [56], SHREC-watertight [27], SMAL [70], SHREC '19 [44] and KIDS [57]. We down-sample all meshes to about at most 1000 faces. We do not perform post-processing on the obtained matchings. The energy \mathbb{E} of problem (ILP-SM) is computed analogously to [67].

Shape matching algorithms. Since our main objective is to improve the computational performance of the best existing solver for problem (ILP-SM), as a baseline we reim-

plemented the rounding strategy proposed by Windheuser et al. [66]² based on the state-of-the-art LP solver Gurobi [29]. For further details we refer to the Appendix.

In addition, we also compare our solver for problem (ILP-SM) with two recent state-of-the-art methods that rely on other shape matching formalisms. Among them is a method based on smooth shells (Eisenberger et al.) [18], and a method based on a discrete functional map optimization framework (Ren et al.) [53].

6.1. Combinatorial Solvers for Problem (ILP-SM)

First, we compare against the directly related approach of Windheuser et al. [66], which solves the same problem (ILP-SM) as ours.

In Fig. 1 (right) we show the scalability of the solver of Windheuser et al. and ours depending on the number of triangles per shape. We find that while Windheuser et al. already takes 1 h for low-resolution shapes with ≈ 200 triangles (leading to a total of $\approx 8 \cdot 10^5$ binary variables in problem (ILP-SM)), our method scales significantly better and can handle shapes with substantially higher resolutions. Our method has a linear memory consumption (to the problem size, which is quadratic in the shape resolution). We note that the bump in the graph stems from the heuristically determined recomputation of the min-marginals which may

²The original code is not available.




Figure 5. **Qualitative comparison** of the method by Eisenberger et al. [18] (second row), Ren et al. [53] (third row) and Ours (last row). While Eisenberger et al. and Ren et al. do not guarantee orientation preservation they lead to erroneous matchings (e.g. left-right flips, see red arrows), whereas our method leads to smooth and orientation-preserving matchings.

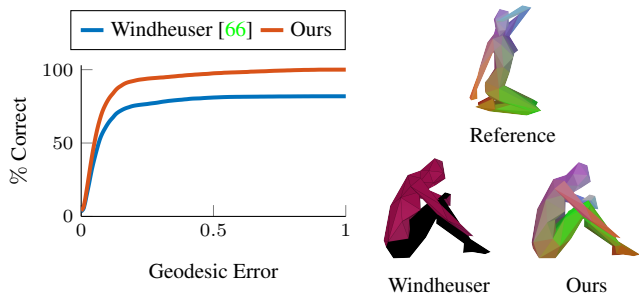


Figure 6. Comparison of the average percentage of correct matchings for the entire TOSCA dataset of **Windheuser et al. vs. Ours** (left). The horizontal axis shows the geodesic error threshold, and the vertical axis shows the percentage of matches that are smaller than or equal to this error. For Windheuser et al. we allow the solver to take $10\times$ more time than our method needed (per shape matching instance) – even then the curve of Windheuser et al. is low because it is unable to find good matchings within the given time budget, see the qualitative example on the right (black shows unmatched parts, all shapes have 175 triangles).

vary for individual matching problems (see Sec. A3 in Appendix).

In Fig. 6 we show quantitative and qualitative results of both solvers on the full TOSCA dataset, where we have

found that our method performs significantly better with an average area under the curve (higher is better) of 0.91 vs. 0.72 for Windheuser et al. (see Appendix for more details).

6.2. Comparison to State of the Art

Next, we compare our method to recent state-of-the-art shape matching approaches. Since Eisenberger et al. [18] show that their recent method substantially outperforms a range of other methods (BCICP [55], Zoomout [45], KM [65], FM [50], BIM [34]) on various datasets, for our comparison we only focus on the method by Eisenberger et al. [18], and in addition also include results of the more recent method of Ren et al. [53].

In Fig. 5 we show qualitative results for various non-rigid 3D shape matching instances from the datasets TOSCA, SHREC-watertight, SMAL and KIDS. The method of Ren et al. suffers from matchings that are not geometrically consistent, which thus leads to nonsmooth matchings (e.g. the pliers in the sixth column, or the kid in the eighth column). Moreover, both Ren et al. and Eisenberger et al. do not guarantee orientation preservation, so that for example left-right flips occur (e.g. the lion in the fifth column for Eisenberger et al., or the cat in the second-last column for Ren et al.). In contrast, our method obtains reliable matchings in these cases which are smooth and preserve the orientation.

In Fig. 8 we show quantitative results on the TOSCA,




Figure 7. Our method can handle the difficult case of **matching pairs of partial shapes without availability of complete shapes**, in which the orientation-preserving diffeomorphism serves as powerful constraint (boundaries are shown as red lines). (Best viewed magnified on screen)

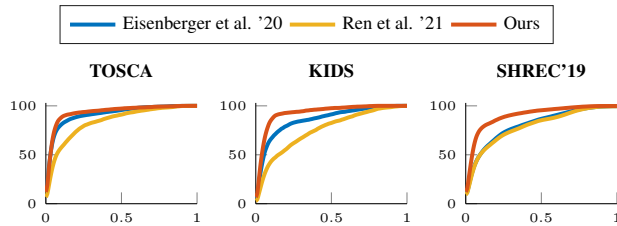


Figure 8. Comparison of the percentage of correct matchings for the **TOSCA dataset** (left), the **KIDS dataset** (middle) and the **SHREC'19 dataset** (right). The horizontal axis shows the geodesic error threshold, and the vertical axis shows the percentage of matches smaller than or equal to this error. Ours obtains the highest area under curve (\uparrow) (**0.929** for ours vs. 0.836 for Ren et al. and 0.910 for Eisenberger et al. on TOSCA).

KIDS and SHREC'19 datasets. Despite the fact that our method aims to directly solve a high-dimensional ILP with up to $2 \cdot 10^7$ binary variables (see problem (ILP-SM)), our method outperforms the baselines in terms of solution quality. Due to the large number of binary variables, our approach requires 22.68 min on average to compute a matching (opposed to few seconds for Ren et al. and Eisenberger et al.). For some individual shape classes we have found that ours has slightly worse performance (see Appendix), which stems from poor local optima for individual shape matching instances.

6.3. Partial-to-Partial Non-Rigid Matching

In Fig. 7 we showcase that our proposed solver is even able to handle difficult non-rigid partial-to-partial shape matching problems. Although partial shape matching has attracted a lot of attention recently [3, 40, 41, 45, 56, 65, 68], most existing works typically consider the case of matching a partial shape to a complete shape. There are also some partial-partial shape matching methods [1, 12] that build upon machine learning, which we consider to be orthogonal to our method. In contrast, for the first time we utilize orientation-preserving diffeomorphisms to constrain the challenging problem of non-rigidly matching a pair of

partial shapes without availability of complete shapes.

7. Discussion & Limitations

Our experimental analysis has confirmed that we can compute matchings involving up to $2 \cdot 10^7$ variables on a range of datasets within about 2h. Although our solver does not guarantee to find globally optimal solutions and thus may lead to suboptimal results in some cases (see Appendix), our experiments confirm that in most cases we produce high-quality matchings, and that we can even handle partial-to-partial shape matching.

8. Conclusion

We proposed a novel combinatorial solver for efficiently computing solutions to the mathematically elegant integer linear programming formulation of 3D shape matching introduced in [66]. Our solver consists of a primal-dual approach where the primal step makes use of min-marginals computed globally for the full problem. The original solver of [66] could only handle shapes of up to 250 triangles and therefore had to be applied in a heuristic coarse-to-fine strategy. In contrast, the proposed method leads to a drastic speedup and can therefore handle ILPs with millions of variables and 3D shapes of practically relevant resolution. This allows us to globally optimize over all orientation preserving geometrically consistent non-rigid mappings providing state-of-the-art matching results. Moreover, we extend this combinatorial matching approach to challenging scenarios like partial-to-partial shape matching without availability of the complete shapes. We argue that also generalizations, e.g. to more than two shapes, should be practically solvable through a similar combinatorial optimization approach.

References

- [1] Souhaib Attaiki, Gautam Pai, and Maks Ovsjanikov. Dpfm: Deep partial functional maps. In *3DV*, 2021. 8
- [2] Mokhtar. S. Bazaraa and A. N. Elshafei. An exact branch-and-bound procedure for the quadratic-assignment problem.

- Naval Research Logistics Quarterly*, 26(1):109–121, Mar. 1979. 2
- [3] Florian Bernard, Zeeshan Khan Suri, and Christian Theobalt. MINA: Convex Mixed-Integer Programming for Non-Rigid Shape Alignment. In *CVPR*, 2020. 1, 2, 8
- [4] Florian Bernard, Christian Theobalt, and Michael Moeller. DS*: Tighter Lifting-Free Convex Relaxations for Quadratic Matching Problems. In *CVPR*, 2018. 2
- [5] Florian Bernard, Johan Thunberg, Paul Swoboda, and Christian Theobalt. HiPPI: Higher-order projected power iterations for scalable multi-matching. In *ICCV*, 2019. 2
- [6] Paul J. Besl and Neil D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, Feb. 1992. 2
- [7] Seth Billings and Russell Taylor. Iterative Most Likely Oriented Point Registration. In *MICCAI*, 2014. 2
- [8] Alexander Bronstein, Michael Bronstein, and Ron Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer Science & Business Media, 2008. 6
- [9] Alexander M Bronstein, Michael M Bronstein, Ron Kimmel, Mona Mahmoudi, and Guillermo Sapiro. A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision*, 89(2-3):266–286, 2010. 2
- [10] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *CVPR*, 2017. 3
- [11] Qifeng Chen and Vladlen Koltun. Robust Nonrigid Registration by Convex Optimization. In *ICCV*, 2015. 2
- [12] Luca Cosmo, Emanuele Rodolà, Jonathan Masci, Andrea Torsello, and Michael M. Bronstein. Matching deformable objects in clutter. In *3DV*, 2016. 8
- [13] James Coughlan, Alan Yuille, Camper English, and Dan Snow. Efficient Deformable Template Detection and Localization without User Initialization. *Computer Vision and Image Understanding*, 78(3):303–319, June 2000. 2
- [14] Nicolas Donati, Abhishek Sharma, and Maks Ovsjanikov. Deep Geometric Functional Maps: Robust Feature Learning for Shape Correspondence. In *CVPR*, 2020. 3
- [15] Nadav Dym and Shahar Ziv Kovalsky. Linearly converging quasi branch and bound algorithms for global rigid registration. In *ICCV*, 2019. 2
- [16] Nadav Dym, Haggai Maron, and Yaron Lipman. DS++: A flexible, scalable and provably tight relaxation for matching problems. *ACM Transactions on Graphics*, 36(6):1–14, Nov. 2017. 2
- [17] Jonathan Eckstein, Dimitri P Bertsekas, et al. An alternating direction method for linear programming. 1990. 4
- [18] Marvin Eisenberger, Zorah Lahner, and Daniel Cremers. Smooth Shells: Multi-Scale Shape Registration With Functional Maps. In *CVPR*, 2020. 2, 6, 7
- [19] Marvin Eisenberger, David Novotny, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. NeuroMorph: Unsupervised Shape Interpolation and Correspondence in One Go. In *CVPR*, 2021. 3
- [20] Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep Shells: Unsupervised Shape Correspondence with Optimal Transport. In *NeurIPS*, 2020. 3
- [21] Raúl San José Estépar, Anders Brun, and Carl-Fredrik Westin. Robust Generalized Total Least Squares Iterative Closest Point Registration. In *MICCAI*, 2004. 2
- [22] Danielle Ezuz and Mirela Ben-Chen. Deblurring and Denoising of Maps between Shapes. *Computer Graphics Forum*, 36(5):165–174, Aug. 2017. 2
- [23] Danielle Ezuz, Behrend Heeren, Omri Azencot, Martin Rumpf, and Mirela Ben-Chen. Elastic Correspondence between Triangle Meshes. *Computer Graphics Forum*, 38(2):121–134, 2019. 2
- [24] Pedro F Felzenszwalb. Representation and Detection of Deformable Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, page 32, 2005. 2
- [25] Fajwel Fogel, Rodolphe Jenatton, Francis Bach, and Alexandre d’Aspremont. Convex Relaxations for Permutation Problems. *SIAM Journal on Matrix Analysis and Applications*, 36(4):1465–1488, Jan. 2015. 2
- [26] Maolin Gao, Zorah Lahner, Johan Thunberg, Daniel Cremers, and Florian Bernard. Isometric Multi-Shape Matching. In *CVPR*, 2021. 2
- [27] Daniela Giorgi, Silvia Biasotti, and Laura Paraboschi. SHape RETrieval Contest 2007: Watertight Models Track. page 25, 2007. 6
- [28] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, and Mathieu Aubry. Unsupervised cycle-consistent deformation for shape matching. *arXiv:1907.03165*, July 2019. 3
- [29] Gurobi Optimization, LLC. Gurobi optimizer reference manual, 2021. 4, 6
- [30] Oshri Halimi, Or Litany, Emanuele Rodola Rodola, Alex M. Bronstein, and Ron Kimmel. Unsupervised Learning of Dense Shape Correspondence. In *CVPR*, 2019. 3
- [31] Benjamin Holzs Schuh, Zorah Lahner, and Daniel Cremers. Simulated Annealing for 3D Shape Correspondence. In *3DV*, 2020. 2
- [32] Ling Hu, Qinsong Li, Shengjun Liu, and Xinru Liu. Efficient Deformable Shape Correspondence via Multiscale Spectral Manifold Wavelets Preservation. In *CVPR*, 2021. 2
- [33] Ruqi Huang, Jing Ren, Peter Wonka, and Maks Ovsjanikov. Consistent zoomout: Efficient spectral map synchronization. In *Computer Graphics Forum*, volume 39, pages 265–278, 2020. 2
- [34] Vladimir G. Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. In *Transactions on Graphics (TOG)*, 2011. 7
- [35] 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
- [36] Zorah Lahner, Emanuele Rodola, Frank R Schmidt, Michael M Bronstein, and Daniel Cremers. Efficient globally optimal 2d-to-3d deformable shape matching. In *CVPR*, 2016. 2

- [37] Jan-Hendrik Lange and Paul Swoboda. Efficient Message Passing for 0–1 ILPs with Binary Decision Diagrams. In *ICML*, 2021. [5](#)
- [38] D. Khue Le-Huu and Nikos Paragios. Alternating Direction Graph Matching. In *CVPR*, 2017. [2](#)
- [39] Marius Leordeanu, Martial Hebert, and Rahul Sukthankar. An integer projected fixed point method for graph matching and map inference. In *NIPS*, 2009. [2](#)
- [40] Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep Functional Maps: Structured Prediction for Dense Shape Correspondence. In *ICCV*, 2017. [3](#), [8](#)
- [41] Or Litany, Emanuele Rodolà, Alexander M Bronstein, Michael M Bronstein, and Daniel Cremers. Non-rigid puzzles. In *Computer Graphics Forum*, volume 35, pages 135–143, 2016. [8](#)
- [42] Manish Mandad, David Cohen-Steiner, Leif Kobbelt, Pierre Alliez, and Mathieu Desbrun. Variance-minimizing transport plans for inter-surface mapping. *ACM Transactions on Graphics*, 36(4):1–14, July 2017. [2](#)
- [43] Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, and Yaron Lipman. Point registration via efficient convex relaxation. *ACM Transactions on Graphics*, 35(4):1–12, July 2016. [2](#)
- [44] Simone Melzi, Riccardo Marin, Emanuele Rodolà, Umberto Castellani, Jing Ren, Adrien Poulencard, Peter Wonka, and Maks Ovsjanikov. Shrec 2019: Matching humans with different connectivity. In *Eurographics Workshop on 3D Object Retrieval*, 2019. [6](#)
- [45] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019. [2](#), [7](#), [8](#)
- [46] Facundo Mémoli and Guillermo Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics*, 5(3):313–347, 2005. [2](#)
- [47] James Munkres. Algorithms for the Assignment and Transportation Problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, Mar. 1957. [1](#)
- [48] Andriy Myronenko and Xubo Song. Point Set Registration: Coherent Point Drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, Dec. 2010. [2](#)
- [49] Carl Olsson, Fredrik Kahl, and Magnus Oskarsson. Branch-and-Bound Methods for Euclidean Registration Problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):783–794, May 2009. [2](#)
- [50] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):1–11, 2012. [2](#), [7](#)
- [51] Panos M Pardalos, Franz Rendl, and Henry Wolkowicz. The Quadratic Assignment Problem - A Survey and Recent Developments. *DIMACS Series in Discrete Mathematics*, 1993. [1](#), [2](#)
- [52] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, volume 30, 2017. [3](#)
- [53] Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Discrete Optimization for Shape Matching. *Computer Graphics Forum*, 40(5):81–96, Aug. 2021. [2](#), [6](#), [7](#)
- [54] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured Regularization of Functional Map Computations. *Computer Graphics Forum*, 38(5):39–53, Aug. 2019. [2](#)
- [55] Jing Ren, Adrien Poulencard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics*, 37(6):1–16, Jan. 2018. [2](#), [7](#)
- [56] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial Functional Correspondence: Partial Functional Correspondence. *Computer Graphics Forum*, 36(1):222–236, Jan. 2017. [2](#), [6](#), [8](#)
- [57] Emanuele Rodola, Samuel Rota Buló, Thomas Windheuser, Matthias Vestner, and Daniel Cremers. Dense Non-rigid Shape Correspondence Using Random Forests. In *CVPR*, 2014. [3](#), [6](#)
- [58] Jean-Michel Roufousse, Abhishek Sharma, and Maks Ovsjanikov. Unsupervised deep learning for structured shape matching. In *ICCV*, 2019. [3](#)
- [59] Yusuf Sahillioğlu. Recent advances in shape correspondence. *The Visual Computer*, 36(8):1705–1721, 2020. [2](#)
- [60] F. R. Schmidt, T. Windheuser, U. Schlickewei, and D. Cremers. Dense elastic 3d shape matching, 2014. [2](#), [4](#)
- [61] Peter H. Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, Mar. 1966. [2](#)
- [62] Robert W. Sumner and Jovan Popović. Deformation transfer for triangle meshes. *ACM Trans. Graph.*, 23(3):399–405, Aug. 2004. [2](#)
- [63] Paul Swoboda, Carsten Rother, Hassan Abu Alhaija, Dagmar Kainmuller, and Bogdan Savchynskyy. A study of lagrangean decompositions and dual ascent solvers for graph matching. In *CVPR*, 2017. [2](#)
- [64] Oliver van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A Survey on Shape Correspondence: O. van Kaick et al. / A Survey on Shape Correspondence. *Computer Graphics Forum*, 30(6):1681–1707, Sept. 2011. [2](#)
- [65] Matthias Vestner, Zorah Lahner, 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](#), [7](#), [8](#)
- [66] Thomas Windheuser, Ulrich Schlickewei, Frank R. Schmidt, and Daniel Cremers. Geometrically consistent elastic matching of 3D shapes: A linear programming solution. In *ICCV*, 2011. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [67] Thomas Windheuser, Ulrich Schlickewei, Frank R. Schmidt, and Daniel Cremers. Large-Scale Integer Linear Programming for Orientation Preserving 3D Shape Matching. *Computer Graphics Forum*, 30(5):1471–1480, Aug. 2011. [1](#), [2](#), [4](#), [6](#)

- [68] Rui Xiang, Rongjie Lai, and Hongkai Zhao. A Dual Iterative Refinement Method for Non-Rigid Shape Matching. In *CVPR*, 2021. [2](#), [8](#)
- [69] Yiming Zeng, Yue Qian, Zhiyu Zhu, Junhui Hou, Hui Yuan, and Ying He. CorrNet3D: Unsupervised End-to-End Learning of Dense Correspondence for 3D Point Clouds. In *CVPR*, 2021. [3](#)
- [70] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *CVPR*, 2017. [6](#)