

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/157089/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Yao, Yuxin, Deng, Bailin , Xu, Weiwei and Zhang, Juyong 2023. Fast and robust non-rigid registration using accelerated majorization-minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45 (8) , pp. 9681-9698.
10.1109/TPAMI.2023.3247603

Publishers page: <https://doi.org/10.1109/TPAMI.2023.3247603>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Fast and Robust Non-Rigid Registration Using Accelerated Majorization-Minimization

Yuxin Yao, Bailin Deng, *Member, IEEE*, Weiwei Xu, *Member, IEEE*, Juyong Zhang[†], *Member, IEEE*

Abstract—Non-rigid 3D registration, which deforms a source 3D shape in a non-rigid way to align with a target 3D shape, is a classical problem in computer vision. Such problems can be challenging because of imperfect data (noise, outliers and partial overlap) and high degrees of freedom. Existing methods typically adopt the ℓ_p type robust norm to measure the alignment error and regularize the smoothness of deformation, and use a proximal algorithm to solve the resulting non-smooth optimization problem. However, the slow convergence of such algorithms limits their wide applications. In this paper, we propose a formulation for robust non-rigid registration based on a globally smooth robust norm for alignment and regularization, which can effectively handle outliers and partial overlaps. The problem is solved using the majorization-minimization algorithm, which reduces each iteration to a convex quadratic problem with a closed-form solution. We further apply Anderson acceleration to speed up the convergence of the solver, enabling the solver to run efficiently on devices with limited compute capability. Extensive experiments demonstrate the effectiveness of our method for non-rigid alignment between two shapes with outliers and partial overlaps, with quantitative evaluation showing that it outperforms state-of-the-art methods in terms of registration accuracy and computational speed. The source code is available at https://github.com/yaoyx689/AMM_NRR.

Index Terms—Non-rigid registration, Robust Estimator, Welsch’s function, Anderson acceleration.



1 INTRODUCTION

WITH the popularity of depth acquisition devices such as Kinect, PrimeSense and depth sensors on smartphones, techniques for 3D object tracking and reconstruction from point clouds have enabled various applications. Non-rigid registration is a fundamental problem for such techniques, especially for reconstructing dynamic objects. Since depth maps obtained from structured light or time-of-flight cameras often contain outliers and holes, a robust non-rigid registration algorithm is needed to handle such data. Moreover, real-time applications require high computational efficiency for non-rigid registration.

Given a source surface and a target surface, each represented as a point cloud or a mesh, non-rigid registration aims to find a deformation field for the source surface to align it with the target surface. This problem is typically solved via optimization. The objective energy often includes alignment terms that measure the deviation between the target surface and deformed source surface, as well as regularization terms that enforce the smoothness of the deformation field. Many existing methods formulate these terms using the ℓ_2 -norm, which penalizes alignment and smoothness errors across the whole surface [1], [2], [3]. On the other hand, the ground-truth alignment may induce large errors for these terms in some local regions due to noise, outliers, partial overlaps, or articulated motions. The ℓ_2

formulations can inhibit such large localized errors and lead to erroneous alignment. To improve the alignment accuracy, recent works have utilized sparsity-promoting norms for these terms, such as the ℓ_1 -norm [4], [5], [6] and the ℓ_0 -norm [7]. The sparsity optimization enforces small error values on most parts of the surface while allowing for large errors in some local regions, improving the robustness of the registration process. However, the resulting optimization problem can be non-smooth and more challenging to solve. Existing methods often use proximal algorithms such as the alternating direction method of multipliers (ADMM), which can suffer from slow convergence to high-accuracy solutions [8].

In this paper, we propose a new approach to robust non-rigid registration with fast convergence. The key idea is to enforce sparsity using Welsch’s function [9], which has been utilized for robust filtering of images [10] and meshes [11], as well as robust rigid registration [12]. We formulate an optimization that applies Welsch’s function to both the alignment term and the regularization term. Unlike the ℓ_p -norms, Welsch’s function is globally smooth and avoids non-smooth optimization. We solve the optimization problem using the majorization-minimization (MM) algorithm [13]. It iteratively constructs a surrogate function for the objective energy based on the current variable values and minimizes the surrogate function to update the variables, and is guaranteed to converge to a local minimum. This leads to a convex quadratic problem in each iteration, which can be solved efficiently in closed form. To speed up the convergence of the MM algorithm, we regard it as a fixed-point iteration and apply Anderson acceleration [14], [15], a well-established technique to accelerate the convergence of fixed-point iterations. Experimental results verify the robustness of our method as well as its superior performance

- Y. Yao and J. Zhang are with School of Mathematical Sciences, University of Science and Technology of China.
- B. Deng is with School of Computer Science and Informatics, Cardiff University.
- W. Xu is with State Key Lab of CAD & CG, Department of Computer science, Zhejiang University.

[†]Corresponding author. Email: juyong@ustc.edu.cn.

compared to existing robust registration approaches.

In summary, our main contributions include:

- We formulate an optimization problem for non-rigid registration, using Welsch’s function to induce sparsity for alignment error and transformation smoothness. The proposed formulation effectively improves the robustness and accuracy of the results.
- We propose an MM algorithm to solve the optimization problem, which only involves minimizing a convex quadratic surrogate function in each iteration; we also apply Anderson acceleration to speed up its convergence. The simple form of the surrogate function, together with the effectiveness of Anderson acceleration, helps to improve the computational efficiency of robust non-rigid registration compared to existing approaches.

A preliminary version of this work appeared in [16]. The main additions in the current version include: (1) a more robust method for the deformation graph construction; (2) a more efficient optimization solver that combines the MM algorithm with Anderson acceleration; (3) a more extensive evaluation of the performance.

2 RELATED WORK

Non-rigid registration has been extensively studied in computer vision, medical image processing and computer graphics. The reader can refer to the surveys in [17] and [18] for a more complete overview of past research. In the following, we focus on works that are closely related to our method.

Non-rigid Registration Models. Various optimization-based methods perform non-rigid registration by minimizing the alignment error, often in conjunction with other constraints. Extending the classical iterative closest point (ICP) algorithm for rigid registration, Amberg et al. [1] proposed a non-rigid ICP algorithm that incrementally deforms the source model to align with the target model. Li et al. [2] adopted an embedded deformation approach [19] to express a non-rigid deformation using a deformation graph, and simultaneously optimized the correspondence between source and target scans, the confidence weights for the correspondence, and a warping field that aligns the source with the target. In addition, various geometric constraint terms have been introduced to preserve the local shape of the source model during deformation. Examples include preservation of Laplacian operators [20], [21], [22], as-conformal-as-possible deformations [23], and as-rigid-as-possible deformations [24].

Other methods tackle the registration problem from a statistical perspective. By treating point cloud alignment as a probability density estimation problem, Myronenko et al. [25] proposed the Coherent Point Drift (CPD) algorithm, which encourages displacement vectors to point in similar directions to improve the coherence of the transformation. Hirose [26] reformulated CPD in a Bayesian setting to improve robustness with a better guarantee of convergence. Hontani et al. [27] incorporated a statistical shape model and a noise model into the non-rigid ICP framework, and detected outliers based on their sparsity. Jian et al. [28], [29] represented each point set as a mixture of Gaussians and

treated point set registration as a problem of aligning two mixtures. Also with a statistical framework, Wand et al. [30] used a meshless deformation model to perform the pairwise alignment. Ma et al. [31] proposed an L_2E estimator to build more reliable sparse and dense correspondences. These methods typically require calculations for each possible pair of points from the source and the target, which is time-consuming for problems with a large number of points.

Some dynamic reconstruction and tracking frameworks involve non-rigid registration between a template model and the live frames. The registration is performed by alternately updating the correspondences and solving the transformation on the GPU to produce real-time results [32], [33], [34], [35], [36]. For such applications, the non-rigid registration tasks are usually less challenging due to the small deviation between the source and target models, and real-time performance is favored over high-precision solutions.

When the transformation between the source and target models is nearly isometric, some approaches utilize the geodesic distance of the surface as an invariant metric to directly optimize the global correspondences [37], [38], [39]. However, these methods may not be suitable for noisy mesh surfaces, as the noise can distort the geodesic distance and lead to erroneous alignment.

With the growing popularity of deep learning in computer graphics and vision tasks, it has also been adopted for non-rigid registration problems. Shimada et al. [40] represented the source and target models using regular voxel grids, and learned the displacement field using datasets of deformable objects with known correspondence. For non-rigid 3D reconstruction, Božič et al. [41] extracted features from the input RGB-D frames and matched them to obtain sparse correspondences, and used these initial correspondences to improve the accuracy of registration. For non-rigid tracking, Li et al. [42] proposed an alignment term based on deep features learned through CNN to improve the robustness of the optimization formulation, and introduced a learning-based preconditioner to speed up the convergence of the PCG method used in the optimization solver. Božič et al. [43] combined a dense correspondences prediction module based on an optical flow network of RGB images, and a differentiable alignment module to achieve more accurate real-time non-rigid tracking. These methods improve the robustness and/or speed compared with optimization-based methods, by training on carefully constructed datasets. Feng et al. [44] represented the non-rigid transformation as a point-wise combination of several rigid transformations. They proposed an unsupervised framework to learn the alignment, using a differentiable loss that measures similarity between 3D shapes via their depth images from multiple views. However, this method cannot handle partial overlaps.

Robust Metrics. Since the data generated from a 3D acquisition process often contains noise, outliers and partial overlaps, various robust metrics have been applied in optimization-based approaches to better handle such cases. Unlike the classical ℓ_2 error metrics that require the error to be small across the whole shape, such robust metrics promote the sparsity of the error and allow for large errors in some local regions to accommodate data imperfections. Yang et al. [4] introduced an ℓ_1 term for the smoothness of the

transformation, which is suitable for models with articulated motions. Li et al. [5] additionally utilized an ℓ_1 -norm for the alignment error, to improve robustness against noises and outliers. Guo et al. [7] proposed an ℓ_0 -based motion regularizer for robust motion tracking and surface reconstruction. Zampogiannis et al. [45] applied the Huber- L_1 loss to the difference between the transformation parameters at neighboring deformation graph nodes, to promote piecewise smooth deformation fields. Other works used the Geman-McClure function for the non-rigid alignment term [35], [46], [47]. In this paper, we adopt Welsch’s function for robust non-rigid registration, which results in a smooth optimization problem that allows for an efficient numerical solver.

Anderson Acceleration. Anderson acceleration [14] was initially designed for the iterative solution of nonlinear integral equations. In recent years, it has been shown effective in accelerating the convergence of fixed-point iterations [15], [48], [49]. In each iteration, Anderson acceleration utilizes m previous iterates to construct a new iterate that converges faster, which can be considered as a quasi-Newton method for finding a root of the residual function [50]. It has been adopted to accelerate first-order solvers in computer graphics [51], [52], [53], as well as ICP-based methods for rigid registration [12], [54]. In this paper, we apply Anderson acceleration to speed up the convergence of our non-rigid registration solver.

3 MODEL

Given a source surface and a target surface, non-rigid registration computes a deformation field on the source surface to align it with the target surface. We assume that the source surface is represented with a set of sample points $\mathcal{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_{|\mathcal{V}|} \in \mathbb{R}^3\}$ connected by a set of edges \mathcal{E} that define their neighboring relations, and $\mathcal{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_{|\mathcal{U}|} \in \mathbb{R}^3\}$ is a set of sample points on the target surface. In this paper, we focus on the case where the source surface is a triangle mesh, so that \mathcal{V} and \mathcal{E} are the mesh vertices and edges, respectively. However, the formulation is also applicable when the input surface is a point cloud and \mathcal{V} are the points: in this case, the edge set \mathcal{E} can be constructed by connecting each point with its nearest neighbors. Following [19], we define the deformation field using an embedded deformation graph $\mathcal{G} = (\mathcal{V}_{\mathcal{G}}, \mathcal{E}_{\mathcal{G}})$, where $\mathcal{V}_{\mathcal{G}} = \{\mathbf{p}_1, \dots, \mathbf{p}_{|\mathcal{V}_{\mathcal{G}}|} \in \mathbb{R}^3\} \subset \mathcal{V}$ is the set of graph nodes that lie on the source surface, and $\mathcal{E}_{\mathcal{G}}$ is the set of edges. Each graph node is associated with an affine transformation $\mathbf{X}_j = [\mathbf{A}_j, \mathbf{t}_j]$, where $\mathbf{A}_j \in \mathbb{R}^{3 \times 3}$ and $\mathbf{t}_j \in \mathbb{R}^3$. In the following, we use

$$\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_{|\mathcal{V}_{\mathcal{G}}|}] \in \mathbb{R}^{4|\mathcal{V}_{\mathcal{G}}| \times 3}$$

to denote the concatenation of all \mathbf{X}_j . Using the deformation graph, each sample point \mathbf{v}_i on the source surface is influenced by a set $\mathcal{I}(\mathbf{v}_i)$ of nearby deformation graph nodes:

$$\mathcal{I}(\mathbf{v}_i) = \{\mathbf{p}_j \in \mathcal{V}_{\mathcal{G}} \mid D(\mathbf{v}_i, \mathbf{p}_j) < R\}, \quad (1)$$

where $D(\cdot, \cdot)$ denotes the geodesic distance on the source surface, and R is a given radius parameter. The new position of the sample point \mathbf{v}_i is computed as

$$\hat{\mathbf{v}}_i = \sum_{\mathbf{p}_j \in \mathcal{I}(\mathbf{v}_i)} w_{ij} \cdot (\mathbf{A}_j(\mathbf{v}_i - \mathbf{p}_j) + \mathbf{p}_j + \mathbf{t}_j), \quad (2)$$

where w_{ij} is a distance-dependent normalized weight [3]:

$$w_{ij} = \frac{(1 - D^2(\mathbf{v}_i, \mathbf{p}_j)/R^2)^3}{\sum_{\mathbf{p}_k \in \mathcal{I}(\mathbf{v}_i)} (1 - D^2(\mathbf{v}_i, \mathbf{p}_k)/R^2)^3}. \quad (3)$$

From the given source and target surfaces, we first construct a deformation graph, and then optimize the transformations associated with the graph nodes to deform the source surface and align it with the target surface. Compared with existing non-rigid registration methods that define the transformation field on each source vertex such as [5], the use of a deformation graph can effectively reduce the number of variables for the optimization problem and significantly reduce the computational time. The remainder of this section explains each step of our methods in detail.

3.1 Construction of deformation graph

To avoid excessive degrees of freedom for the deformation, we would like to reduce the number of deformation graph nodes while ensuring that for each source point \mathbf{v}_i the set of influencing nodes $\mathcal{I}(\mathbf{v}_i)$ (defined in Eq. (1)) is not empty. To this end, we simultaneously construct the set $\mathcal{V}_{\mathcal{G}}$ of deformation graph nodes and the set $\mathcal{I}(\mathbf{v}_i)$ of influencing nodes for each point as follows. We first initialize $\mathcal{V}_{\mathcal{G}}$ and all $\mathcal{I}(\mathbf{v}_i)$ to empty sets. Then we perform PCA on all source points $\{\mathbf{v}_i\}$, and sort them based on their projections onto the axis corresponding to the largest eigenvalue of the covariance matrix. We go through all the points according to this order, and add a source point \mathbf{v}_j to the deformation graph nodes $\mathcal{V}_{\mathcal{G}}$ if either $\mathcal{V}_{\mathcal{G}}$ or $\mathcal{I}(\mathbf{v}_j)$ is empty. Once a point \mathbf{v}_j is added to the deformation graph nodes, we also add \mathbf{v}_j to the influencing node set for any point \mathbf{v}_i that is within geodesic distance R from \mathbf{v}_j (i.e., $D(\mathbf{v}_i, \mathbf{v}_j) < R$). To locate all such points efficiently, we first start from \mathbf{v}_j and perform a breadth-first search on the surface to identify the largest neighborhood $\mathcal{B}(\mathbf{v}_j)$ of points whose Euclidean distance to \mathbf{v}_j is smaller than R . Since the Euclidean distance is no larger than the geodesic distance, any point with a geodesic distance smaller than R must belong to $\mathcal{B}(\mathbf{v}_j)$. Hence we compute the geodesic distance from \mathbf{v}_j to each point $\mathbf{q} \in \mathcal{B}(\mathbf{v}_j)$, and add \mathbf{v}_j to $\mathcal{I}(\mathbf{q})$ if $D(\mathbf{v}_j, \mathbf{q}) < R$. To compute the geodesic distance, we construct a sub-mesh consisting of all faces on the source surface mesh that are incident with the vertices from $\mathcal{B}(\mathbf{v}_j)$, and apply the exact geodesic algorithm of [55] (VTP) on this sub-mesh. We use VTP for the computation because it can produce accurate results while being efficient on small meshes. After the set of deformation graph nodes is constructed, we connect two nodes with an edge if there exists a point that is influenced by both of them. Algorithm 1 summarizes our method for constructing the deformation graph. Note that in the preliminary version [16] of this paper, the deformation graph construction method needs to compute the geodesic distance from each graph node to all the source points; in comparison, our new method only requires the geodesic distance to the source points within the sphere centered at the graph node and with a radius R , which involves a significantly smaller number of geodesic distance values. This allows us to compute the exact geodesic distance using VTP to achieve more robust results, whereas the method in [16] needs to use the fast marching method (FMM) to compute an approximate geodesic distance

Algorithm 1: Construction of the deformation graph and the influencing node sets.

Input: The source surface mesh with vertex set $\{\mathcal{V}\}$; the radius parameter R .

Result: The deformation graph $\{\mathcal{V}_G, \mathcal{E}_G\}$; The set of influencing nodes $\mathcal{I}(\mathbf{v}_i)$ for each $\mathbf{v}_i \in \mathcal{V}$.

```

1 Initialize  $\mathcal{V}_G = \emptyset, \mathcal{E}_G = \emptyset$ , and  $\mathcal{I}(\mathbf{v}_i) = \emptyset$  for all
   $\mathbf{v}_i \in \mathcal{V}$ ;
2 Perform PCA on all points in  $\mathcal{V}$  to obtain the axis  $\mathbf{a}$ 
  corresponding to the largest eigenvalue of the
  covariance matrix;
3 Let  $\mathbf{v}_{\sigma(1)}, \mathbf{v}_{\sigma(2)}, \dots, \mathbf{v}_{\sigma(|\mathcal{V}|)}$  be the ordering of points
  in  $\mathcal{V}$  according to their projections to  $\mathbf{a}$ ;
  // Construct deformation graph nodes
  and influencing node sets
4 for  $i = 1, 2, \dots, |\mathcal{V}|$  do
5   if  $\mathcal{V}_G$  is empty OR  $\mathcal{I}(\mathbf{v}_{\sigma(i)})$  is empty then
6     Add  $\mathbf{v}_{\sigma(i)}$  to  $\mathcal{V}_G$ ;
7     Perform breadth-first search from  $\mathbf{v}_{\sigma(i)}$  to find
      the largest neighborhood  $\mathcal{B}(\mathbf{v}_{\sigma(i)})$  where
       $\|\mathbf{v}_{\sigma(i)} - \mathbf{q}\| < R \forall \mathbf{q} \in \mathcal{B}(\mathbf{v}_{\sigma(i)})$ ;
8     Use VTP [55] to compute the geodesic distance
       $D(\mathbf{v}_{\sigma(i)}, \mathbf{q})$  for all  $\mathbf{q} \in \mathcal{B}(\mathbf{v}_{\sigma(i)})$ ;
9     for each  $\mathbf{q} \in \mathcal{B}(\mathbf{v}_{\sigma(i)})$  do
10      if  $D(\mathbf{v}_{\sigma(i)}, \mathbf{q}) < R$  then
11        Add  $\mathbf{v}_{\sigma(i)}$  to  $\mathcal{I}(\mathbf{q})$ ;
12      end
13    end
14  end
15 end
  // Construct deformation graph edges
16 for  $\mathbf{v}_i \in \mathcal{V}$  do
17   for any  $\mathbf{p}_j, \mathbf{p}_k \in \mathcal{I}(\mathbf{v}_i)$  with  $j < k$  do
18     Add  $(\mathbf{p}_j, \mathbf{p}_k)$  to  $\mathcal{E}_G$ ;
19   end
20 end

```

to avoid excessive computational costs. Fig. 1 shows some examples of deformation graphs constructed using our new method and the method from [16] and their computational time, where the FMM geodesic distance is computed using an open-source implementation¹. We can see that our method is notably faster when using the same radius parameter.

Remark. An alternative way of constructing the deformation graph nodes is to repeatedly sample the source points and add them to the graph using farthest point sampling [56], until the geodesic distance between graph nodes and the farthest point is smaller than the radius parameter R or the number of graph nodes reaches a certain threshold. However, this would require computing geodesic distance between faraway points on the surface, which tends to be much slower than our approach where the geodesic distance is only evaluated between nearby points.

3.2 Problem formulation

Similar to existing non-rigid registration methods, we would like to align the source and target surfaces using a smooth

¹<https://github.com/sywe1/geodesic-computation>




Fig. 1. Comparison of deformation graph construction by [16] and our method on “crane” (left) and “march2” (right) datasets of [57].

transformation field that is locally close to rigid deformation. Therefore, we minimize the following target function to determine the transformation \mathbf{X} that is defined over the deformation graph nodes:

$$E(\mathbf{X}) = E_{\text{align}}(\mathbf{X}) + \alpha E_{\text{reg}}(\mathbf{X}) + \beta E_{\text{rot}}(\mathbf{X}). \quad (4)$$

Here E_{align} measures the alignment error, E_{reg} is a smoothness regularization term for the transformation field, and E_{rot} penalizes the deviation between the transformation \mathbf{X}_j on each node and a rigid transformation. α and β are user-specified weights. The definition for each term is explained below.

Alignment term. For each transformed source point $\hat{\mathbf{v}}_i$, we find its closest point $\mathbf{u}_{\rho(i)}$ on the target surface as the corresponding point:

$$\mathbf{u}_{\rho(i)} := \arg \min_{\mathbf{u} \in \mathcal{U}} \|\mathbf{u} - \hat{\mathbf{v}}_i\|.$$

Traditional methods such as [1] measure the alignment error using the sum of squared distance from all the transformed source points to their corresponding points. However, outliers and partial overlaps can result in large distances between some point pairs for the ground-truth




Fig. 2. Left: function ψ_ν with different values of the parameter ν . Right: different surrogate functions for the function ψ_ν with $\nu = 1$.

transformation, which can lead to a large value of such ℓ_2 alignment measures and prevent the optimization from finding the ground-truth alignment. Recently, robust metrics such as the ℓ_p -norm ($p \leq 1$) [4], [5], [7], [58] have been utilized instead for the alignment error. It is less sensitive to noises and partial overlaps, since such ℓ_p -norm minimization allows for large distances at some points. However, due to the non-smoothness of such ℓ_p -norm formulations, their numerical minimization can be much more expensive than the ℓ_2 -norm. For example, the problem is solved in [58] with an iterative algorithm that alternately updates the point correspondence and the transformation; the transformation update step minimizes a target function that involves the ℓ_p -norm and is performed using an ADMM solver, which is time-consuming and lacks a convergence guarantee due to the non-convexity of the problem. In this paper, inspired by recent works in robust image filtering [10] and mesh processing [11], we adopt Welsch's function [9] as a robust metric and derive the following alignment error term:

$$E_{\text{align}}(\mathbf{X}) = \sum_{\mathbf{v}_i \in \mathcal{V}} \psi_{\nu_a}(\|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}\|), \quad (5)$$

where $\nu_a > 0$ is a user-specified parameter, and

$$\psi_\nu(x) = 1 - \exp\left(-\frac{x^2}{2\nu^2}\right). \quad (6)$$

As shown in Fig. 2, the function ψ_ν is monotonically increasing on $[0, +\infty)$ while being bounded from above by 1. In addition, when $\nu_a \rightarrow 0$, the error measure in Eq. (5) approaches the ℓ_0 -norm of the distance values between the corresponding points. Thus it penalizes the distance between corresponding points, without being overly sensitive to large distances caused by outliers and partial overlaps. Moreover, as shown later in Section 3.3, the use of Welsch's function enables us to design an efficient optimization solver with guaranteed convergence.

Regularization term. Ideally, the transformation induced by two neighboring deformation graph nodes \mathbf{p}_i and \mathbf{p}_j should be consistent on their overlapping influenced region. We measure such consistency at \mathbf{p}_i using the following term:

$$\mathbf{D}_{ij} = r_{ij}(\mathbf{A}_j(\mathbf{p}_i - \mathbf{p}_j) + \mathbf{p}_j + \mathbf{t}_j - (\mathbf{p}_i + \mathbf{t}_i)). \quad (7)$$

Here the term $\mathbf{A}_j(\mathbf{p}_i - \mathbf{p}_j) + \mathbf{p}_j + \mathbf{t}_j - (\mathbf{p}_i + \mathbf{t}_i)$ is the difference between transformations induced by \mathbf{X}_i and \mathbf{X}_j at \mathbf{p}_i , which is commonly used for the regularization of transformation fields defined on a deformation graph [3],

[19]. r_{ij} is a normalization weight that accounts for the distance between \mathbf{p}_i and \mathbf{p}_j :

$$r_{ij} = \frac{2|\mathcal{E}_G| \cdot \|\mathbf{p}_i - \mathbf{p}_j\|^{-1}}{\sum_{\mathbf{p}_i \in \mathcal{V}_G} \sum_{\mathbf{p}_j \in \mathcal{N}(\mathbf{p}_i)} \|\mathbf{p}_i - \mathbf{p}_j\|^{-1}}, \quad (8)$$

where $\mathcal{N}(\mathbf{p}_i)$ denotes the set of neighboring deformation graph nodes for the node \mathbf{p}_i .

For non-rigid registration, the magnitude of $\|\mathbf{D}_{ij}\|$ should be small for most parts of the deformation graph. On the other hand, the optimal transformation field may induce large magnitudes of $\|\mathbf{D}_{ij}\|$ in some local regions such as the joints of a human body due to articulated motion. Therefore, we apply the Welsch's function to $\|\mathbf{D}_{ij}\|$ to penalize its magnitude across the deformation graph while allowing for large magnitudes in some regions, resulting in the following regularization term:

$$E_{\text{reg}} = \sum_{\mathbf{p}_i \in \mathcal{V}_G} \sum_{\mathbf{p}_j \in \mathcal{N}(\mathbf{p}_i)} \psi_{\nu_r}(\|\mathbf{D}_{ij}\|), \quad (9)$$

where $\nu_r > 0$ is a user-specified parameter.

Rotation Matrix term. To preserve local surface shapes during the deformation, we would like each transformation \mathbf{X}_i to be close to a rigid transformation. We enforce this property by penalizing the difference between the affine transformation matrix \mathbf{A}_i and its closest projection onto the rotation matrix group $\mathcal{R} = \{\mathbf{R} \in \mathbb{R}^{3 \times 3} | \mathbf{R}\mathbf{R}^T = \mathbf{I}, \det(\mathbf{R}) > 0\}$, and define the term E_{rot} as

$$E_{\text{rot}} = \sum_{\mathbf{p}_i \in \mathcal{V}_G} \|\mathbf{A}_i - \text{proj}_{\mathcal{R}}(\mathbf{A}_i)\|_F^2, \quad (10)$$

where $\text{proj}_{\mathcal{R}}(\cdot)$ is the projection operator:

$$\text{proj}_{\mathcal{R}}(\mathbf{A}) = \arg \min_{\mathbf{R} \in \mathcal{R}} \|\mathbf{R} - \mathbf{A}\|.$$

3.3 Numerical optimization

The target function of the optimization problem (4) is non-linear and non-convex. Thanks to the use of Welsch's function, it can be solved efficiently using the majorization-minimization (MM) algorithm [13]. Specifically, given the variable values $\mathbf{X}^{(k)}$ in the current iteration, the MM algorithm constructs a surrogate function $E(\mathbf{X}|\mathbf{X}^{(k)})$ for the target function E such that

$$\begin{aligned} E(\mathbf{X}^{(k)}|\mathbf{X}^{(k)}) &= E(\mathbf{X}^{(k)}), \\ E(\mathbf{X}|\mathbf{X}^{(k)}) &\geq E(\mathbf{X}) \quad \forall \mathbf{X} \neq \mathbf{X}^{(k)}. \end{aligned} \quad (11)$$

Then the variables are updated by minimizing the surrogate function

$$\mathbf{X}^{(k+1)} = \arg \min_{\mathbf{X}} E(\mathbf{X}|\mathbf{X}^{(k)}). \quad (12)$$

This guarantees the decrease of the target function in each iteration, since Equations (11) and (12) imply that

$$E(\mathbf{X}^{(k+1)}) \leq E(\mathbf{X}^{(k+1)}|\mathbf{X}^{(k)}) \leq E(\mathbf{X}^{(k)}|\mathbf{X}^{(k)}) = E(\mathbf{X}^{(k)}).$$

As a result, the iterations are guaranteed to converge to a local minimum [13]. In comparison, existing solvers for minimizing non-convex ℓ_p -norms such as ADMM [58] or iteratively reweighted least squares [59] either lack a convergence guarantee or rely on additional strong assumptions to guarantee convergence. In the following, we explain the construction of the surrogate function and its numerical minimization.

3.3.1 Surrogate function

To construct the surrogate function $E(\mathbf{X}|\mathbf{X}^{(k)})$, we note that there is a convex quadratic surrogate function for the Welsch's function ψ_ν at y [10] (see Fig. 2 right):

$$\psi_\nu(x|y) = \psi_\nu(y) + \frac{1 - \psi_\nu(y)}{2\nu^2} (x^2 - y^2). \quad (13)$$

This function bounds the Welsch's function from above, and the two function graphs touch at $x = y$. Applying this to E_{reg} in Eq. (9), we obtain a convex quadratic surrogate function

$$\begin{aligned} & E_{\text{reg}}(\mathbf{X}|\mathbf{X}^{(k)}) \\ &= \frac{1}{2\nu_r^2} \sum_{\mathbf{p}_i \in \mathcal{V}_G} \sum_{\mathbf{p}_j \in \mathcal{N}(\mathbf{p}_i)} \exp\left(-\frac{\|\mathbf{D}_{ij}^{(k)}\|^2}{2\nu_r^2}\right) \|\mathbf{D}_{ij}\|^2, \end{aligned} \quad (14)$$

where $\mathbf{D}_{ij}^{(k)}$ is evaluated using Eq. (7) at $\mathbf{X}^{(k)}$. Note that in Eq. (14) we have ignored some constant terms as they do not affect the optimization. Similarly, we can apply Eq. (13) to E_{align} and ignore some constant terms, to obtain a surrogate function

$$\frac{1}{2\nu_a^2} \sum_{\mathbf{v}_i \in \mathcal{V}} \exp\left(-\frac{\|\hat{\mathbf{v}}_i^{(k)} - \mathbf{u}_{\rho(i)}^{(k)}\|^2}{2\nu_a^2}\right) \|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}\|^2, \quad (15)$$

where the point $\hat{\mathbf{v}}_i^{(k)}$ is calculated with Eq. (2) at $\mathbf{X}^{(k)}$, and $\mathbf{u}_{\rho(i)}^{(k)}$ is the closest point to $\hat{\mathbf{v}}_i^{(k)}$ on the target surface. However, this is not a quadratic function of the new position $\hat{\mathbf{v}}_i$, since the point $\mathbf{u}_{\rho(i)}$ depends non-linearly on $\hat{\mathbf{v}}_i$. To obtain a more simple form, we note that the term $\|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}\|^2$ has a quadratic surrogate function $\|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}^{(k)}\|^2$ at $\mathbf{X}^{(k)}$, since by definition the closest point $\mathbf{u}_{\rho(i)}$ satisfies $\|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}\|^2 \leq \|\hat{\mathbf{v}}_i - \mathbf{u}\|^2$ for all $\mathbf{u} \in \mathbb{R}^3$. Applying this to Eq. (15), we obtain the following convex quadratic surrogate function for E_{align} at $\mathbf{X}^{(k)}$ up to an additional constant:

$$\begin{aligned} & E_{\text{align}}(\mathbf{X}|\mathbf{X}^{(k)}) \\ &= \frac{1}{2\nu_a^2} \sum_{\mathbf{v}_i \in \mathcal{V}} \exp\left(-\frac{\|\hat{\mathbf{v}}_i^{(k)} - \mathbf{u}_{\rho(i)}^{(k)}\|^2}{2\nu_a^2}\right) \|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}^{(k)}\|^2, \end{aligned} \quad (16)$$

For the rotation term E_{rot} , we can similarly derive a quadratic surrogate function as

$$E_{\text{rot}}(\mathbf{X}|\mathbf{X}^{(k)}) = \sum_{\mathbf{p}_i \in \mathcal{V}_G} \|\mathbf{A}_i - \text{proj}_{\mathcal{R}}(\mathbf{A}_i^{(k)})\|. \quad (17)$$

In total, the surrogate function of $E(\mathbf{X})$ can be written as

$$\begin{aligned} & E(\mathbf{X}|\mathbf{X}^{(k)}) \\ &= E_{\text{align}}(\mathbf{X}|\mathbf{X}^{(k)}) + \alpha E_{\text{reg}}(\mathbf{X}|\mathbf{X}^{(k)}) + \beta E_{\text{rot}}(\mathbf{X}|\mathbf{X}^{(k)}). \end{aligned} \quad (18)$$

This surrogate function is minimized to obtain the updated variables $\mathbf{X}^{(k+1)}$.

3.3.2 Numerical minimization

Since $E(\mathbf{X}|\mathbf{X}^{(k)})$ is a convex quadratic function, we can minimize it by solving a linear system. To derive the system, we first rewrite $E(\mathbf{X}|\mathbf{X}^{(k)})$ in a matrix form. The term $E_{\text{align}}(\mathbf{X}|\mathbf{X}^{(k)})$ can be written as

$$E_{\text{align}}(\mathbf{X}|\mathbf{X}^{(k)}) = \|\mathbf{W}_a(\mathbf{F}\mathbf{X} - \mathbf{Q})\|_F^2, \quad (19)$$

where $\mathbf{W}_a = \text{diag}(\sqrt{w_1^a}, \dots, \sqrt{w_{|\mathcal{V}|}^a}) \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ with

$$w_i^a = \frac{1}{2\nu_a^2} \exp\left(-\frac{\|\hat{\mathbf{v}}_i^{(k)} - \mathbf{u}_{\rho(i)}^{(k)}\|^2}{2\nu_a^2}\right),$$

\mathbf{F} is a block matrix $\{\mathbf{F}_{ij}\}_{\substack{1 \leq i \leq |\mathcal{V}| \\ 1 \leq j \leq |\mathcal{V}_G|}} \in \mathbb{R}^{|\mathcal{V}| \times 4|\mathcal{V}_G|}$ with

$$\mathbf{F}_{ij} = \begin{cases} w_{ij} \cdot [\mathbf{v}_i^T - \mathbf{p}_j^T, 1] & \text{if } \mathbf{p}_j \in \mathcal{I}(\mathbf{v}_i) \\ \mathbf{0} & \text{otherwise} \end{cases},$$

and $\mathbf{Q} = [\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_{|\mathcal{V}|}]^T \in \mathbb{R}^{|\mathcal{V}| \times 3}$ with

$$\mathbf{Q}_i = \mathbf{u}_{\rho(i)}^{(k)} - \sum_{\mathbf{p}_j \in \mathcal{I}(\mathbf{v}_i)} w_{ij} \mathbf{p}_j.$$

Similarly, the term $E_{\text{reg}}(\mathbf{X}|\mathbf{X}^{(k)})$ can be written as

$$E_{\text{reg}}(\mathbf{X}|\mathbf{X}^{(k)}) = \|\mathbf{W}_r(\mathbf{H}\mathbf{X} - \mathbf{Y})\|_F^2, \quad (20)$$

where the matrices $\mathbf{H} \in \mathbb{R}^{2|\mathcal{E}_G| \times 4|\mathcal{V}_G|}$ and $\mathbf{Y} \in \mathbb{R}^{2|\mathcal{E}_G| \times 3}$ encode the computation of a term \mathbf{D}_{ij} in the same row: the row in \mathbf{H} has two non-zero blocks $[r_{ij}(\mathbf{p}_i^T - \mathbf{p}_j^T), r_{ij}]$ and $[0, 0, 0, -r_{ij}]$ corresponding to \mathbf{X}_j and \mathbf{X}_i respectively, whereas the row in \mathbf{Y} has elements $[r_{ij}(\mathbf{p}_i^T - \mathbf{p}_j^T)]$. The diagonal matrix $\mathbf{W}_r \in \mathbb{R}^{2|\mathcal{E}_G| \times 2|\mathcal{E}_G|}$ has an element $\sqrt{\frac{1}{2\nu_r^2} \exp\left(-\frac{\|\mathbf{D}_{ij}^{(k)}\|^2}{2\nu_r^2}\right)}$ in the row corresponding to \mathbf{D}_{ij} . Finally, the term $E_{\text{rot}}(\mathbf{X}|\mathbf{X}^{(k)})$ can be written as

$$E_{\text{rot}}(\mathbf{X}|\mathbf{X}^{(k)}) = \|\mathbf{J}\mathbf{X} - \mathbf{Z}\|_F^2,$$

where

$$\begin{aligned} \mathbf{J} &= \text{diag}(1, 1, 1, 0, 1, 1, 1, 0, \dots, 1, 1, 1, 0) \in \mathbb{R}^{4|\mathcal{V}_G| \times 4|\mathcal{V}_G|}, \\ \mathbf{Z} &= [\text{proj}_{\mathcal{R}}(\mathbf{A}_1), \mathbf{0}, \dots, \text{proj}_{\mathcal{R}}(\mathbf{A}_{|\mathcal{V}_G|}), \mathbf{0}]^T \in \mathbb{R}^{4|\mathcal{V}_G| \times 3}. \end{aligned}$$

Using the matrix forms, the updated variable $\mathbf{X}^{(k+1)}$ is computed via the linear system

$$\mathbf{K}^{(k)} \mathbf{X}^{(k+1)} = \mathbf{B}^{(k)}, \quad (21)$$

where

$$\begin{aligned} \mathbf{K}^{(k)} &= \mathbf{F}^T \mathbf{W}_a^T \mathbf{W}_a \mathbf{F} + \alpha \mathbf{H}^T \mathbf{W}_r^T \mathbf{W}_r \mathbf{H} + \beta \mathbf{J}^T \mathbf{J}, \\ \mathbf{B}^{(k)} &= \mathbf{F}^T \mathbf{W}_a^T \mathbf{W}_a \mathbf{Q} + \alpha \mathbf{H}^T \mathbf{W}_r^T \mathbf{W}_r \mathbf{Y} + \beta \mathbf{J}^T \mathbf{Z}. \end{aligned}$$

We repeat this step until either the maximum number of iterations I_{max} is reached, or the deformed source model converges (indicated by the condition $\max_i \|\hat{\mathbf{v}}_i^{(k+1)} - \hat{\mathbf{v}}_i^{(k)}\| < \epsilon$ where ϵ is a user-specified threshold).

3.3.3 Choice of parameters ν_a and ν_r

The parameters ν_a and ν_r play an important role in the robustness of our method. The surrogate functions in Eqs. (14) and (16) show that our solver actually minimizes a quadratic error function, where the weights for the alignment error terms and the regularization terms are updated dynamically as a Gaussian function of the current error norms with variance ν_a^2 and ν_r^2 respectively. On one hand, we would like ν_a and ν_r to be sufficiently small in the final stage of the optimization, such that the Gaussian weights can effectively filter out the influence of alignment and rotation terms with large errors due to outliers, partial overlaps and articulated motions. On the other hand, fixing ν_a^2 and ν_r^2 to small values can reduce the effectiveness of the optimization,

as many terms would have large errors at the beginning of the optimization and be excluded due to their small Gaussian weights. Therefore, we gradually decrease the values of ν_a and ν_r during the optimization. Specifically, we initialize ν_a and ν_r to values $\bar{\nu}_a$ and $\bar{\nu}_r$ that are large enough to accommodate a large number of terms, and use these values ν_a, ν_r to perform optimization until convergence. Then we reduce ν_a and ν_r by half, and use the previous optimized variables as initial values to continue the optimization until convergence. We repeat this process until the ν_a reaches a user-specified lower bound $\underline{\nu}_a$. We choose $\bar{\nu}_a = \bar{d}$ and $\bar{\nu}_r = \bar{l}/\sqrt{3}$ following [12], and set $\bar{\nu}_r = 3\bar{l}$, where \bar{d} is the median value of the initial distance between each source point and its closest point on the target surface, and \bar{l} denotes the average edge length of the source model.

3.3.4 Anderson acceleration

The MM solver monotonically decreases the target function and converges quickly to an approximate solution, but it can take a long time to converge to a high-accuracy solution. To speed up the final convergence, we adopt Anderson acceleration [15], a well-established approach for accelerating fixed-point iterations. To do so, we consider our MM solver as a fixed point iteration

$$\mathbf{X}^{(k+1)} = G(\mathbf{X}^{(k)}) \quad (22)$$

with a residual function

$$F(\mathbf{X}) = G(\mathbf{X}) - \mathbf{X}.$$

If the MM solver converges to a solution \mathbf{X}^* , then \mathbf{X}^* is a fixed point of the iteration (22) and $F(\mathbf{X}^*) = \mathbf{0}$. In each iteration, Anderson acceleration uses the current iterate $\mathbf{X}^{(k)}$ and m previous iterates $\mathbf{X}^{(k-1)}, \dots, \mathbf{X}^{(k-m)}$ to compute an accelerated iterate \mathbf{X}_{AA} as an affine combination of $G(\mathbf{X}^{(k)}), G(\mathbf{X}^{(k-1)}), \dots, G(\mathbf{X}^{(k-m)})$:

$$\mathbf{X}_{AA} = G(\mathbf{X}^{(k)}) - \sum_{j=1}^m \theta_j^* (G(\mathbf{X}^{(k-j+1)}) - G(\mathbf{X}^{(k-j)})), \quad (23)$$

with the coefficients $\{\theta_j^*\}$ computed by solving a linear least-squares problem:

$$(\theta_1^*, \dots, \theta_m^*) = \arg \min_{\theta_1, \dots, \theta_m} \|F^{(k)} - \sum_{j=1}^m \theta_j (F^{(k-j+1)} - F^{(k-j)})\|^2,$$

where $F^{(k)} = F(\mathbf{G}^{(k)})$. Since the accelerated iterate \mathbf{X}_{AA} may suffer from instability and lead to stagnation, we follow the stabilization approach from [51] and only accept it as the new iterate $\mathbf{X}^{(k+1)}$ if it decreases the target function; otherwise, we revert to the original MM update $G(\mathbf{X}^{(k)})$ which is guaranteed to decrease the target function, i.e.,

$$\mathbf{X}^{(k+1)} = \begin{cases} \mathbf{X}_{AA} & \text{if } E(\mathbf{X}_{AA}) < E(\mathbf{X}^{(k)}), \\ G(\mathbf{X}^{(k)}) & \text{otherwise.} \end{cases}$$

This approach guarantees a monotonic decrease of the target function. Similar to [51], some intermediate results that are computed during the evaluation of $E(\mathbf{X}_{AA})$ — in particular the closest points $\{\mathbf{u}_{\rho(i)}\}$ and the rotation matrices $\{\text{proj}_{\mathcal{R}}(\mathbf{A}_i)\}$ — can be reused for the MM step in the next iteration if \mathbf{X}_{AA} is accepted. This helps to alleviate the computational overhead of the acceptance mechanism.

Algorithm 2: Robust non-rigid registration using Welsch's function and Anderson acceleration.

Input: $\{\mathcal{V}, \mathcal{E}\}, \mathcal{U}$: source and target models;
Result: \mathbf{X}^* converging to a local minimum of $E(\mathbf{X})$.

```

1 Initialize deformation graph  $\{\mathcal{V}_G, \mathcal{E}_G\}$  by Alg. 1.
2 Initialize  $\mathbf{X}^{(0)}$  with identity transformations;
3  $\nu_a = \bar{\nu}_a$ ;  $\nu_r = \bar{\nu}_r$ ;
4  $k = 1$ ;
5 while TRUE do
6    $k_{\text{start}} = k - 1$ ;  $E_{\text{prev}} = +\infty$ ;
7   repeat
8     Find corresponding point  $\mathbf{u}_{\rho(i)}^{(k)}$  for each  $\hat{\mathbf{v}}_i^{(k)}$ ;
9     Update the projection  $\text{proj}_{\mathcal{R}}(\mathbf{A}_j^{(k)})$ ;
10    Calculate objective function  $E$  by Eq. (4);
11    if  $E > E_{\text{prev}}$  then
12       $\mathbf{X}^{(k)} = \mathbf{X}^l$ ;
13      continue;
14    end
15     $E_{\text{prev}} = E$ ;
16    Update weight matrices  $\mathbf{W}_a$  and  $\mathbf{W}_r$ ;
17     $\mathbf{X}^l = \arg \min_{\mathbf{X}} E(\mathbf{X} | \mathbf{X}^{(k)})$  by solving Eq. (21);
18    // Anderson acceleration
19     $G^{(k)} = \mathbf{X}^l$ ;  $F^{(k)} = G^{(k)} - \mathbf{X}^{(k)}$ ;
20     $m_k = \min(k - k_{\text{start}}, m)$ ;
21     $(\theta_1^*, \dots, \theta_{m_k}^*) =$ 
22       $\arg \min \|F^{(k)} - \sum_{j=1}^{m_k} \theta_j (F^{(k-j+1)} - F^{(k-j)})\|^2$ ;
23     $\mathbf{X}^{(k+1)} = G^{(k)} - \sum_{j=1}^{m_k} \theta_j^* (G^{(k-j+1)} - G^{(k-j)})$ ;
24     $k = k + 1$ ;
25    until  $\max_i \|\hat{\mathbf{v}}_i^{(k+1)} - \hat{\mathbf{v}}_i^{(k)}\| < \epsilon$  or
26       $(k - k_{\text{start}}) == I_{\text{max}}$ ;
27    if  $\nu_a == \underline{\nu}_a$  then return  $\mathbf{T}^{(k)}$ ;
28     $\nu_a = \max(\nu_a/2, \underline{\nu}_a)$ ;  $\nu_r = \nu_r/2$ ;
29     $k = k + 1$ ;
30  end

```

Algorithm 2 summarizes our approach that combines the MM update with Anderson acceleration.

Remark. In the preliminary version [16] of the current paper, the same optimization problem is solved using a different approach, which iteratively constructs a non-quadratic surrogate function and minimizes it using an inner L-BFGS solver. One benefit of our new approach is that it involves fewer parameters for the solver. In particular, the inner L-BFGS solver in [16] requires a threshold parameter for the sufficient decrease condition used in its line search, as well as two parameters for its termination criteria. These parameters need to be chosen properly to achieve the best performance of the solver. In comparison, our new solver does not involve a line search or an inner solver, and hence does not require the tuning of such parameters. More numerical comparisons between the two approaches can be found in Section 4.3.

4 RESULTS

We conduct extensive experiments to evaluate the effectiveness of our method and compare its performance with

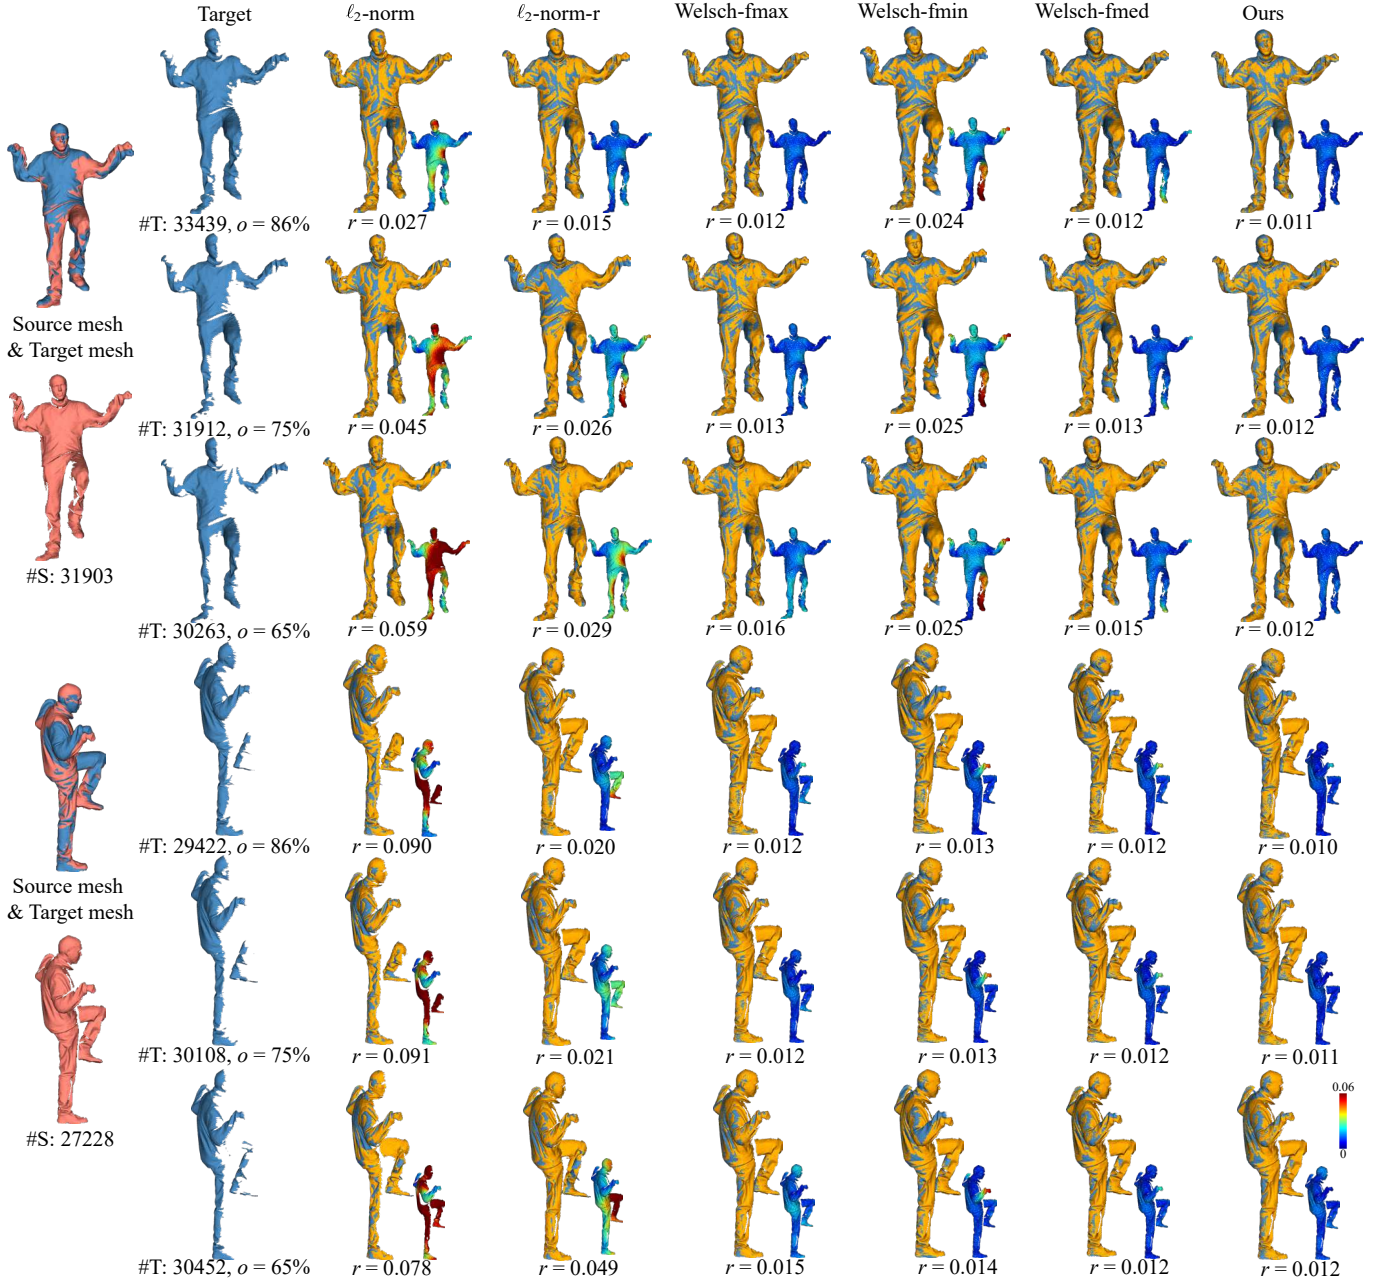


Fig. 3. Comparison on partially overlapping data constructed from the “crane” dataset [57]. We set $k_\alpha = 100$, $k_\beta = 10$ for all variants.

existing approaches.

4.1 Implementation and metrics

We implement our method in C++, using Eigen [60] for linear algebra operations and OpenMP for parallelization. Unless stated otherwise, the experiments are run on a PC with 16GB of RAM and a 6-core CPU at 3.60GHz. Each pair of surfaces are pre-processed by scaling them to achieve a unit-length diagonal for their bounding box. By default we set $\epsilon = 10^{-5}$, $I_{\max} = 100$, and the sampling radius $R = 5\bar{l}$ where \bar{l} is the average edge length of the source model. The source code of our method is available at https://github.com/yaoyx689/AMM_NRR.

To evaluate the registration accuracy of a result, we first compute the error for each source point \mathbf{v}_i using the

distance from its transformed position $\hat{\mathbf{v}}_i$ to its ground-truth corresponding point $\mathbf{u}_{\rho^*(i)}$ on the target surface:

$$D_i = \|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho^*(i)}\|. \quad (24)$$

The registration accuracy is then measured by the root mean square error (RMSE):

$$r = \sqrt{\frac{\sum_{\mathbf{v}_i \in \mathcal{V}} D_i^2}{|\mathcal{V}|}}, \quad (25)$$

where a lower value of r indicates higher accuracy. The error metrics D_i and r are measured in the original scale of the models, and are in the unit of meters unless stated otherwise. In the following, we use #S and #T to denote the number of sample points on the source surface and the target surface, respectively. We render the initial source model, the target

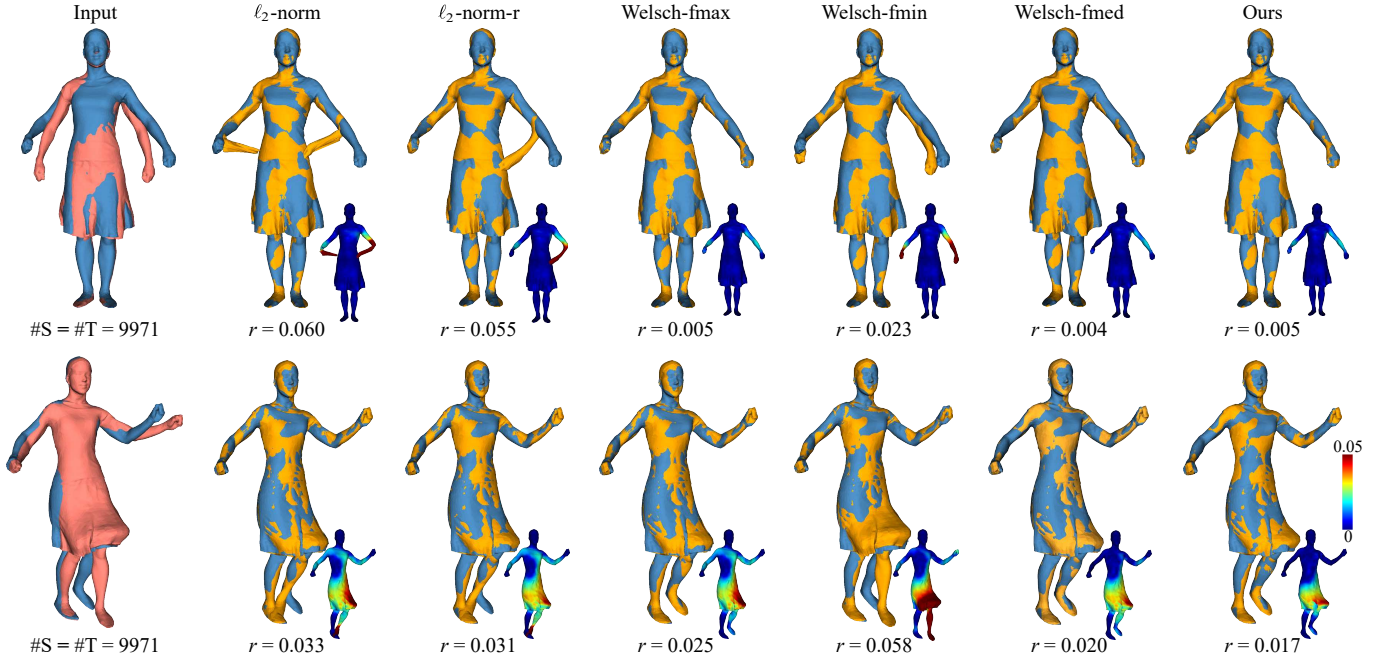


Fig. 4. Comparison on two pairs of meshes from the “samba” dataset [57]. We set $k_\alpha = 10, k_\beta = 1$ for all variants.

model, and the deformed source model in red, blue, and orange respectively, and use color-coding to visualize the point-wise error D_i across the surface.

4.2 Effectiveness of our method

Welsch’s function and parameters update. We first evaluate the effectiveness of our formulation with Welsch’s function and our dynamic update strategies for its parameter. To do so, we consider the following variants of our algorithm:

- We replace the Welsch’s functions in the target function in Eq. (4) with the squared ℓ_2 -norm, such that the target function becomes

$$E_{\ell_2} = \sum_{\mathbf{v}_i \in \mathcal{V}} w_i^d \|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}\|^2 + \alpha \sum_{\mathbf{p}_i \in \mathcal{V}_g} \sum_{\mathbf{p}_j \in \mathcal{N}(\mathbf{p}_i)} \|\mathbf{D}_{ij}\|^2 + \beta E_{\text{rot}},$$

where $\{w_i^d\}$ are weights. We consider two variants with different weighting schemes:

- All weights $\{w_i^d\}$ are set to 1. In the following examples, we refer to this variant as ℓ_2 -norm.
- The weight w_i^d is set to 1 only if $\hat{\mathbf{v}}_i$ is close enough to $\mathbf{u}_{\rho(i)}$. Otherwise, it is set to zero. To mimic the exclusion mechanism of our Welsch formulation, we use the empirical three-sigma rule for the Gaussian function and set $w_i^d = 0$ if $\|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}\| \geq \nu_a$ where ν_a is the Welsch’s function parameter used in our alignment term, i.e., the weight $w_i^{d(k)}$ at the k -th iteration is chosen as

$$w_i^{d(k)} = \begin{cases} 1 & \|\hat{\mathbf{v}}_i - \mathbf{u}_{\rho(i)}^{(k)}\| < 3\nu_a, \\ 0 & \text{others,} \end{cases}$$

We refer to this variant as ℓ_2 -norm-r.

- Instead of dynamically updating ν_a and ν_r , we fix them throughout the optimization. We consider two variants:

- Both parameters are fixed to their maximum values during our dynamic update scheme. We refer to this variant as *Welsch-fmax*.
- Both parameters are fixed to their minimum values during our dynamic update scheme. We refer to this as *Welsch-fmin*.
- Both parameters are fixed to their average values between the minimum values and the maximum values during our dynamic update scheme. We refer to this as *Welsch-fmed*.

In Fig. 3, we compare different variants on problems where the source and target models overlap partially. To this end, we choose two pairs of meshes from the “crane” dataset of [57], each consisting of two consecutive frames from the sequence. For each pair, we choose a view direction and one of the meshes, and use the visible part of the mesh from the view direction as the source model. For the other mesh, we choose three view directions different from the source model, and use its visible parts from these directions to create three target models that overlap partially with the source model. Fig. 3 presents the registration results from different variants on these models. To compute the RMSE (Eq. (25)) in this scenario, we note that the point correspondence between the complete models is known from the dataset; thus for each point on the source model we find its corresponding point on the complete mesh of the target model as the ground-truth correspondence for the RMSE evaluation. For each target model, we also evaluate its overlap ratio o with the source model as the percentage of source points whose ground-truth correspondence lies on the target model. Fig. 3 shows that the ℓ_2 -based variants result in lower accuracy than other variants, due to the sensitivity of ℓ_2 -norm to large errors, while the Welsch-based variants with fixed ν_a and ν_r are less effective in distinguishing non-overlapping points. Our Welsch-based formulation with a dynamic update of (ν_a, ν_r)




Fig. 5. Deformation graphs and registration results using our method with different values of the radius parameter R , on a pair of models from the MPI Faust dataset [61]. We set $k_\alpha = 0.1, k_\beta = 0.001$. Here $\#n$ denotes the number of the deformation graph nodes, t_i and t_o denote the computational time for the deformation graph construction and the numerical optimization respectively, and $\#i$ denotes the number of iterations required for the numerical optimization solver to converge.

achieves better accuracy than these variants.

In Fig. 4, we further compare the variants on two pairs of meshes from the “samba” dataset [57] with a relatively large deformation between the source and the target. Our approach also achieves better accuracy than the variants in this scenario, thanks to its use of Welsh’s function which helps to attenuate the influence of incorrect correspondence obtained from the closest point, as well as its dynamic parameter update strategy that helps to avoid undesirable local minima of the optimization.

Radius parameter for deformation graph. The radius parameter R for the deformation graph construction affects the number of graph nodes. It can be used to trade off the degrees of freedom and computational efficiency. Fig. 5 shows results using different values of R on a pair of models from the MPI Faust dataset [61]. We can see that smaller R leads to more nodes in the deformation graph, which provides more degrees of freedom for the deformation and allow for more accurate alignment. On the other hand, a larger number of nodes increase the number of variables for the optimization, which requires a longer computational time for the deformation graph construction and the numerical solver. In addition, having too many variables may cause

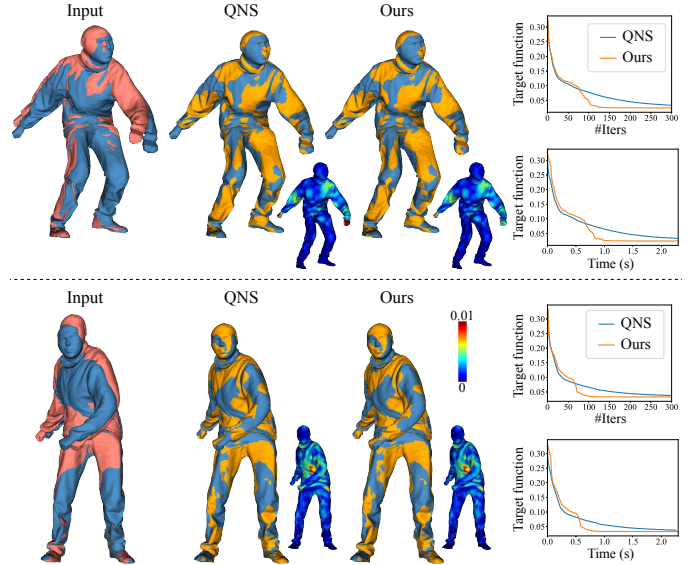


Fig. 6. Comparisons between the solver from [16] and our solver for the same optimization problem on the “jumping” dataset of [57]. We set $k_\alpha = 10, k_\beta = 100$ for both solvers.

the optimization to converge to a suboptimal solution, as it becomes more difficult to maintain the local shape (e.g., see the right leg of the result from $R = 3\bar{l}$). In our experiments, our default choice $R = 5\bar{l}$ usually achieves a good balance between accuracy and efficiency.

4.3 Comparison with the solver in [16]

A preliminary version of our method was presented in [16], where the same optimization problem is solved using a quasi-Newton approach instead. Similar to our solver, the solver in [16] constructs a surrogate for the target function and minimizes it to update the variables, with a dynamic update of the parameters (ν_a, ν_r) . However, their surrogate function is not quadratic and requires an inner L-BFGS solver for its minimization. The two solvers achieve similar results but differ in their efficiency, as shown in Fig. 6. Here we choose two pairs of adjacent frame data from the “jumping” dataset of [57], and perform registration with fixed parameters (ν_a, ν_r) for a fair comparison between the solvers for their speed of convergence. The plots of the target function show that our solver converges faster than [16]. This is partly because the surrogate function is only tight around the current variable values; the L-BFGS inner solver in [16] spends multiple iterations to minimize a loose upper bound function away from the current variable values, which is less efficient in reducing the target function. In comparison, our solver immediately updates the surrogate after each iteration to maintain a tight upper bound function for minimization, which improves efficiency. Moreover, the use of Anderson acceleration helps our solver to converge faster to the fixed point and avoid the slow decrease of target function as seen in the final phase for the solver of [16].

4.4 Comparison with existing methods

We compare our method with the ℓ_2 -based approach from [1] (N-ICP), the ℓ_1 -based approach from [5] (RPTS), and the

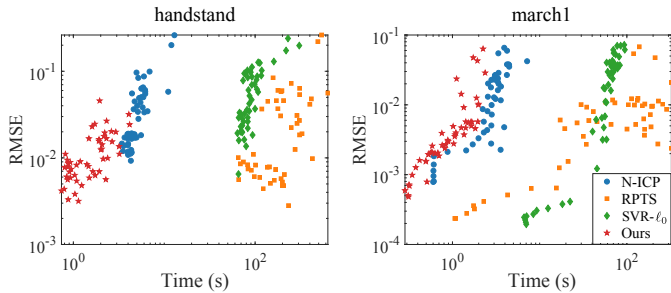


Fig. 7. Visualization of the computational time and RMSE using different methods on 50 pairs of models from the “handstand” and “march1” datasets [57].

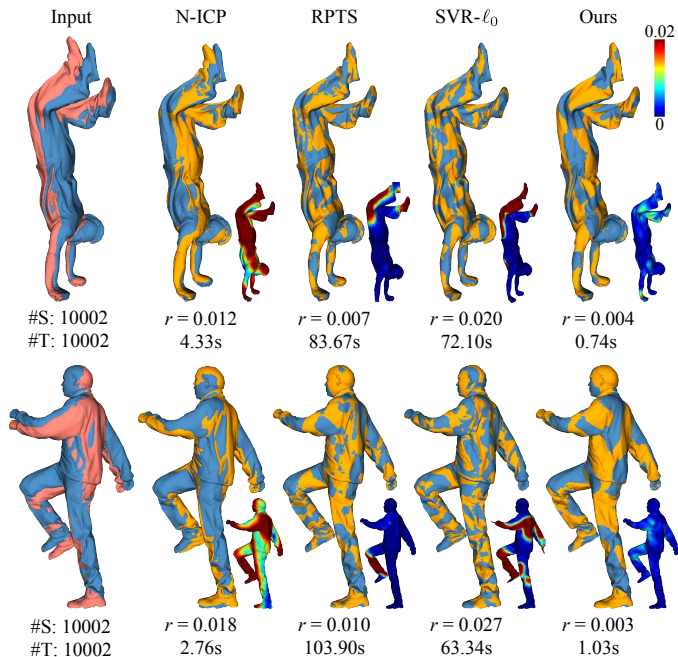


Fig. 8. Comparison on mesh pairs from the “handstand” dataset (top) and the “march1” dataset (bottom) [57]. We set $\alpha = 10$ for N-ICP, $\alpha = 10, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1$ for SVR- ℓ_0 , and $k_\alpha = 100, k_\beta = 1$ for our method.

ℓ_0 -based approach from [7] (SVR- ℓ_0). We implement N-ICP and RPTS by ourselves², while the implementation of SVR- ℓ_0 is provided by the authors. All these methods solve an optimization problem for the transformation, where the target function is a weighted sum of an alignment term, a regularization term, and (for some methods) a local rotation/rigidity term. To unify notation, for each method we set the weight for the alignment term to 1 unless stated otherwise, and use α and β to denote the weight for the regularization term and the rotation/rigidity term, respectively. For a fair comparison, we choose the best parameters for each method. For our method, to reduce sensitivity to the mesh resolution and the Welsch’s function parameters ν_a, ν_r , we set the weights α, β as

$$\alpha = k_\alpha \cdot \frac{|\mathcal{V}|}{|\mathcal{E}_G|} \cdot \frac{\nu_r^2}{\nu_a^2}, \quad \beta = k_\beta \cdot \frac{|\mathcal{V}|}{|\mathcal{V}_G|} \cdot \frac{1}{2\nu_a^2}, \quad (26)$$

²https://github.com/Juyong/Fast_RNRR

TABLE 1

Mean/median RMSE ($\times 10^{-2}$) and average computational time (s) using different methods on 50 pairs of models from the “handstand” and “march1” datasets [57]. We set $\alpha = 10$ for N-ICP, $\alpha = 100, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1$ for SVR- ℓ_0 , and $k_\alpha = 100, k_\beta = 1$ for our method.

Data		N-ICP	RPTS	SVR- ℓ_0	Ours
handstand	Mean RMSE	4.31	3.30	6.45	1.21
	Median RMSE	3.09	1.17	4.95	0.90
	Average Time	5.31	215.47	96.61	1.70
march1	Mean RMSE	1.78	0.93	2.49	0.75
	Median RMSE	1.46	0.60	2.51	0.33
	Average Time	2.86	103.12	55.36	1.06

and control α, β using the parameters k_α and k_β . For N-ICP, we ignore the parameter adjustment for the stiffness and landmark terms for a fair comparison. For SVR- ℓ_0 , we extract the ℓ_0 -optimization part and replace the corresponding points with the closest points. For both N-ICP and RPTS, we terminate the solver if it reaches 100 iterations or $\max_i \|\hat{\mathbf{v}}_i^{(k+1)} - \hat{\mathbf{v}}_i^{(k)}\| < 1 \times 10^{-3}$. RPTS also requires an inner ADMM solver to compute all affine transformations \mathbf{X} , and it is terminated when reaching 20 iterations or $\|\mathbf{X}^{(k+1)} - \mathbf{X}^{(k)}\|_F < 1 \times 10^{-3}$. For SVR- ℓ_0 , we terminate the solver when it reaches 100 iterations or $\|\mathbf{X}^{(k+1)} - \mathbf{X}^{(k)}\|_F < 1 \times 10^{-3}$. We compare the methods on the human motion datasets from [57], the MPI Faust dataset [61], the TOSCA dataset [62], the DeepDeform dataset [41], and real-world data acquired by a Kinect camera. The results are presented in detail in the following paragraphs. The iterative registration process of different optimization-based methods on some problem instances is also demonstrated in the supplementary video.

Clean data with small deformation. In Figs. 7 and 8 and Tab. 1, we test the methods on the “handstand” and “march1” datasets from [57]. Each dataset contains a sequence of meshes for the motion of a human performer. We register the i -th frame to the $(i+2)$ -th frame for each dataset. In this setting, the deformation between the source and the target models is relatively small. Fig. 7 shows the computational time and RMSEs (Eq. (25)) of 50 pairs of such models, while Fig. 8 shows the results on one example pair for each dataset. Tab. 1 further shows the mean and median values of RMSE for each dataset. The results show that overall our method can achieve higher accuracy with a shorter computational time.

Clean data with large deformation. In Tab. 2 and Fig. 9, we compare the methods on the MPI Faust dataset [61]. We select 10 subjects with nine poses, and use the first pose of each subject as the source surface and the other poses of the same subject as its targets. Due to the large difference between the source and the target models, we use the SHOT feature [63] with diffusion pruning [64] to obtain initial corresponding points. In Fig. 9, we visualize the initial correspondences using blue lines connecting the corresponding points. Fig. 9 shows a visualization of the computational time and RMSE for each pose pair using different methods, as well as the results for a particular pair on one subject. Tab. 2 further shows the average RMSE and computational time among all subjects for each pose

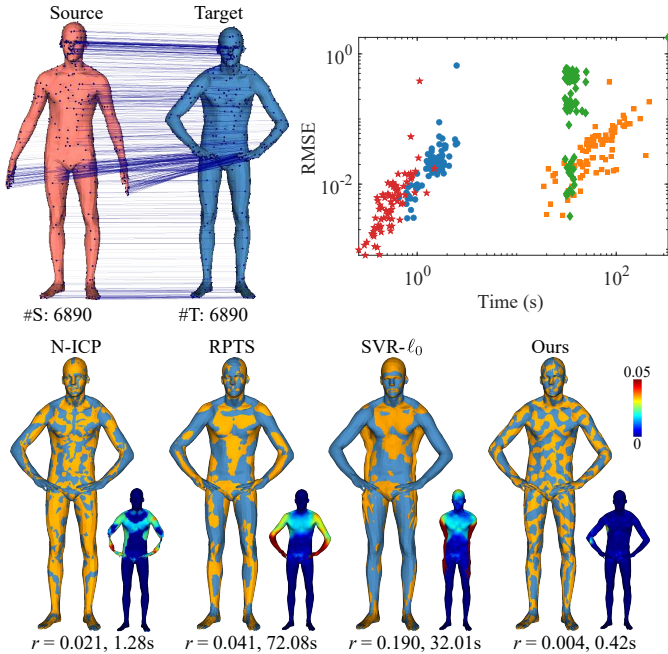


Fig. 9. Comparisons on the MPI Faust dataset [61]. The plot in the upper right visualizes the computational time and RMSE using different methods on 80 pairs of models from the dataset, while the rendered images show the results on a particular pair.

TABLE 2

Mean RMSE ($\times 10^{-3}$) and average computational time (s) using different methods on models from the MPI Faust dataset [61]. We set $\alpha = 0.1$ for N-ICP, $\alpha = 0.1, \beta = 0.001$ for RPTS, $\alpha = 0.001, \beta = 1$ for SVR- ℓ_0 , and $k_\alpha = 0.1, k_\beta = 0.001$ for our method.

Pose pair	N-ICP		RPTS		SVR- ℓ_0		Ours	
	Time	RMSE	Time	RMSE	Time	RMSE	Time	RMSE
1	1.47	21.30	67.18	43.77	35.84	170.75	0.52	4.74
2	0.97	9.69	35.52	14.54	35.13	34.08	0.40	2.21
3	1.44	26.29	88.58	65.32	36.60	406.23	0.65	7.36
4	1.85	20.69	86.69	23.83	35.02	166.94	0.59	4.50
5	0.84	5.60	29.64	6.96	34.31	52.13	0.34	1.25
6	1.59	19.85	56.75	23.68	63.55	541.82	0.61	7.22
7	2.05	110.15	125.90	132.49	34.84	487.12	0.96	55.75
8	1.53	22.87	66.29	47.80	36.25	542.07	0.57	8.46
Mean	1.47	29.56	69.57	44.80	38.94	300.14	0.58	11.44
Median	1.50	21.00	66.74	33.80	35.49	288.49	0.58	5.98

pair. We can see that our method requires significantly less computational time while achieving better accuracy. A major factor for the efficiency of our method is the adoption of a deformation graph, which only requires optimizing one affine transformation per graph vertex. In comparison, N-ICP and RPTS require optimizing one affine transformation per mesh vertex, which significantly increases the number of variables as well as computational time. In addition, the RMSE of SVR- ℓ_0 is notably higher than the other methods, as it was originally designed for aligning nearby frames for motion tracking and becomes less effective on models with large deformation.

Noisy data. To evaluate the effectiveness of our method on noisy data, we synthesize noisy models using the “jumping” dataset from [57]. Specifically, we collect 50 pairs of models, where each pair consists of the i -th frame from the dataset

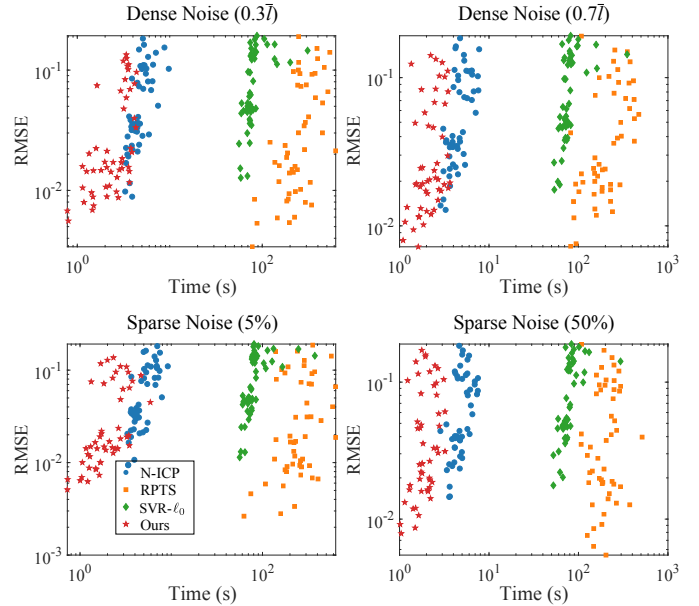


Fig. 10. Visualization of computational time and RMSE resulting from different methods on 50 pairs of models constructed from the “jumping” dataset [57] with different profiles of added noise.

TABLE 3

Mean/median RMSE ($\times 10^{-2}$) and average computational time (s) using different methods on 50 pairs of models constructed from the “jumping” dataset [57] with different noise profiles. We set $\alpha = 10$ for N-ICP, $\alpha = 100, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1$ for SVR- ℓ_0 , and $k_\alpha = 100, k_\beta = 1$ for our method.

Data		N-ICP	RPTS	SVR- ℓ_0	Ours
Dense Noise (0.3 \bar{l})	Mean RMSE	6.39	4.47	8.67	3.62
	Median RMSE	3.96	2.16	6.13	1.64
	Average Time	4.88	268.07	86.19	2.60
Dense Noise (0.7 \bar{l})	Mean RMSE	6.62	5.06	8.77	4.04
	Median RMSE	4.16	2.74	6.25	2.02
	Average Time	4.74	224.94	83.50	2.27
Sparse Noise (5%)	Mean RMSE	6.41	4.47	8.66	3.71
	Median RMSE	3.83	1.82	6.23	1.59
	Average Time	5.09	270.34	89.42	2.17
Sparse Noise (50%)	Mean RMSE	6.96	5.37	8.83	5.23
	Median RMSE	4.47	3.44	6.55	3.08
	Average Time	4.95	206.29	82.87	1.99

as the source model, and the $(i+2)$ -th frame with added Gaussian noise as the target model. For each pair, we generate four noisy variants of the target model as follows:

- Dense Noise (0.3 \bar{l}): we add noises with a standard deviation $\sigma = 0.3\bar{l}$ to all vertices;
- Dense Noise (0.7 \bar{l}): we add noises with a standard deviation $\sigma = 0.7\bar{l}$ to all vertices;
- Sparse Noise (5%): we add noises with a standard deviation $\sigma = \bar{l}$ to 5% of the vertices.
- Sparse Noise (50%): we add noises with a standard deviation $\sigma = \bar{l}$ to 50% of the vertices.

Fig. 10 visualizes the computational time and RMSE for each method on the 50 pairs of models for each noise profile, while Fig. 11 shows the results on a particular model pair. Tab. 3 shows for each method the average RMSE and computational time for all model pairs under each noise profile. We can see that our method achieves better overall performance in

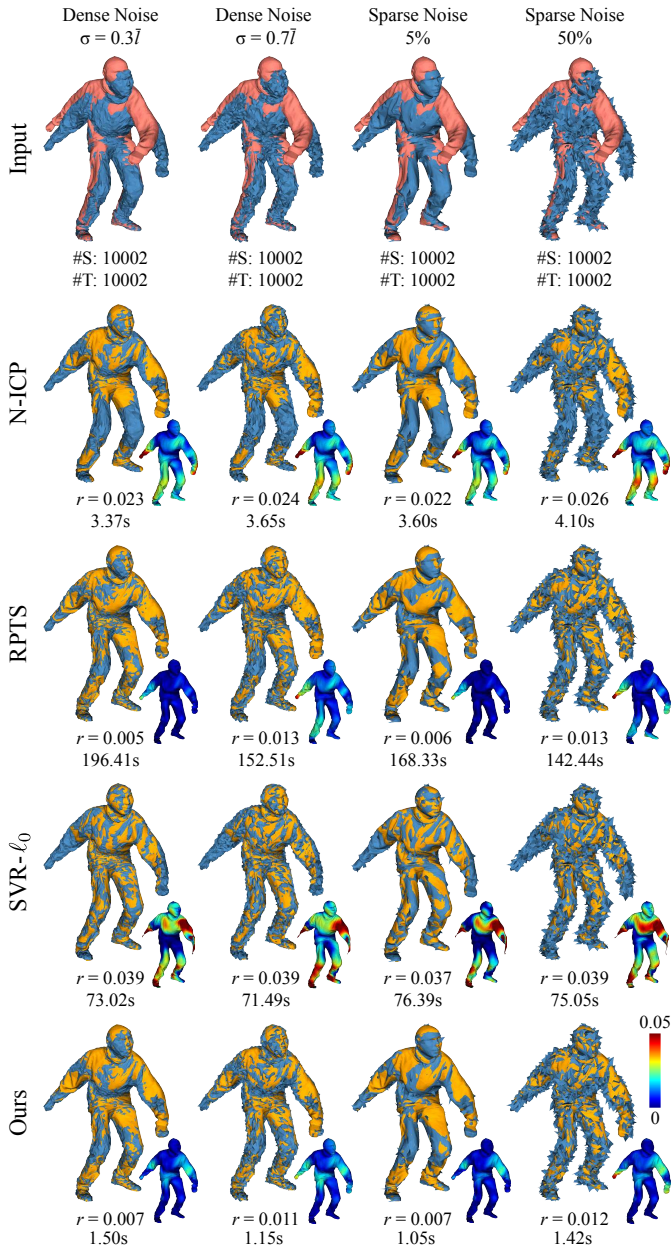


Fig. 11. Comparison on a pair of models constructed from the “jumping” dataset [57] using different profiles of added noise. We set $\alpha = 10$ for N-ICP, $\alpha = 100, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1$ for SVR- ℓ_0 , and $\alpha = 100, \beta = 1$ for our method.

terms of both accuracy and efficiency.

Partial overlaps. We also evaluate different methods on models that overlap partially. We choose three pairs of models from the “march2” dataset and three pairs of models from the “squat1” dataset from [57], where the models in each pair are two frames apart. Then for each pair, we construct two different view directions for the two models respectively, and use their visible parts from these directions to construct source and target models that overlap partially when aligned. We also construct five pairs of models from the TOSCA dataset [62] in a similar way. Figs. 12 and 13 show the results from different methods on these partially overlapping models. N-ICP performs poorly on such models, since its

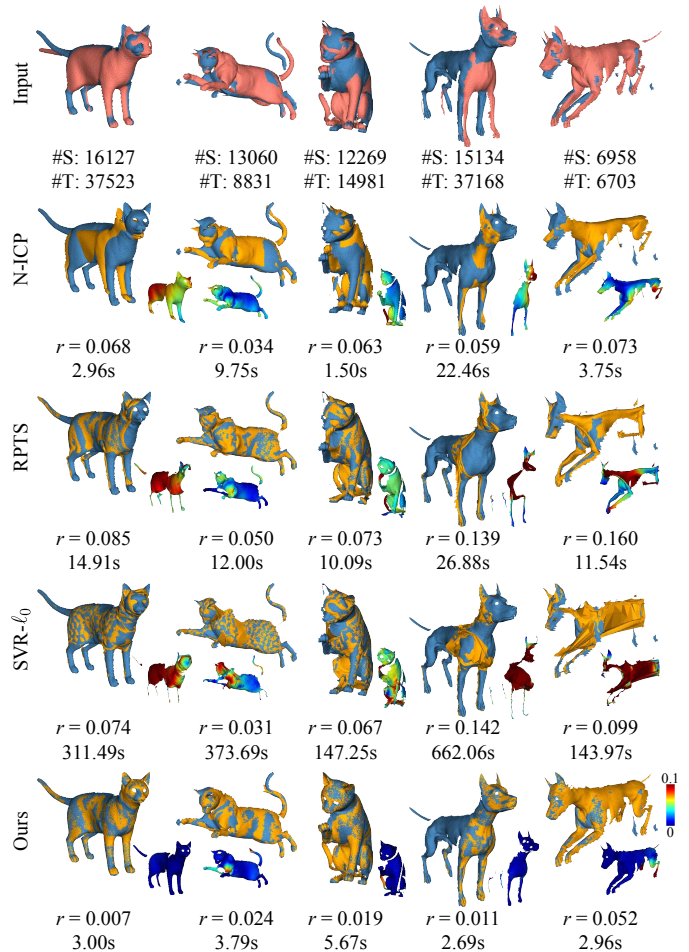


Fig. 12. Comparison on partially overlapping data constructed from the TOSCA dataset [62]. We set $\alpha = 100$ for N-ICP, $\alpha = 1000, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1$ for SVR- ℓ_0 , and $\alpha = 100, \beta = 10$ for our method. We set $R = 8l$ because the mesh is uneven.

ℓ_2 -based optimization can be sensitive to target function terms with large residuals. The ℓ_1 -based formulation in RPTS is more robust than an ℓ_2 formulation, but it still attempts to reduce large residual terms, which can lead to incorrect alignments. SVR- ℓ_0 applies an ℓ_0 -based regularization term but its alignment term is still based on the ℓ_2 -norm, which can affect its robustness. In comparison, our use of the Welsch’s function on both the alignment term and the regularization term helps to improve robustness and achieve overall better results in these examples.

Real scanned data. Our method is also effective on real-world data. To demonstrate this, we conduct experiments on the real scanned data from the DeepDeform dataset [41]. This dataset provides RGB-D frame pairs with estimated optical flows and scene flows and manually annotated sparse corresponding point pairs. To evaluate the accuracy of the result, we adopt the scene flows provided by the dataset as the ground-truth deformation, and compute an error measure as follows:

$$r_s = \sqrt{\frac{1}{|\mathcal{S}|} \sum_{\mathbf{t}_i^s \in \mathcal{S}} \|\hat{\mathbf{v}}_i - (\mathbf{v}_i + \mathbf{t}_i^s)\|^2},$$

where \mathcal{S} is the set of scene flows for the source surface points.



Fig. 13. Comparison on six pairs of partially overlapping models constructed from the “march2” (left) and “squat1” (right) datasets [57]. We set $\alpha = 10$ for N-ICP, $\alpha = 1000, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1$ for SVR- ℓ_0 , and $\alpha = 100, \beta = 10$ for our method.

We use 349 pairs in the validation set for testing, as the validation set includes the scene flows required for computing r_s . Besides N-ICP, RPTS and SVR- ℓ_0 , we also compare with four recent deep learning methods, Leopard+NICP [65], NDP [66], RMA-Net [44], and SyNoRiM [67] using their open-source implementations^{3,4,5}. For each frame, we first downsample the depth map and obtain a point cloud according to the intrinsic camera parameters, then remove the background using the ground-truth segmentation masks provided by the dataset. Finally, we construct a triangular mesh from the point cloud according to their adjacency in the depth

map, and extract the largest connected part. We then use either the resulting mesh or its vertices as the input to a method depending on its requirement. Since SyNoRiM requires 8192 points on each input shape, when the number of points on an input shape is different from 8192 we either downsample the points or add repeated points to achieve the required number. To improve the performance of RMA-Net, we fine-tune the model on the 2120 data pairs from the training set of the DeepDeform dataset. For other learning-based methods, since they already test their performance on the DeepDeform dataset, we directly evaluate their performance using their pre-trained models. We perform the same processing on the training data as mentioned above for the testing data. The training and testing of the

³<https://github.com/rabbity1/DeformationPyramid>

⁴<https://github.com/WanquanF/RMA-Net>

⁵<https://github.com/huangjh-pub/synorim>

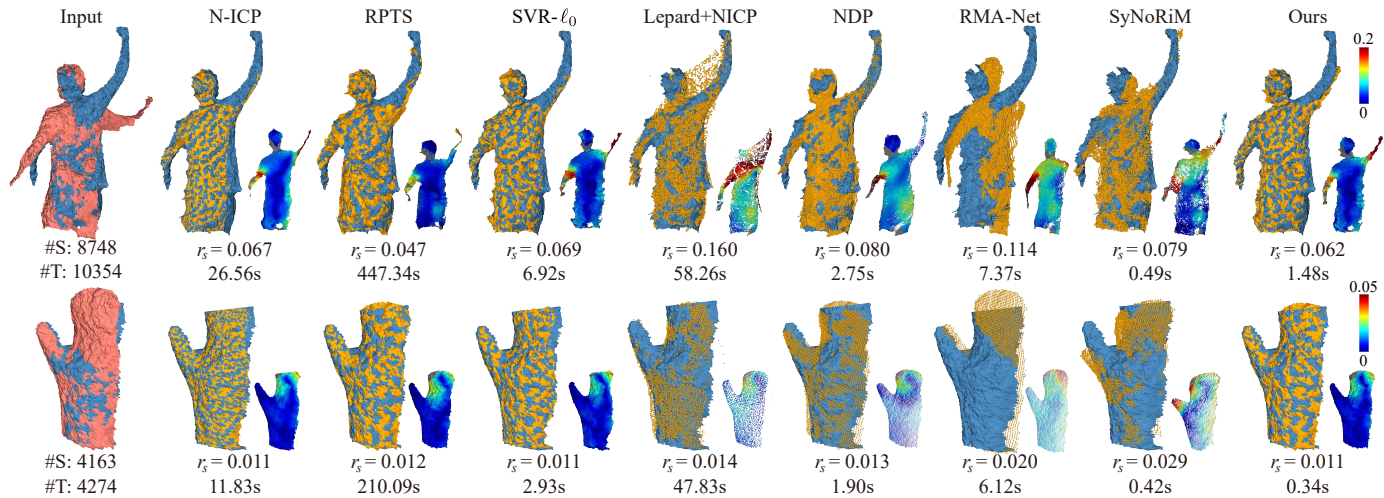


Fig. 14. Comparison between different methods on the DeepDeform dataset [41]. We set $\alpha = 10$ for N-ICP, $\alpha = 100, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1, R = 8\bar{l}$ for SVR- ℓ_0 , and $k_\alpha = 10, k_\beta = 1, R = 8\bar{l}$ for ours. In the visualized error map, we show the points without ground-truth scene flows in gray.

TABLE 4

Mean/median r_s ($\times 10^{-2}$) and average computational time (s) using different methods on 349 pairs of models from DeepDeform dataset [41]. For Lepard+NICP, there are only 219 valid results. We set $\alpha = 10$ for N-ICP, $\alpha = 100, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1, R = 8\bar{l}$ for SVR- ℓ_0 , and $k_\alpha = 10, k_\beta = 1, R = 8\bar{l}$ for our method.

Method	Mean r_s	Median r_s	Average Time
N-ICP	6.58	2.00	74.61
RPTS	6.49	1.88	1050.64
SVR- ℓ_0	6.67	2.07	10.47
Lepard+NICP	8.04	4.40	54.55
NDP	6.66	1.86	2.04
RMA-Net	15.72	7.63	7.56
SyNoRiM	7.90	2.75	0.56
Ours	6.49	1.93	1.39




Fig. 15. Comparison between different methods on a Raspberry Pi for DeepDeform dataset [41]. We set $\alpha = 10$ for N-ICP, $\alpha = 100, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1, R = 8\bar{l}$ for SVR- ℓ_0 , and $k_\alpha = 10, k_\beta = 1, R = 8\bar{l}$ for our method. In the visualized error map, we mark the points without ground-truth scene flows in gray.

learning-based methods are run on an NVIDIA GeForce RTX 3090 GPU. For both SVR- ℓ_0 and our method, we set the sampling radius $R = 8\bar{l}$ for the deformation graph construction for better efficiency. For the optimization-based methods, we first perform a coarse alignment using the rigid registration method from [12] before running each method for non-rigid registration. Tab. 4 shows the mean and median error measures and average computational time for each method, while Fig. 14 shows some example results from the methods. Overall, the optimization-based methods achieve lower registration errors than the learning-based methods, which is likely due to the capability of optimization-based methods to locally refine the alignment. RMA-Net has worse accuracy than other learning-based methods, because its training does not consider partial overlaps between the input shapes. Our method achieves the lowest mean value of r_s and the third lowest median value of r_s (and is only 3% higher than the lowest median r_s value), which shows the good accuracy of our method. On the other hand, the learning-based methods are overall faster than the optimization-based methods, thanks to the efficiency of running a neural network model on the GPU. Nevertheless, our method is only slower than SyNoRiM and is significantly faster than other optimization-based methods.

Performance on mobile devices. Furthermore, due to the use of the deformation graph and our efficient solver, our method can be deployed on devices with limited compute capabilities. To this end, we compare our method and other optimization-based methods on a Raspberry Pi 4B with 8GB of RAM and a quad-core Cortex-A72 at 1.5GHz running Ubuntu 20.04. In Fig. 15, we show the comparison on a problem instance from the the DeepDeform dataset. In Fig. 16, we compare the methods on data pairs obtained by scanning a self-rotating human subject with a Kinect camera; the scans are converted to triangle meshes for registration, by removing isolated points and connecting the remaining points according to their adjacency in the depth map. Due to the lower compute power of Raspberry Pi, we relax the termination condition of our method to $\epsilon = 3 \times 10^{-4}$, which causes no notable change to the result while still being tighter than other methods. We can see that our method is significantly faster than other methods and achieves similar or better accuracy.

5 CONCLUSION

In this paper, we proposed a robust non-rigid registration model based on Welsch's function. Applying the Welsch's




Fig. 16. Comparison between different methods on a Raspberry Pi for real data collected with Kinect. We set $\alpha = 100$ for N-ICP, $\alpha = 10, \beta = 1$ for RPTS, $\alpha = 1, \beta = 1$ for SVR- ℓ_0 , and $k_\alpha = 10, k_\beta = 100$ for our method.

function to the alignment term and the regularization term makes the formulation robust to noises and partial overlaps. To efficiently solve this problem, we apply majorization-minimization to transform the nonlinear and non-convex problem into a sequence of simple sub-problems. To speed up the convergence, we regard the MM algorithm as a fixed-point iteration and use Anderson acceleration to accelerate the solution. Extensive experiments demonstrate the effectiveness of our method and its efficiency compared to existing approaches.

Although our method achieves good results on many examples, it still has some limitations. First, although our use of the robust metric helps to mitigate the impact of incorrect correspondence, our non-convex formulation still requires proper initialization to avoid converging to an undesirable local minimum. In particular, when there is a significant difference between the source and target shapes, the closest-point correspondence may be erroneous on a large part of the surface and lead to inaccurate alignment (see Fig. 17 for an example of such failure cases). The issue can potentially be addressed by adopting a more sophisticated point correspondence method that also accounts for the local shape features. Secondly, our deformation graph construction method relies on a global radius parameter that leads to an approximately uniform density of graph nodes. Such a deformation graph may not be suitable for meshes with highly non-uniform samples. This can be improved using an adaptive radius parameter that depends on the local curvatures and sampling densities, which we will leave as future work.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (No. 62122071), the Youth Innovation




Fig. 17. A failure case of our method on a problem instance from the DeepDeform dataset [41], where the significant difference between the source and target models leads to inaccurate alignment.

Promotion Association CAS (No. 2018495), and the Fundamental Research Funds for the Central Universities (No. WK3470000021).

REFERENCES

- [1] B. Amberg, S. Romdhani, and T. Vetter, "Optimal step nonrigid icp algorithms for surface registration," in *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2007, pp. 1–8.
- [2] H. Li, R. W. Sumner, and M. Pauly, "Global correspondence optimization for non-rigid registration of depth scans," *Computer Graphics Forum*, vol. 27, no. 5, pp. 1421–1430, 2008.
- [3] H. Li, B. Adams, L. J. Guibas, and M. Pauly, "Robust single-view geometry and motion reconstruction," *ACM Transactions on Graphics (ToG)*, vol. 28, no. 5, p. 175, 2009.
- [4] J. Yang, K. Li, K. Li, and Y.-K. Lai, "Sparse non-rigid registration of 3d shapes," *Computer Graphics Forum*, vol. 34, no. 5, pp. 89–99, 2015.
- [5] K. Li, J. Yang, Y.-K. Lai, and D. Guo, "Robust non-rigid registration with reweighted position and transformation sparsity," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 6, pp. 2255–2269, 2018.
- [6] T. Jiang, X. Yang, J. Zhang, F. Tian, S. Liu, N. Xiang, and K. Qian, "Huber- l_1 -based non-isometric surface registration," *The Visual Computer*, vol. 35, no. 6-8, pp. 935–948, 2019.
- [7] K. Guo, F. Xu, Y. Wang, Y. Liu, and Q. Dai, "Robust non-rigid motion tracking and surface reconstruction using l_0 regularization," in *IEEE International Conference on Computer Vision*, 2015, pp. 3083–3091.
- [8] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein *et al.*, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [9] P. W. Holland and R. E. Welsh, "Robust regression using iteratively reweighted least-squares," *Communications in Statistics-theory and Methods*, vol. 6, no. 9, pp. 813–827, 1977.
- [10] B. Ham, M. Cho, and J. Ponce, "Robust image filtering using joint static and dynamic guidance," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 4823–4831.
- [11] J. Zhang, B. Deng, Y. Hong, Y. Peng, W. Qin, and L. Liu, "Static/dynamic filtering for mesh geometry," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 4, pp. 1774–1787, 2018.
- [12] J. Zhang, Y. Yao, and B. Deng, "Fast and robust iterative closest point," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [13] K. Lange, *MM optimization algorithms*. SIAM, 2016.
- [14] D. G. Anderson, "Iterative procedures for nonlinear integral equations," *Journal of the ACM (JACM)*, vol. 12, no. 4, pp. 547–560, 1965.
- [15] H. F. Walker and P. Ni, "Anderson acceleration for fixed-point iterations," *SIAM Journal on Numerical Analysis*, vol. 49, no. 4, pp. 1715–1735, 2011.
- [16] Y. Yao, B. Deng, W. Xu, and J. Zhang, "Quasi-newton solver for robust non-rigid registration," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7600–7609.

- [17] G. K. Tam, Z.-Q. Cheng, Y.-K. Lai, F. C. Langbein, Y. Liu, D. Marshall, R. R. Martin, X.-F. Sun, and P. L. Rosin, "Registration of 3d point clouds and meshes: A survey from rigid to nonrigid," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 7, pp. 1199–1217, 2013.
- [18] B. Deng, Y. Yao, R. M. Dyke, and J. Zhang, "A survey of non-rigid 3d registration," *Computer Graphics Forum*, vol. 41, no. 2, pp. 559–589, 2022.
- [19] R. W. Sumner, J. Schmid, and M. Pauly, "Embedded deformation for shape manipulation," *ACM Transactions on Graphics (TOG)*, vol. 26, no. 3, p. 80, 2007.
- [20] M. Wand, P. Jenke, Q. Huang, M. Bokeloh, L. Guibas, and A. Schilling, "Reconstruction of deforming geometry from time-varying point clouds," in *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*. Eurographics Association, 2007, p. 49–58.
- [21] M. Liao, Q. Zhang, H. Wang, R. Yang, and M. Gong, "Modeling deformable objects from a single depth camera," in *2009 IEEE 12th International Conference on Computer Vision*. IEEE, 2009, pp. 167–174.
- [22] P. Huang, C. Budd, and A. Hilton, "Global temporal registration of multiple non-rigid surface sequences," in *CVPR 2011*. IEEE, 2011, pp. 3473–3480.
- [23] Z. Wu, K. Li, Y.-K. Lai, and J. Yang, "Global as-conformal-as-possible non-rigid registration of multi-view scans," in *IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2019, pp. 308–313.
- [24] J. Yang, D. Guo, K. Li, Z. Wu, and Y.-K. Lai, "Global 3d non-rigid registration of deformable objects using a single rgb-d camera," *IEEE Transactions on Image Processing*, vol. 28, no. 10, pp. 4746–4761, 2019.
- [25] M. Andriy and S. Xubo, "Point set registration: coherent point drift," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 12, pp. 2262–2275, 2010.
- [26] O. Hirose, "A bayesian formulation of coherent point drift," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 7, pp. 2269–2286, 2021.
- [27] H. Hontani, T. Matsuno, and Y. Sawada, "Robust nonrigid icp using outlier-sparsity regularization," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2012, pp. 174–181.
- [28] B. Jian and B. C. Vemuri, "A robust algorithm for point set registration using mixture of gaussians," in *IEEE International Conference on Computer Vision (ICCV)*, vol. 2. IEEE, 2005, pp. 1246–1251.
- [29] —, "Robust point set registration using gaussian mixture models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 8, pp. 1633–1645, 2011.
- [30] M. Wand, B. Adams, M. Ovsjanikov, A. Berner, M. Bokeloh, P. Jenke, L. Guibas, H. P. Seidel, and A. Schilling, "Efficient reconstruction of nonrigid shape and motion from real-time 3d scanner data," *ACM Transactions on Graphics (TOG)*, vol. 28, no. 2, pp. 1–15, 2009.
- [31] J. Ma, W. Qiu, J. Zhao, Y. Ma, A. L. Yuille, and Z. Tu, "Robust L_2E estimation of transformation for non-rigid registration," *IEEE Transactions on Signal Processing*, vol. 63, no. 5, pp. 1115–1129, 2015.
- [32] M. Zollhöfer, M. Nießner, S. Izadi, C. Rehmann, C. Zach, M. Fisher, C. Wu, A. Fitzgibbon, C. Loop, C. Theobalt *et al.*, "Real-time non-rigid reconstruction using an rgb-d camera," *ACM Transactions on Graphics (ToG)*, vol. 33, no. 4, pp. 1–12, 2014.
- [33] R. A. Newcombe, D. Fox, and S. M. Seitz, "Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 343–352.
- [34] T. Yu, K. Guo, F. Xu, Y. Dong, Z. Su, J. Zhao, J. Li, Q. Dai, and Y. Liu, "Bodyfusion: Real-time capture of human motion and surface geometry using a single depth camera," in *IEEE International Conference on Computer Vision*, 2017, pp. 910–919.
- [35] T. Yu, Z. Zheng, K. Guo, J. Zhao, Q. Dai, H. Li, G. Pons-Moll, and Y. Liu, "DoubleFusion: Real-time capture of human performances with inner body shapes from a single depth sensor," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7287–7296.
- [36] Z. Li, T. Yu, C. Pan, Z. Zheng, and Y. Liu, "Robust 3d self-portraits in seconds," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1344–1353.
- [37] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, "Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching," *Proceedings of the National Academy of Sciences*, vol. 103, no. 5, pp. 1168–1172, 2006.
- [38] A. Dubrovina and R. Kimmel, "Approximately isometric shape correspondence by matching pointwise spectral features and global geodesic structures," *Advances in Adaptive Data Analysis*, vol. 3, no. 01n02, pp. 203–228, 2011.
- [39] Q. Chen and V. Koltun, "Robust nonrigid registration by convex optimization," in *IEEE International Conference on Computer Vision*, 2015, pp. 2039–2047.
- [40] S. Shimada, V. Golyanik, E. Tretschk, D. Stricker, and C. Theobalt, "DispVoxNets: Non-rigid point set alignment with supervised learning proxies," in *International Conference on 3D Vision (3DV)*, 2019, pp. 27–36.
- [41] A. Božič, M. Zollhofer, C. Theobalt, and M. Nießner, "DeepDeform: Learning non-rigid RGB-D reconstruction with semi-supervised data," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7002–7012.
- [42] Y. Li, A. Božič, T. Zhang, Y. Ji, T. Harada, and M. Nießner, "Learning to optimize non-rigid tracking," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4910–4918.
- [43] A. Božič, P. Palafox, M. Zollhöfer, A. Dai, J. Thies, and M. Nießner, "Neural non-rigid tracking," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33, 2020, pp. 18 727–18 737.
- [44] W. Feng, J. Zhang, H. Cai, H. Xu, J. Hou, and H. Bao, "Recurrent multi-view alignment network for unsupervised surface registration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 297–10 307.
- [45] K. Zampogiannis, C. Fermüller, and Y. Aloimonos, "Topology-aware non-rigid point cloud registration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 3, pp. 1056–1069, 2021.
- [46] L. Xu, Z. Su, L. Han, T. Yu, Y. Liu, and L. Fang, "UnstructuredFusion: realtime 4D geometry and texture reconstruction using commercial RGBD cameras," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 10, pp. 2508–2522, 2020.
- [47] Z. Su, L. Xu, Z. Zheng, T. Yu, Y. Liu, and L. Fang, "RobustFusion: Human volumetric capture with data-driven visual cues using a RGBD camera," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*. Springer, 2020, pp. 246–264.
- [48] A. Toth and C. Kelley, "Convergence analysis for anderson acceleration," *SIAM Journal on Numerical Analysis*, vol. 53, no. 2, pp. 805–819, 2015.
- [49] C. Evans, S. Pollock, L. G. Rebholz, and M. Xiao, "A proof that Anderson acceleration improves the convergence rate in linearly converging fixed-point methods (but not in those converging quadratically)," *SIAM Journal on Numerical Analysis*, vol. 58, no. 1, pp. 788–810, 2020.
- [50] H.-r. Fang and Y. Saad, "Two classes of multisection methods for nonlinear acceleration," *Numerical Linear Algebra with Applications*, vol. 16, no. 3, pp. 197–221, 2009.
- [51] Y. Peng, B. Deng, J. Zhang, F. Geng, W. Qin, and L. Liu, "Anderson acceleration for geometry optimization and physics simulation," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–14, 2018.
- [52] J. Zhang, Y. Peng, W. Ouyang, and B. Deng, "Accelerating ADMM for efficient simulation and optimization," *ACM Trans. Graph.*, vol. 38, no. 6, pp. 163:1–163:21, 2019.
- [53] W. Ouyang, Y. Peng, Y. Yao, J. Zhang, and B. Deng, "Anderson acceleration for nonconvex ADMM based on Douglas-Rachford splitting," *Computer Graphics Forum*, vol. 39, no. 5, pp. 221–239, 2020.
- [54] A. L. Pavlov, G. V. Ovchinnikov, D. Y. Derbyshev, D. Tsetserukou, and I. V. Oseledets, "AA-ICP: iterative closest point with anderson acceleration," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 1–6.
- [55] Y. Qin, X. Han, H. Yu, Y. Yu, and J. Zhang, "Fast and exact discrete geodesic computation based on triangle-oriented wavefront propagation," *ACM Trans. Graph.*, vol. 35, no. 4, 2016.
- [56] C. Moenning and N. A. Dodgson, "Fast Marching farthest point sampling," University of Cambridge, Computer Laboratory, Tech. Rep. UCAM-CL-TR-562, Apr. 2003. [Online]. Available: <https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-562.pdf>
- [57] D. Vlasic, I. Baran, W. Matusik, and J. Popović, "Articulated mesh animation from multi-view silhouettes," in *ACM SIGGRAPH 2008 papers*, 2008, pp. 1–9.

- [58] S. Bouaziz, A. Tagliasacchi, and M. Pauly, "Sparse iterative closest point," in *Computer graphics forum*, vol. 32, no. 5. Wiley Online Library, 2013, pp. 113–123.
- [59] I. Daubechies, R. DeVore, M. Fornasier, and C. S. Güntürk, "Iteratively reweighted least squares minimization for sparse recovery," *Communications on Pure and Applied Mathematics*, vol. 63, no. 1, pp. 1–38, 2010.
- [60] G. Guennebaud, B. Jacob *et al.*, "Eigen v3," <http://eigen.tuxfamily.org>, 2010.
- [61] F. Bogo, J. Romero, M. Loper, and M. J. Black, "Faust: Dataset and evaluation for 3d mesh registration," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 3794–3801.
- [62] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008.
- [63] F. Tombari, S. Salti, and L. Di Stefano, "Unique signatures of histograms for local surface description," in *European conference on computer vision*. Springer, 2010, pp. 356–369.
- [64] G. K. Tam, R. R. Martin, P. L. Rosin, and Y. Lai, "Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry," *ACM Transactions on Graphics*, vol. 33, no. 1, 2014.
- [65] Y. Li and T. Harada, "Leopard: Learning partial point cloud matching in rigid and deformable scenes." *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [66] —, "Non-rigid point cloud registration with neural deformation pyramid." *arXiv preprint arXiv:2205.12796*, 2022.
- [67] J. Huang, T. Birdal, Z. Gojcic, L. J. Guibas, and S.-M. Hu, "Multiway non-rigid point cloud registration via learned functional map synchronization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.



Juyong Zhang is a professor in the School of Mathematical Sciences at University of Science and Technology of China. He received the BS degree from the University of Science and Technology of China in 2006, and the PhD degree from Nanyang Technological University, Singapore. His research interests include computer graphics, computer vision, and numerical optimization. He is an associate editor of *IEEE Transactions on Multimedia* and *The Visual Computer*.



Yuxin Yao is currently working toward a PhD degree in the School of Mathematical Sciences, University of Science and Technology of China. Her research interests include computer vision, computer graphics and numerical optimization.



Bailin Deng is a Senior Lecturer in the School of Computer Science and Informatics at Cardiff University. He received the BEng degree in computer software (2005) and the MSc degree in computer science (2008) from Tsinghua University (China), and the PhD degree in technical mathematics (2011) from Vienna University of Technology (Austria). His research interests include geometry processing, numerical optimization, computational design, and digital fabrication. He is a member of the IEEE.



Weiwei Xu is a researcher with the State Key Lab of CAD & CG, College of Computer Science, Zhejiang University, awardee of NSFC Excellent Young Scholars Program in 2013. His main research interests include the digital geometry processing, physical simulation, computer vision and virtual reality. He has published around 70 papers on international graphics journals and conferences, including 20 papers on ACM TOG. He is a member of the IEEE and ACM.