Cronfa - Swansea University Open Access Repository

_____

This is an author produced version of a paper published in:
Computer Aided Geometric Design

_____

Cronfa URL for this paper:

http://cronfa.swan.ac.uk/Record/cronfa49861

_____

http://www.swansea.ac.uk/library/researchsupport/ris-support/

Non-rigid registration under anisotropic deformations

Roberto M. Dyke [a],[*], Yu-Kun Lai [a], Paul L. Rosin [a], Gary K.L. Tam [b]

[a] *School of Computer Science & Informatics, Cardiff University, United Kingdom*
[b] *Department of Computer Science, Swansea University, United Kingdom*

A B S T R A C T

Non-rigid registration of deformed 3D shapes is a challenging and fundamental task in geometric processing, which aims to non-rigidly deform a source shape into alignment with a target shape. Current state-of-the-art methods assume deformations to be near-isometric. This assumption does not reflect real-world conditions, for example in large-scale deformation, where moderate anisotropic deformations (e.g., stretches) are common. In this paper we propose two significant changes to a typical registration pipeline to address such challenging deformations. First, we introduce a method to estimate anisotropic non-isometric deformations and incorporate this into an iterative non-rigid registration pipeline. Second, we compute additional correspondences in non-isometrically deforming regions using reliable correspondences as landmarks and prune inconsistent correspondences. We compare the performance of our proposed algorithm to several state-of-the-art methods using existing benchmarks. Experimental results show that our method outperforms existing methods.

1. Introduction

Surface registration is a fundamental problem in the domains of computer graphics and vision, in which the aim is to find a transformation that best aligns two input surfaces. Surface registration algorithms underline computational solutions to many prevailing problems, such as 3D acquisition/reconstruction, statistical shape analysis and shape retrieval. With the increasing ubiquity of 3D scanners and application of 3D scanning in real-world scenarios, the importance of accurate registration algorithms is continuing to rise.

In many real-life scenarios, surfaces are often non-rigidly deformed. Non-rigid surface registration is therefore required to find the non-rigid transformation between them. Extending from the well-known Iterative Closest Point (ICP) approach for rigid registration Besl and McKay (1992); Chen and Medioni (1992), Non-rigid ICP (N-ICP) methods Bouaziz and Pauly (2013) achieve registration for non-rigid surfaces by alternating between two steps. In the first step, a set of correspondences is computed using a closest point criterion, and then the second step identifies a non-rigid transformation that minimises an error metric. Generally, the associated cost function decreases after each iteration, converging monotonically to a local minimum. Because of the simple way correspondences are generated, N-ICP is fast enough to be used in some real-time applications; though alone, it is incapable of coping with large-scale deformations. N-ICP methods thus typically require an

initial set of correspondences generated through alternative means, such as automatic/manual markers (e.g., Amberg et al. (2007)), so as to achieve good registration results.

In the N-ICP framework, dense correspondences are obtained using the closest point criterion. However, it is only effective when two surfaces are reasonably close. Additional (sparse) correspondences are often needed to cope with shapes with large deformations. Since shapes to be registered may only have partial overlaps (e.g., due to occlusion when two views are captured), correspondences are most often generated by feature matching local shape descriptors of source and target shapes.

This requirement exhibits two challenges: First, it can lead to false matches between areas that appear the same locally, but belong to different regions (e.g., any local regions on a cylindrical surface). This could be accounted for by ensuring the correspondences are consistent. Second, in cases where the surface has been warped, they may not match (unless the shape signatures are insensitive to the particular warp, which is not generally possible). Typically, there is little contingency built into these N-ICP methods, with most relying on such areas being sufficiently insignificant so as to not affect the final result.

Most recent correspondence methods (e.g., Huang et al. (2008); Tam et al. (2014)) used in non-rigid registration account for a certain degree of near-isometric deformation by employing a geodesic-based consistency measure. However, for surfaces with large deformations, non-isometric deformations commonly exist (typically in joint areas of articulation, or on deforming parts), making the global isometry assumption less useful. In the literature, most surface registration methods do not directly address non-isometric deformation. They simply penalise any deformation that is non-isometric. To cope with large deformations, we propose a new consistency measure that takes into account anisotropic non-isometric deformation explicitly.

Some previous work addresses a related problem of establishing correspondences between shapes which involve potentially large deformations Kim et al. (2011); Litany et al. (2017b). However, fundamental differences exist. Non-rigid registration aims to identify non-rigid deformation (usually in the form of a set of local deformations) to deform the source mesh to align with the target, whereas correspondence methods only identify point-to-point correspondences between surfaces. On the one hand, in order to directly derive deformation with the latter approach, complete or dense per-vertex correspondences would be needed. The computation of such dense correspondences can be slow. When two surfaces are only partially overlapping, such complete correspondences may not be defined. On the other hand, these methods may have significant restrictions on topology (complete genus-zero models) Kim et al. (2011) or may generate globally inconsistent correspondences due to the lack of a global constraint Litany et al. (2017b). Compared to these methods, non-rigid registration can often be driven by a sparse set of reliable correspondences which is easier and more efficient to obtain. It also does not require any strong topology assumptions.

Due to the typical sparsity of the correspondences used, non-rigid registration usually requires some regularisation to ensure that the surface deforms appropriately during the N-ICP iterations. Notable methods define the local regularisation neighbourhood to be 1-ring (i.e., vertices connected to the vertex of interest by an edge) Bouaziz and Pauly (2013) using an as-rigid-as-possible (ARAP) formulation Sorkine and Alexa (2007). This however is insufficient for cases with large deformations, as shown in our experiments. Following Chen et al. (2017), which uses r-ring ARAP energy for deformation with controllable stiffness, we incorporate ARAP-based regularisation over larger neighbourhoods, allowing large deformation to be handled effectively in the N-ICP framework.

Technical contribution In this paper, we propose a novel non-rigid registration technique capable of registering large-scale and non-isometric deformations. Addressing large-scale and non-isometric deformation is a fundamental challenge of existing non-rigid registration techniques. To the best of our knowledge, this challenge has not been addressed before. More specifically, our main technical contributions are:

- A novel method to estimate anisotropic deformations on a discrete mesh by using the principal scaling factor.
- A correspondence generation and correspondence pruning method based on local geodesics that copes with anisotropic deformations; this makes use of our anisotropic deformation estimate.
- We introduce the r-ring ARAP formulation for regularisation in non-rigid registration, which effectively handles challenging large deformations where existing registration methods fail.

We further perform both qualitative and quantitative evaluation using public benchmark datasets. Our results show that our method outperforms the state-of-the-art methods in non-rigid registration.

2. Related work

Surface registration Over the past decades, many improvements of the original ICP techniques Besl and McKay (1992); Chen and Medioni (1992) were proposed. The underlying mechanism is similar and consists of identifying correspondences using closest point criterion, and an iterative alignment to handle both rigidly Rusinkiewicz and Levoy (2001) and some non-rigidly deformed surfaces Tam et al. (2013). Pottmann et al. (2004) uses a local quadratic approximation that provides fast convergence. While these methods are highly efficient, the closest point criterion used to generate correspondences in N-ICP methods is often inappropriate for registering surfaces undergoing large-scale deformations.

Non-rigid registration methods were initially developed to register range scans where the scanning devices used had been incorrectly calibrated Li et al. (2009). Brown and Rusinkiewicz (2004) proposes an efficient N-ICP based on thin plate spline deformation. Amberg et al. (2007) describes a semi-automatic template-based approach that uses a sparse set of hand-picked correspondences on the target surface to initialise registration. Huang et al. (2008) presents a technique capable of registering large-scale deformations. Near-rigidly deforming regions of a shape are segmented into clusters, which are then deformed and registered. The method thus assumes a mostly piecewise-rigid deformation from the source to the target shape. To obtain reliable correspondences, they adapt the spectral matching technique Leordeanu and Hebert (2005), which computes a globally consistent set of correspondences by assuming a near-isometric deformation.

Feature descriptors   Local feature descriptors may be used to establish correspondences between keypoints with similar local geometry. Over the years, many robust feature descriptors have been proposed Johnson and Hebert (1999); Frome et al. (2004); Rusu et al. (2008, 2009); Tombari et al. (2010). They have been extensively evaluated in Guo et al. (2016). Among those, the Signature of Histograms of Orientations (SHOT) Tombari et al. (2010) is a highly popular descriptor, because of its computational efficiency and state-of-the-art performance. More recently Zeng et al. (2017) proposed a volumetric feature descriptor using a Convolutional Neural Network (CNN) architecture. Other descriptors based on CNNs have also been proposed Wei et al. (2016); Monti et al. (2017). However, these techniques may still generate inconsistent correspondences and usually require a large amount of data for training. Since identifying the optimal feature descriptor for non-rigid surfaces is not the main focus of this work, we simply use SHOT signatures Tombari et al. (2010) to establish initial correspondences due to its robust performance.

Shape correspondences   A variety of methods have been proposed for establishing correspondences between shapes Kim et al. (2011); Pokrass et al. (2013); Kovnatsky et al. (2013); Litany et al. (2017b). These methods however focus on producing (dense) correspondences rather than non-rigid registration. Therefore, they can be slow with dense meshes. They may also assume the shapes are (nearly) complete Kim et al. (2011); Pokrass et al. (2013); Kovnatsky et al. (2013) or may produce incorrect correspondences due to the lack of global constraints Litany et al. (2017b).

For non-rigid registration, it is usually sufficient to have a sparse set of correspondences. Leordeanu and Hebert (2005) first uses spectral matching to acquire a consistent set of correspondences between images. Huang et al. (2008) adapts the idea to address large-scale, but piecewise-rigid deformations. They further use an assumption of global isometry to obtain correspondences. However, the computation of geodesic distances is slow, and using an isometry measure based on global geodesic distances may not adequately model real-life deformation well. Due to these drawbacks, Tam et al. (2014) proposes to use a more flexible local isometry global consistency model, along with a more efficient diffusion framework to select reliable correspondences. All these methods assume (near-)isometric deformation. However, these techniques still cannot model real-life deformation well. Here, to the best of our knowledge, we are the first to explicitly model anisotropic non-isometric deformation in a pruning and registration technique, allowing reliable correspondences and non-rigid alignment to be obtained even under large-scale deformation.

Data-driven non-rigid registration   Similar to many other geometry processing problems, non-rigid registration also benefits from a data-driven approach where additional constraints can be enforced by the availability of a collection of relevant shapes. Hirshberg et al. (2012) develops a co-registration technique for multiple shapes. Methods have also been developed for registration of human faces Bouaziz et al. (2013) and human bodies Loper et al. (2015) by exploiting pre-built (linear) shape models. For more general cases Gao et al. (2016) proposes a new data-driven deformation representation and exploits given examples to constrain plausible deformations in non-rigid registration. Although these methods produce competitive results, they rely on a sufficiently large collection of shapes to train the underlying model. Furthermore, reliable collections of shapes may not be easily available in many practical situations.

Shape deformation   Typically, non-rigid registration methods minimise an objective function, comprising of a data and a regularisation term. The regularisation term restricts arbitrary deformations, and ensures a natural deformation. Sorkine (2006) first proposes a Laplacian deformation framework for meshes. Sorkine and Alexa (2007) builds upon the concept, and proposes a simplified approach for mesh editing. The method computes the optimal local rotations and translations for each vertex's one-ring neighbourhood. Kilian et al. (2007) ensures near-isometric shape deformation by preserving geodesics between surface points. Sauvage et al. (2008) proposes a deformation framework for B-spline surfaces that incorporates a volume constraint. Popa et al. (2006) proposes a deformation model capable of handling anisotropic materials through the use of local co-ordinate frames. Achenbach et al. (2015) successfully applies an anisotropic regularisation energy with edge Laplacians to a template face model, to improve alignment for anisotropic facial regions such as wrinkles. Such anisotropic regularisation is only used in the later refinement stage and therefore the method is not designed for large-scale anisotropic deformations, which we address in this paper. Chen et al. (2017) proposes a simple formulation that extends Sorkine and Alexa (2007), enlarging the neighbourhood size (to r-ring neighbours) to change the local rigidity of deformations. In this work, we introduce the r-ring ARAP energy to the N-ICP framework and demonstrate its effectiveness in handling large-scale deformations. This approach may seem counter-intuitive, as it penalises anisotropic deformation. However when it comes to challenging large-scale deformations where reliable correspondences may not be readily available, such regularisation is necessary to ensure the deformation is not arbitrary and does not overfit to noise.

Non-rigid registration   Non-rigid registration aims to align surfaces undergoing non-rigid deformation. Much research has been conducted due to its general applicability Tam et al. (2013). However, most research makes strong assumptions regarding the type of deformation, e.g., piecewise-rigid Huang et al. (2008) and isometric Yang et al. (2015); Li et al. (2018). Although these assumptions are reasonable when the deformations are small or near-isometric, it is still a challenge to register surfaces with large-scale non-isometric deformation, in which case anisotropic deformation is common. To our knowledge, this problem has not been addressed in the existing non-rigid registration research.

Functional maps   The *functional map* framework, first proposed by Ovsjanikov et al. (2012), represents correspondences between two functional spaces as a linear map. However, this approach can only work on topologically consistent, isometric shapes. Rodolà et al. (2017) considers the problem of part-to-whole shape matching, observing that a sub-matrix of the Laplacian of a whole shape may correspond closely with the eigenvectors of the Laplacian of a part of a similar shape. The method alternately optimises for the part on the whole shape and the functional correspondences. This allows near-isometric, partial shapes to be matched. Litany et al. (2017b) further enhances Rodolà et al. (2017), to enable it to be solved solely in the spectral domain, supporting non-isometries and partial shapes. Huang et al. (2014) proposes functional map networks that prove robust for matching between heterogeneous shape collections that may be partial and non-isometric. Kovnatsky et al. (2015) obtains functional correspondences by incorporating a geometric structure prior. This method is capable of handling missing parts, non-isometries, and geometric and topological noise. Vestner et al. (2017) adopt *heat kernels* rather than geodesic distances to improve robustness to topological noise and non-isometric deformation. Recently, Litany et al. (2017a) proposed an intrinsic deep learning-based functional map framework. Since Laplacian eigenfunctions are known to be sensitive to topological changes, the use of dense and invariant feature descriptors in the data term is necessary. Typically, descriptors such as SHOT (Rodolà et al. (2017); Litany et al. (2017a,b)), HKS (heat kernel signatures), WKS (wave kernel signatures) or a combination of the three are incorporated (Kovnatsky et al. (2015); Vestner et al. (2017)). These techniques are able to produce correspondences, but not local deformations to align non-rigidly deformed surfaces, and therefore do not solve the same problem as ours.

3. Methodology

Our N-ICP algorithm (Algorithm 1) consists of four key components: modelling of anisotropic non-isometric deformation (Section 3.1), correspondence computation (Section 3.2), correspondence pruning (Section 3.3) and deformation optimisation (Section 3.4). These components are underpinned by our incorporation of an anisotropic geodesic method. In this section, we first present an overview of the algorithm, and subsequently detail each component in the following subsections.

Algorithm 1 Algorithm overview.

    Input: Source and target surfaces (X, Y).
    Output: deformed surface (X') and correspondences $(C_{\text{N-ICP}})$
1:   $d_X \leftarrow \text{ExactGeodesics}(X)$
2:   $d_Y \leftarrow \text{ExactGeodesics}(Y)$
3:   $C_{\text{FeatureMatch}} \leftarrow \text{FeatureMatch}(X, Y)$ ▷ Local feature matching.
4:   $C_{\text{pruned}} \leftarrow \text{prune}(d_X, d_Y, C_{\text{FeatureMatch}})$
5:   $X', C_{\text{N-ICP}} \leftarrow \text{N-ICP}(X, Y, C_{\text{pruned}})$
6: **while** X' has sufficient changes over the last iteration **do**
7:      $E_X \leftarrow \text{PrincipalScalingFactor}(X, Y, d_X, d_Y, C_{\text{N-ICP}})$ ▷ Output source edge lengths.
8:      $d'_X \leftarrow \text{AnisotropicGeodesics}(X, E_X)$
9:      $C_{\text{pruned}} \leftarrow \text{prune}(d'_X, d_Y, C_{\text{N-ICP}})$
10:     $C_{\text{landmark}} \leftarrow \text{landmark}(X, Y, d'_X, d_Y, C_{\text{pruned}})$
11:     $C_{\text{pruned}} \leftarrow \text{prune}(d'_X, d_Y, C_{\text{landmark}} \cup C_{\text{pruned}})$ ▷ Use scaling factor in consistency measure.
12:     $X', C_{\text{N-ICP}} \leftarrow \text{N-ICP}(X, Y, C_{\text{pruned}})$
13: **end while**

3.1. Anisotropy estimation

In real-life non-rigid scenarios, anisotropic non-isometric deformation is common. For example, the contraction of the muscle bicep would cause the arm to bend at the elbow (as shown in Fig. 7). Such muscle contraction involves shortening in longitudinal and lengthening in lateral directions of the muscle. It is a typical anisotropic non-isometric deformation. However, most existing work assumes isotropic near-isometric deformation only. Such an assumption is inadequate to capture good correspondence in a non-rigidly deforming region.

To better model these anisotropic non-isometric deformations for a registration technique, we develop a local anisotropy metric, and incorporate it in the local geodesic computation. This allows more and better correspondences to be found in non-isometrically deforming areas, supporting more robust registration results. In the following, we first discuss the anisotropy metric in Section 3.1.1, then discuss how to adopt the metric for geodesic computation in Section 3.1.2.

(a) Anisotropic deformation.　　(b) Two pairs of eigenvectors　　(c) New basis $\boldsymbol{B} = (\boldsymbol{e}_1, \boldsymbol{e}_2, \boldsymbol{e}_3)$.

Fig. 1. Illustrations that demonstrate the idea of estimating anisotropic non-isometric deformation. a) An anisotropic deformation of points between two shapes (X & Y) may be approximately modelled by an ellipse. The surrounding unlabelled points represent neighbouring points on the source shape $\boldsymbol{s}_j \in N_X(\boldsymbol{s}_i)$ and the predicted correspondence of $\boldsymbol{t}_v$ on surface Y. b) Two pairs of eigenvectors ($\hat{\boldsymbol{e}}_{1a}, \hat{\boldsymbol{e}}_{2a}, \hat{\boldsymbol{e}}_{1b}, \hat{\boldsymbol{e}}_{2b} \in \mathbb{R}^2$) represent stretching for an edge $\boldsymbol{v}_{ab}$. c) $\boldsymbol{v}_{ab}$ transforming into new basis $\boldsymbol{B} = (\boldsymbol{e}_1, \boldsymbol{e}_2, \boldsymbol{e}_3)$.

3.1.1. Local anisotropic deformation metric

To model the local anisotropic deformation at each vertex, we propose an anisotropy metric – a pair of local eigenvectors that defines the *Principal Scaling Factor* (PSF) of deformation.

We first project local points in the source mesh onto the tangential plane of the vertex of concern. Then we can estimate the local scaling factor by comparing these 2D coordinates with the corresponding geodesic distances from the target mesh (illustrated in Fig. 1a). This allows us to extract the scaling directions and magnitudes via eigen-analysis, inspired by the idea of principal curvatures Euler (1760). These constitute one part of a 2-order tensor field that describes changes in scaling at a per-vertex level of the source shape.

In order to estimate PSF, it is necessary to have some points with known correspondences. This leads to a chicken-and-egg problem. We address this issue with an iterative approach. We first compute initial correspondences by matching features between shapes, as per step 3 of Algorithm 1. Consider a discrete vertex $\boldsymbol{s}_i$ on the source shape X, with neighbouring vertices $\boldsymbol{s}_j \in N_X(\boldsymbol{s}_i)$ where the correspondences of $\boldsymbol{s}_i$ and all $\boldsymbol{s}_j$ are known. We project $\boldsymbol{s}_i$ onto a 2D plane with an arbitrary orientation that is tangential to the normal of $\boldsymbol{s}_i$. Let $\hat{\boldsymbol{s}}_i$ and $\hat{\boldsymbol{s}}_j$ be the projection of vertices $\boldsymbol{s}_i$ and $\boldsymbol{s}_j$ onto the 2D plane, respectively. The corresponding points on the target shape for $\boldsymbol{s}_i$ and $\boldsymbol{s}_j$ on the source shape are denoted by $\boldsymbol{t}_u$ and $\boldsymbol{t}_v$.

For simpler formulation, we set $\hat{\boldsymbol{s}}_i = (0,0)$ as the origin, and denote $\hat{\boldsymbol{s}}_j = (\hat{x}_j, \hat{y}_j)$. We also define $d_j = d_g(\boldsymbol{t}_u, \boldsymbol{t}_v)$ as the geodesic distance on the target shape between $\boldsymbol{t}_u$ and $\boldsymbol{t}_v$. To ensure robustness, we formulate the problem as a least squares fitting, and use a quadratic function $f(\hat{x}, \hat{y})$ to approximate d_j^2 since low-order polynomials are more stable. The squared distance should have its centre of symmetry at the origin and $f(0,0) = 0$, at which the x term, y term and constant thus should all be 0, leading to

$$f(\hat{x}, \hat{y}) = a\hat{x}^2 + b\hat{x}\hat{y} + c\hat{y}^2, \tag{1}$$

where a, b and c are coefficients to be determined. We formulate the least squares problem to minimise the following

$$F(a, b, c) = \sum_{\boldsymbol{s}_j \in N_X(\boldsymbol{s}_i)} \left(a\hat{x}_j^2 + b\hat{x}_j\hat{y}_j + c\hat{y}_j^2 - d_j^2 \right)^2. \tag{2}$$

Setting $\frac{\partial F}{\partial (a,b,c)} = 0$ leads to a linear system, which can be easily solved. Let $\hat{\boldsymbol{x}} = (\hat{x}, \hat{y})$. We construct a matrix $\boldsymbol{A}_q$

$$\boldsymbol{A}_q = \begin{bmatrix} a & \frac{1}{2}b \\ \frac{1}{2}b & c \end{bmatrix}, \tag{3}$$

and $f(\hat{\boldsymbol{x}}) = \hat{\boldsymbol{x}}^T \boldsymbol{A}_q \hat{\boldsymbol{x}}$. We then find the first and second eigenvectors of $\boldsymbol{A}_q$, i.e., $\hat{\boldsymbol{e}}_1$, $\hat{\boldsymbol{e}}_2$ and their respective eigenvalues ($\|\hat{\boldsymbol{e}}_1\|$, $\|\hat{\boldsymbol{e}}_2\|$). These eigenvectors represent the anisotropy metric at a specific vertex. We may calculate $\boldsymbol{e}_1, \boldsymbol{e}_2 \in \mathbb{R}^3$ by mapping the tangential plane that $\boldsymbol{e}_1, \boldsymbol{e}_2$ lie on back into $\mathbb{R}^3$. This process is repeated for each vertex in X with a correspondence on Y.

3.1.2. Measuring anisotropic geodesic distance

To measure geodesic distance we use the method proposed by Liu et al. (2017). This method requires the anisotropy metric to be encoded into the predicted edge lengths of the shape. We update the estimated lengths of the edge set for the source shape using the following procedure.

Let $\boldsymbol{v}_{ab}$ be a vector that represents an edge between two adjacent vertices on the source shape (i.e., $\boldsymbol{v}_{ab} = \boldsymbol{s}_b - \boldsymbol{s}_a$). At each vertex, there is a pair of eigenvectors, a principal and secondary vector ($\hat{\boldsymbol{e}}_1$, $\hat{\boldsymbol{e}}_2$). Thus, for each edge, there are two

Fig. 2. Example problematic region between two shapes with no correspondences.

pairs of eigenvectors ($\hat{\boldsymbol{e}}_{1a}$, $\hat{\boldsymbol{e}}_{2a}$) & ($\hat{\boldsymbol{e}}_{1b}$, $\hat{\boldsymbol{e}}_{2b}$). The eigenvalue is encoded into the length of each eigenvector $\|\hat{\boldsymbol{e}}_{\alpha}\|$. An example of this in the 2D tangential plane is illustrated in Fig. 1b.

In practice, $\boldsymbol{e}_{\alpha}$ and $\boldsymbol{v}_{ab}$ are embedded in $\mathbb{R}^3$. We apply the following calculation for both points a and b. For simplicity, we only describe the process for one point. We assign the normal vector at the point of concern to $\boldsymbol{e}_3 = \frac{\boldsymbol{e}_1 \times \boldsymbol{e}_2}{\|\boldsymbol{e}_1 \times \boldsymbol{e}_2\|}$. We use the three normalised vectors as the current basis (i.e., $\boldsymbol{B} = (\frac{\boldsymbol{e}_1}{\|\boldsymbol{e}_1\|}, \frac{\boldsymbol{e}_2}{\|\boldsymbol{e}_2\|}, \boldsymbol{e}_3)$, as $\boldsymbol{e}_3$ is already normalised). Transforming the subject edge into this basis ($\tilde{\boldsymbol{v}}_{ab} = \boldsymbol{B}^{-1} \boldsymbol{v}_{ab}$), the components of each axis now represent the amount of stretching in each direction (n.b. $\boldsymbol{e}_3$ always has a stretching factor of 1). We illustrate this new basis w.r.t. point a in 3D in Fig. 1c. We can then assign each of our scaling factors to an axis as follows: $x : \|\boldsymbol{e}_1\|$, $y : \|\boldsymbol{e}_2\|$ and $z : 1$. We can apply our scaling factor to each respective axis before transforming the stretched edge back into the original global basis ($\tilde{\boldsymbol{v}}_{ab}^* = \boldsymbol{B} \boldsymbol{S} \tilde{\boldsymbol{v}}_{ab}$), where $\boldsymbol{S}$ is a 3-by-3 scaling matrix with $\boldsymbol{S}_{11} = \|\boldsymbol{e}_1\|$, $\boldsymbol{S}_{22} = \|\boldsymbol{e}_2\|$, $\boldsymbol{S}_{33} = 1$ and remaining entries to be zero. This process is repeated for both points a and b using their respective eigenvectors ($\boldsymbol{e}_{1a}, \boldsymbol{e}_{2a}, \boldsymbol{e}_{3a}$ & $\boldsymbol{e}_{1b}, \boldsymbol{e}_{2b}, \boldsymbol{e}_{3b}$). The average length of the resulting two edges is the new edge length (as a discrete approximation to the integration over the edge).

3.2. Correspondence generation

We now describe our method to find correspondences, as per Algorithm 1. In addition to local feature matching, which only works for regions with limited deformation, we further identify correspondences using automatic landmarking in an as-consistent-as-possible formulation. The formulation works even under substantial non-isometric deformation.

3.2.1. Local feature matching

In our implementation we have used SHOT signatures to produce a candidate set of correspondences. To reiterate, we have chosen SHOT for feature matching as it has acceptable performance Tam et al. (2014); Litany et al. (2017b); the focus of this paper is not to seek the best feature descriptor. SHOT signatures are computed at two scales for each vertex on X and Y Tombari et al. (2010). Vertices between the two shapes are matched based on the similarity of their signatures.

As with most low-dimensional feature descriptors, locally symmetric areas may produce similar local signatures. This kind of incorrect correspondence may be identified by checking its consistency with other well matched feature points. For regions with isometric deformations, well established methods based on diffusion pruning Tam et al. (2014) can be used. We will discuss this along with our extended approach for non-isometric deformation later in Section 3.3.

3.2.2. Correspondence identification based on automatic landmarking

In regions with substantial non-isometric deformation, matching of local feature descriptors (e.g., SHOT) often fails. Here, we propose to use reliable correspondences from isometrically deforming regions as landmarks to help identify new candidate correspondences in these regions. The landmarks are automatically selected and updated in each iteration of our N-ICP framework. Given these landmarks, new candidate correspondences are proposed based on the consistency of their geodesic distances to the landmarks. There are three steps to this process, namely: problematic region identification, automatic landmark selection and correspondence matching.

Problematic region identification Initially, regions with substantial non-isometric deformation can be found by applying a pruning technique Tam et al. (2014). Regions with no (locally isometrically consistent) correspondences will be identified as problematic regions (Fig. 2). Each problematic region is then separately analysed. To save computation time, very small regions are disregarded since adequate correspondences will be identified later at the ICP stage.

Automatic selection of landmark correspondences Correspondences returned from the pruning algorithm Tam et al. (2014) are generally reliable, especially those scoring high confidence values. In principle, we can randomly pick any of these correspondences as landmarks. However, landmarks that are far away from problematic regions are not very effective in estimating

new candidate correspondences. Further, in practice, we observe that only a small number of landmarks are necessary to uniquely determine a new candidate correspondence point $\boldsymbol{y}_j \in Y$ for a given point $\boldsymbol{x}_i \in X$. We thus sample 20 landmarks that are closest to the boundary of a problematic region. This set of landmarks L helps to reduce the ambiguity and impact of slight errors in the landmark locations. Although a more sophisticated landmark selection procedure would be desirable, through experiments, we found that our method is insensitive to the number of landmarks, and produces consistent results when this parameter is varied within a reasonable range. We therefore fixed this parameter in our experiments.

Matching of correspondences For each point $\boldsymbol{x}_i$ in a problematic region on surface X, the point $\boldsymbol{y}_j \in Y$ with the most consistent geodesic distances to landmarks is chosen as the corresponding point. Let $\boldsymbol{l}_k = (\boldsymbol{l}_k^X, \boldsymbol{l}_k^Y) \in L$ be the kth landmark correspondence with $\boldsymbol{l}_k^X \in X$, $\boldsymbol{l}_k^Y \in Y$. We measure the consistency of a correspondence $(\boldsymbol{x}_i, \boldsymbol{y}_j)$ w.r.t. a landmark $\boldsymbol{l}_k$ as

$$\mathrm{cons}_k(\boldsymbol{x}_i, \boldsymbol{y}_j) = \min\left(\frac{d_g(\boldsymbol{x}_i, \boldsymbol{l}_k^X)}{d_g(\boldsymbol{y}_j, \boldsymbol{l}_k^Y)}, \frac{d_g(\boldsymbol{y}_j, \boldsymbol{l}_k^Y)}{d_g(\boldsymbol{x}_i, \boldsymbol{l}_k^X)}\right). \tag{4}$$

Similar to Huang et al. (2008); Tam et al. (2014), Eqn. (4) provides a normalised measure of consistency in the range of $[0, 1]$ where 0 is the worst and 1 is the best. $d_g(\cdot, \cdot)$ measures the geodesic distance between two vertices. However, unlike existing work that sets a threshold, which does not generally work for non-isometrically deforming regions, we choose the target corresponding vertex $\boldsymbol{y}_j$ as the one that gives the overall best geodesic consistency. It works even for regions with non-isometric deformations:

$$\underset{j}{\arg\max} \, \underset{k}{\min} \, \mathrm{cons}_k(\boldsymbol{x}_i, \boldsymbol{y}_j). \tag{5}$$

The rationale is to choose correspondences that give good alignment (by maximising the minimum consistency) to all the landmarks, whilst reducing the effect of far away landmarks that are not discriminative enough for identifying a new good correspondence.

3.3. Correspondence pruning

Tam et al. (2014) proposes an efficient algorithm capable of pruning a set of candidate correspondences, preserving only those that are globally consistent. Poor correspondences are pruned based on their low consistency with good local correspondences. The consistency measure k_{ab} for a pair of correspondences $a = (\boldsymbol{s}_i, \boldsymbol{t}_u)$ and $b = (\boldsymbol{s}_j, \boldsymbol{t}_v)$ is defined as:

$$k_{ab} = \min\left(\frac{d_g(\boldsymbol{s}_i, \boldsymbol{s}_j)}{d_g(\boldsymbol{t}_u, \boldsymbol{t}_v)}, \frac{d_g(\boldsymbol{t}_u, \boldsymbol{t}_v)}{d_g(\boldsymbol{s}_i, \boldsymbol{s}_j)}\right), \tag{6}$$

where $\boldsymbol{s}_i, \boldsymbol{s}_j \in X$ and $\boldsymbol{t}_u, \boldsymbol{t}_v \in Y$ are vertices on surfaces X and Y respectively. In their work Huang et al. (2008); Tam et al. (2014), a threshold c_0 (set to 0.7, as suggested by Huang et al. (2008)) is introduced, and a pair of correspondences is considered acceptable if $k_{ab} \geq c_0$ in order to penalise non-isometric deformations. Tam et al. (2014) extends Huang et al. (2008) by considering only local isometry where end points $\boldsymbol{s}_j$ and $\boldsymbol{t}_v$ of a correspondence b must be within a specified geodesic distance to the end points $\boldsymbol{s}_i$ and $\boldsymbol{t}_u$ of a correspondence a. The global consistency of all reliable correspondences can be further inferred from local isometry via diffusion processing. This technique is shown to perform well in near-isometric cases.

However, in general, these techniques Huang et al. (2008); Tam et al. (2014) fail if the surfaces undergo non-isometric deformation, especially the large-scale deformation that we consider in this paper. Our observation is that large-scale deformation often consists of non-isometric (esp. anisotropic) deformation. As these techniques model isotropic (near-)isometric deformation only, they fail to return any good correspondences essential for accurate N-ICP alignment. Here, we develop a local anisotropy metric (Section 3.1) to estimate anisotropic deformation during registration, and apply Eqn. (6) to incorporate an anisotropic geodesic measure Liu et al. (2017) $d_g'(\cdot, \cdot)$. To our knowledge, this is the first effort to explicitly model anisotropic non-isometric deformation in a pruning technique.

3.4. N-ICP with extended ARAP regularisation

Our N-ICP implementation is inspired by Bouaziz and Pauly (2013). At each iteration of N-ICP, we work out the deformed position $\boldsymbol{x}_i'$ for each vertex $\boldsymbol{x}_i$ minimising an objective function involving both the data and regularisation terms, i.e.:

$$E = E_{\mathrm{data}} + \lambda E_{\mathrm{reg}}. \tag{7}$$

For the first 5 iterations, we only consider correspondences from pruning (C_{pruned} in Algorithm 1). This transforms the surface to an initial deformation that is more suitable. Let C be the set of pruned correspondences and $m = |C|$ be the number of correspondences input. Denote each correspondence as $c_k = (\boldsymbol{x}_{c_k}, \boldsymbol{y}_{c_k})$ where $\boldsymbol{x}_{c_k} \in X$ and $\boldsymbol{y}_{c_k} \in Y$ respectively. The data term we used in the standard formulation:

$$E_{\text{data}} = \sum_{k=1}^{m} \left\| \boldsymbol{x}'_{c_k} - \boldsymbol{y}_{c_k} \right\|_2^2. \tag{8}$$

After the 5th iteration, we discard prior correspondences, and use nearest neighbour correspondences to help guide the mesh into alignment. Our data term is now formulated as a combination of point-to-point and point-to-plane distances, where $\boldsymbol{y}_i$ is the closest point to $\boldsymbol{x}'_i$ on shape Y, and $\boldsymbol{n}_i$ is the normal of point $\boldsymbol{y}_i$. w_1 and w_2 are the respective weights of point-to-point and point-to-plane terms.

$$E_{\text{data}} = w_1 \sum_{i=1}^{n} \left\| \boldsymbol{x}'_i - \boldsymbol{y}_i \right\|_2^2 + w_2 \sum_{i=1}^{n} \left\| \boldsymbol{n}_i \left(\boldsymbol{x}'_i - \boldsymbol{y}_i \right) \right\|_2^2 \tag{9}$$

For the regularisation term, traditional N-ICP uses a local ARAP energy in 1-ring neighbourhoods Sorkine and Alexa (2007). We observe that this regularisation is not sufficient for registration of regions with large-scale deformations. For the purpose of shape deformation, Chen et al. (2017) generalises this concept to work with r-ring neighbourhoods. It would provide greater control over how the deformation of a vertex affects its connected vertices. We propose to incorporate r-ring ARAP into our N-ICP registration pipeline to handle large-scale deformations. Let N_i^r be the r-ring neighbourhood for the vertex $\boldsymbol{x}_i$, and $n = |X|$ be the total number of vertices of X.

$$E_{\text{reg}} = \sum_{i=1}^{n} \sum_{j \in N_i^r} \left\| (\boldsymbol{x}'_j - \boldsymbol{x}'_i) - \boldsymbol{R}_i(\boldsymbol{x}_j - \boldsymbol{x}_i) \right\|_2^2 \tag{10}$$

The optimisation of E involves alternating optimisation of the deformed positions $X' = \{\boldsymbol{x}'_i\}$ for each vertex $\boldsymbol{x}_i$ (while fixing $\mathbf{R}_i$), and the calculation of the optimal local rigid rotation $\mathbf{R}_i$ (while fixing X'). Similar to Sorkine and Alexa (2007); Chen et al. (2017), the former (global step) is a least-squares problem and can be solved by solving a linear system, and the latter is optimised for each vertex and can be solved using Singular Value Decomposition Sorkine and Alexa (2007); Chen et al. (2017).

Note that on line 12, Algorithm 1, we pass X as the source shape rather than X', this enables the consistent application of local regularisation and helps ensure better registration results.

4. Evaluation

To evaluate the performance of our proposed method, we benchmark it on a series of synthetic and real datasets. These datasets contain challenging registration scenarios, that include a range of minor and major anisotropic non-isometric deformations. We compare our technique against the notable non-rigid ICP Bouaziz and Pauly (2013), two recent sparse non-rigid registration methods Yang et al. (2015); Li et al. (2018), and a state-of-the-art functional map method Vestner et al. (2017).

To measure the performance of each method on the datasets, we follow the Princeton benchmark protocol Kim et al. (2011) for geodesic error. The normalised geodesic error between a deformed vertex $\boldsymbol{x}'_i$ and the ground truth position $\boldsymbol{g}_i$ is measured as follows:

$$\epsilon_i = \frac{d_g(\boldsymbol{x}'_i, \boldsymbol{g}_i)}{\text{area}(Y)^{1/2}} \tag{11}$$

Yang et al. (2015); Li et al. (2018); Vestner et al. (2017) and our proposed method each rely on an initial set of correspondences. To ensure our tests were fair, we computed an initial set of correspondences using SHOT Tombari et al. (2010). Vestner et al. (2017) was provided a dense set of feature descriptors; while, for Yang et al. (2015); Li et al. (2018) and our method, we subsequently applied pruning Tam et al. (2014) once for each pair of shapes, and used the same set of initial correspondences. To demonstrate the generalisability of each method, we do not perform further parameter tuning between tests, except where specified.

4.1. Implementation

A 352-dimension SHOT feature descriptor was used for Yang et al. (2015); Vestner et al. (2017); Li et al. (2018) and our method. For the underlying N-ICP of our method, we used the following parameters: w_1 = 10, w_2 = 2, λ = 1000, iterations = 30, & r-ring = 2. The proposed PSF and landmark methods were enabled in all tests, except where specified. For the N-ICP method Bouaziz and Pauly (2013), we kept most of the default settings, except the number of iterations, for which we found a value of 50 was sufficient. To run Vestner et al. (2017), we created a low resolution (5,000 vertices) copy of each mesh using *qslim* Garland and Heckbert (1997) and then applied *MeshFix* Attene (2010), providing both a high resolution and low resolution version of each shape during run-time. Shape pairs were rescaled to ensure their area was consistent. We enabled partial matching and set $\alpha = 10^{-7}$ and $t = [500, 323, 209, 135, 87, 56, 36, 23, 15, 10]$, initially solving for 10,000 correspondences with a maximum problem size of 3,000 for all subsequent iterations.
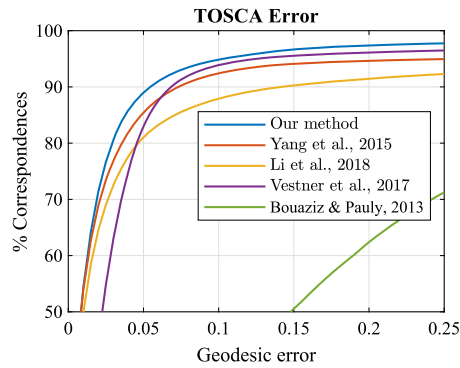
**Fig. 3.** TOSCA high-res dataset results.

Fig. 4. Deformation results on Bouncing dataset Vlasic et al. (2008). a) initial pose source (top) target pose (bottom), b) our method c) sparse non-rigid registration method Yang et al. (2015).

4.2. Synthetic datasets

TOSCA high-resolution dataset   For the TOSCA dataset Bronstein et al. (2008), we generated 69 pairs of shapes with the same class label (i.e., cat to cat, dog to dog, etc.). Each mesh was simplified to have 10,000 faces using *qslim* Garland and Heckbert (1997). Models with flipped normals were excluded from tests. We disabled partial matching for Vestner et al. (2017) on this dataset, and kept other settings (see Section 4.1).

Fig. 3 shows the average registration accuracy of each method on the TOSCA dataset. These graphs show the proportion of vertices that have an error less than the value on the *x*-axis. The techniques corresponding to curves towards the top left of the graph are comparatively better. It should be noted that Bouaziz and Pauly (2013) failed to perform well on the TOSCA dataset, as well as the subsequent datasets. Since this was the simplest dataset, we excluded this method from further tests.

Fig. 10 shows the location of registration errors. For every vertex, we measure the average registration error across a set of deformed models for both our proposed method and Yang et al. (2015). Fig. 11 shows the registration result when dealing with a large non-isometric intra-class deformation on the hind legs of the TOSCA dog.

Bouncing dataset   Fig. 4 shows the results of our proposed method, compared with Yang et al. (2015), on the Bouncing dataset Vlasic et al. (2008). We measure the fitting error between the deformed and target shapes with Hausdorff distance in MeshLab Cignoni et al. (1998). Note that Vestner et al. (2017) only produces correspondences rather than surface alignment, and cannot be used for this evaluation.
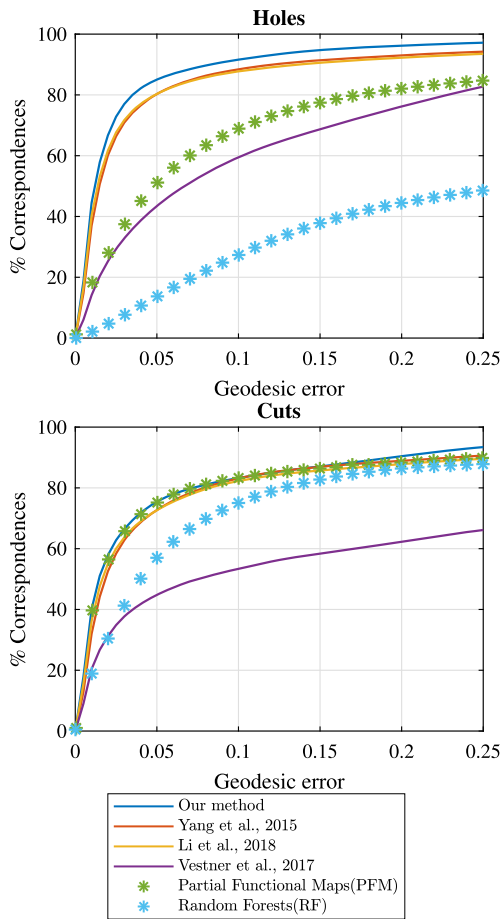
Fig. 5. Results on SHREC'16 partial dataset compared with results (Partial Functional Maps & Random Forests) from dense methods reported in Cosmo et al. (2016).

Fig. 6. Results of our method in different configurations on SHREC'16 partial dataset. In configurations where PSF is off, we use standard geodesics, rather than anisotropic ones (skipping lines 7 & 8 of Algorithm 1). Where landmarking is off, we skip lines 10 & 11 of Algorithm 1.

SHREC'16 partial model dataset We further used the SHREC'16 partial model dataset Cosmo et al. (2016). This dataset is an important derivative work of the TOSCA dataset in which shapes have missing data (holes or cuts) that make non-rigid registration challenging. Fig. 5 illustrates the results of the Partial Functional Maps Rodolà et al. (2017) and Random Forests Rodolà et al. (2014) reported in Cosmo et al. (2016) compared with the results of Yang et al. (2015); Li et al. (2018); Vestner et al. (2017) and our proposed method. This demonstrates that our method is robust to holes and cuts, and clearly outperforms state-of-the-art methods, especially for the cases with holes.

To demonstrate the effectiveness of the individual ideas in our technique (namely, selection of landmarks, principal scaling factor, and *r*-ring ARAP), we ran a series of experiments with different configurations on the SHREC'16 partial model dataset. Results are illustrated in Fig. 6, which show that all the components in our method contribute to improved performance.

In summary, our results demonstrate that our method consistently outperforms Yang et al. (2015); Li et al. (2018); Vestner et al. (2017); Bouaziz and Pauly (2013). In cases where the initial SHOT correspondences were poor, we found that the landmark-based correspondences and *r*-ring ARAP helped improve the alignment of large deformations.

4.3. Real datasets

Partial body scans Results from running our registration pipeline on a collection of partial body scans Allen et al. (2002) are shown in Fig. 7. We also observed that missing data can cause our correspondence generation technique to produce more incorrect correspondences. In our tests, using 2- or 3-ring neighbourhoods helped alleviate this problem and improved registration results. The results presented in Fig. 7 were computed using 3-ring neighbourhoods.

FAUST dataset We used the FAUST dataset Bogo et al. (2014) to help us objectively evaluate our method. FAUST is a collection of real scans captured using a 3D multi-stereo system. Ground-truth correspondences were acquired by covering each

(a)   (b)

Fig. 7. Results of our proposed non-rigid registration technique on partial body scans Allen et al. (2002). a) initial pose source *X* (blue), target *Y* (white), b) overlapped result. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)
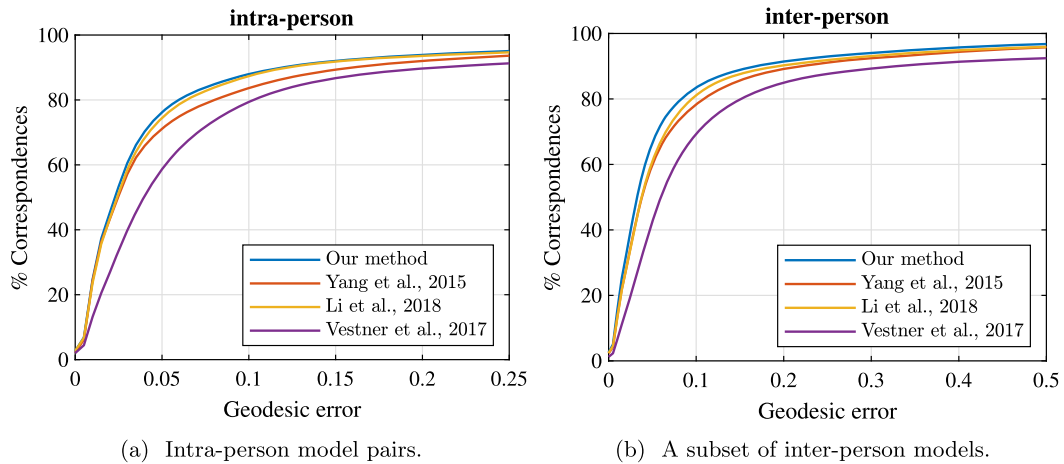

(a) Intra-person model pairs.   (b) A subset of inter-person models.

Fig. 8. Registration results of pairs of models from the FAUST dataset.

subject's body with a series of stamps that were matched using texture-based registration. We compared the performance of our method, Yang et al. (2015), Li et al. (2018) and Vestner et al. (2017).

Since none of the methods require any form of training, we used the training set for our evaluation. In our first experiment we produced an intra-subject data as follows: for each person in the dataset, we registered every mesh of that same individual to their first mesh, creating a set of 90 intra-person shape pairs. We then created our own inter-subject data by pairing each person in the dataset to another person in the dataset in a single, randomly selected pose. This created a set of 90 randomly chosen pairs to evaluate each method with. We used exactly the same set of shape pairs for each method.

Registration results on the FAUST dataset are shown in Fig. 8, and some visual examples in Fig. 9. Our technique consistently shows lower mean geodesic errors compared against the state-of-the-art methods.
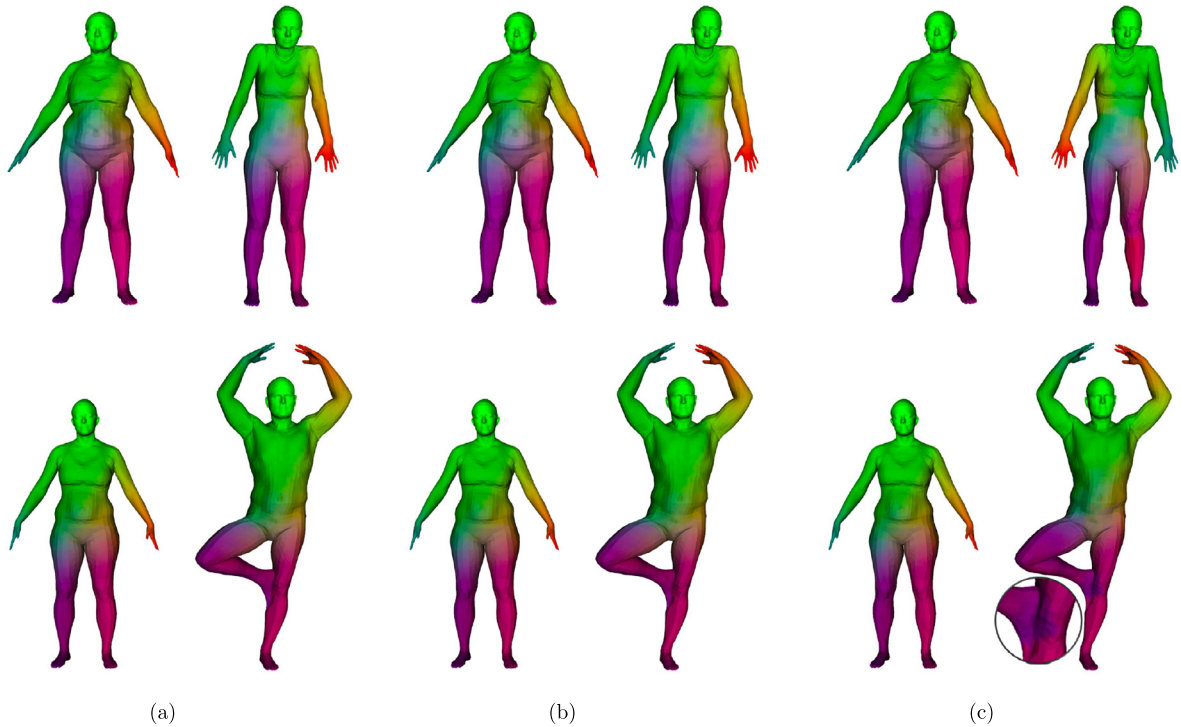
| (a) | (b) | (c) |

Fig. 9. Screenshots of correspondence results of a) our method, b) Li et al. (2018) and c) Vestner et al. (2017) on pairs of inter-person shapes from the FAUST dataset (left source, right target). Colours represent correspondence between shapes (i.e., the same point on each shape should have the same colour). Observe the reflection of correspondences in method (c) in the top row, and the colour bleed at the intersection between legs in the bottom row. The mean geodesic error for each method on the top row is as follows: a) 0.0373, b) 0.0419, c) 0.2971. Similarly, for the bottom row: a) 0.0639, b) 0.0820, c) 0.0840.

It should be noted that the curves in respective figures only show the aggregated performance over the entire datasets. For simpler cases (shape pairs and regions), all the methods work well. When it comes to challenging regions (esp. joints), our technique performs significantly better than the curves show. The superior performance of our method is illustrated by typical examples in the paper.

5. Limitations

With the help of local scaling metric estimation and additional correspondences from landmarks, our method is more robust to large-scale, especially non-isometric deformations in surface registration than existing methods, as demonstrated by extensive comparisons. However, our method still has some limitations. Our method still relies on some correct initial correspondences to start the iterative process, and may perform poorly if the initial correspondences are largely wrong. We currently take a simplified strategy to fix all the parameters. Some parameters, such as the number of landmarks and the r-ring ARAP, may work more effectively when more sophisticated, adaptive algorithms are used; we leave this as future work.

6. Conclusions

In this paper we present a novel pipeline capable of registering two 3D shapes with large and non-isometric non-rigid deformations. We develop a technique to estimate local anisotropic non-isometric deformation and compute reliable correspondences in areas under large-scale deformation. We further develop a new consistency measure and automatic landmark selection method to support non-isometric consistency. We then extend the ARAP regularisation constraint to deform larger surfaces in a more uniform manner. Our experimental results demonstrate that our technique performs well in challenging scenarios. Comparative evaluations also highlight scenarios where state-of-the-art methods fail.

Our technique performs well in non-isometrically deformed regions, and we believe that the technique can be further developed to handle more challenging cases (e.g., heterogeneous objects). In addition, our technique currently only takes a source shape and a target shape as input. It is possible to generalise our method to exploit more example shapes to achieve more accurate registration results.

Fig. 10. Results for the TOSCA cat, gorilla, David & dog sets, a) source/target models used, b) combined registration error of our method, c) combined registration error of Yang et al. (2015).
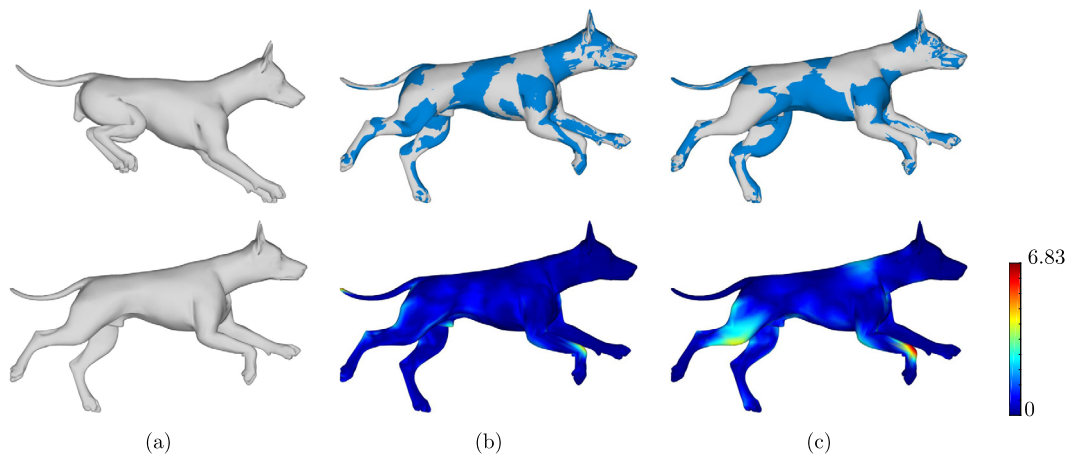

Fig. 11. Registration results on TOSCA dog. a) initial pose source (top) target pose (bottom), b) our method c) Yang et al. (2015).

Acknowledgements

References

Achenbach, J., Zell, E., Botsch, M., 2015. Accurate face reconstruction through anisotropic fitting and eye correction. In: Vision, Modeling & Visualization, the Eurographics Association, pp. 1–8.

Allen, B., Curless, B., Popović, Z., 2002. Articulated body deformation from range scan data. ACM Trans. Graph. 21, 612–619. https://doi.org/10.1145/566570.566626.

Amberg, B., Romdhani, S., Vetter, T., 2007. Optimal step nonrigid ICP algorithms for surface registration. In: Proc. CVPR, pp. 1–8.

Attene, M., 2010. A lightweight approach to repairing digitized polygon meshes. Vis. Comput. 26, 1393–1406. https://doi.org/10.1007/s00371-010-0416-3.

Besl, P.J., McKay, N.D., 1992. A method for registration of 3-d shapes. IEEE Trans. Pattern Anal. Mach. Intell. 14, 239–256. https://doi.org/10.1109/34.121791.

Bogo, F., Romero, J., Loper, M., Black, M.J., 2014. FAUST: dataset and evaluation for 3D mesh registration. In: Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). IEEE, pp. 3794–3801.

Bouaziz, S., Pauly, M., 2013. Dynamic 2D/3D registration for the kinect. In: ACM SIGGRAPH 2013 Courses. ACM, 21.

Bouaziz, S., Wang, Y., Pauly, M., 2013. Online modeling for realtime facial animation. ACM Trans. Graph. 32, 40. https://doi.org/10.1145/2461912.2461976.

Bronstein, A., Bronstein, M., Kimmel, R., 2008. Numerical Geometry of Non-Rigid Shapes, 1 ed. Springer Publishing Company, Incorporated.

Brown, B.J., Rusinkiewicz, S., 2004. Non-rigid range-scan alignment using thin-plate splines. In: Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission. 3DPVT 2004, pp. 759–765.

Chen, S.Y., Gao, L., Lai, Y.K., Xia, S., 2017. Rigidity controllable as-rigid-as-possible shape deformation. Graph. Models 91, 13–21. https://doi.org/10.1016/j.gmod.2017.02.005.

Chen, Y., Medioni, G., 1992. Object modelling by registration of multiple range images. Image Vis. Comput. 10, 145–155. https://doi.org/10.1016/0262-8856(92)90066-C.

Cignoni, P., Rocchini, C., Scopigno, R., 1998. Metro: measuring error on simplified surfaces. In: Computer Graphics Forum, pp. 167–174.

Cosmo, L., Rodolà, E., Bronstein, M.M., Torsello, A., Cremers, D., Sahillioğlu, Y., 2016. Partial matching of deformable shapes. In: Proceedings of the Eurographics 2016 Workshop on 3D Object Retrieval, pp. 61–67.

Euler, L., 1760. Recherches sur la courbure des surfaces. Mém. Acad. Sci. Berlin 16, 9.

Frome, A., Huber, D., Kolluri, R., Bülow, T., Malik, J., 2004. Recognizing Objects in Range Data Using Regional Point Descriptors. Springer Berlin Heidelberg, pp. 224–237.

Gao, L., Lai, Y.K., Liang, D., Chen, S.Y., Xia, S., 2016. Efficient and flexible deformation representation for data-driven surface modeling. ACM Trans. Graph. 35, 157. https://doi.org/10.1145/2908736.

Garland, M., Heckbert, P.S., 1997. Surface simplification using quadric error metrics. In: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques. ACM Press/Addison-Wesley Publishing Co., pp. 209–216.

Guo, Y., Bennamoun, M., Sohel, F., Lu, M., Wan, J., Kwok, N.M., 2016. A comprehensive performance evaluation of 3D local feature descriptors. Int. J. Comput. Vis. 116, 66–89. https://doi.org/10.1007/s11263-015-0824-y.

Hirshberg, D.A., Loper, M., Rachlin, E., Black, M.J., 2012. Coregistration: simultaneous alignment and modeling of articulated 3D shape. In: Proc. ECCV, pp. 242–255.

Huang, Q.X., Adams, B., Wicke, M., Guibas, L.J., 2008. Non-rigid registration under isometric deformations. Comput. Graph. Forum 27, 1449–1457. https://doi.org/10.1111/j.1467-8659.2008.01285.x.

Huang, Q., Wang, F., Guibas, L., 2014. Functional map networks for analyzing and exploring large shape collections. ACM Trans. Graph. 33, 36. https://doi.org/10.1145/2601097.2601111.

Johnson, A.E., Hebert, M., 1999. Using spin images for efficient object recognition in cluttered 3D scenes. IEEE Trans. Pattern Anal. Mach. Intell. 21, 433–449. https://doi.org/10.1109/34.765655.

Kilian, M., Mitra, N.J., Pottmann, H., 2007. Geometric modeling in shape space. In: ACM SIGGRAPH 2007 Papers. ACM, pp. 1–8.

Kim, V.G., Lipman, Y., Funkhouser, T., 2011. Blended intrinsic maps. ACM Trans. Graph. 30, 79. https://doi.org/10.1145/2010324.1964974.

Kovnatsky, A., Bronstein, M.M., Bresson, X., Vandergheynst, P., 2015. Functional correspondence by matrix completion. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 905–914.

Kovnatsky, A., Bronstein, M.M., Bronstein, A.M., Glashoff, K., Kimmel, R., 2013. Coupled quasi-harmonic bases. Comput. Graph. Forum 32, 439–448. https://doi.org/10.1111/cgf.12064.

Lähner, Z., Vestner, M., Boyarski, A., Litany, O., Slossberg, R., Remez, T., Rodolà, E., Bronstein, A.M., Bronstein, M.M., Kimmel, R., Cremers, D., 2017. Efficient deformable shape correspondence via kernel matching. arXiv:1707.08991. https://github.com/zorah/KernelMatching, 2017.

Leordeanu, M., Hebert, M., 2005. A spectral technique for correspondence problems using pairwise constraints. In: Proc. ICCV, vol. 2, pp. 1482–1489.

Li, H., Adams, B., Guibas, L.J., Pauly, M., 2009. Robust single-view geometry and motion reconstruction. ACM Trans. Graph. 28, 175. https://doi.org/10.1145/1661412.1618521.

Li, K., Yang, J., Lai, Y., Guo, D., 2018. Robust non-rigid registration with reweighted position and transformation sparsity. IEEE Trans. Vis. Comput. Graph.. https://doi.org/10.1109/TVCG.2018.2832136.

Litany, O., Remez, T., Rodolà, E., Bronstein, A., Bronstein, M., 2017a. Deep functional maps: structured prediction for dense shape correspondence. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 5660–5668.

Litany, O., Rodolà, E., Bronstein, A.M., Bronstein, M.M., 2017b. Fully spectral partial shape matching. Comput. Graph. Forum. https://doi.org/10.1111/cgf.13123.

Liu, B., Chen, S., Xin, S.Q., He, Y., Liu, Z., Zhao, J., 2017. An optimization-driven approach for computing geodesic paths on triangle meshes. Comput. Aided Des. 90, 105–112. https://doi.org/10.1016/j.cad.2017.05.022.

Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J., 2015. SMPL: a skinned multi-person linear model. ACM Trans. Graph. 34, 248. https://doi.org/10.1145/2816795.2818013.

Monti, F., Boscaini, D., Masci, J., Rodolà, E., Svoboda, J., Bronstein, M.M., 2017. Geometric deep learning on graphs and manifolds using mixture model CNNs. In: Proc. CVPR, pp. 5425–5434.

Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L., 2012. Functional maps: a flexible representation of maps between shapes. ACM Trans. Graph. 31, 30. https://doi.org/10.1145/2185520.2185526.

Pokrass, J., Bronstein, A.M., Bronstein, M.M., Sprechmann, P., Sapiro, G., 2013. Sparse modeling of intrinsic correspondences. Comput. Graph. Forum 32, 459–468. https://doi.org/10.1111/cgf.12066.

Popa, T., Julius, D., Sheffer, A., 2006. Material-aware mesh deformations. In: IEEE International Conference on Shape Modeling and Applications (SMI'06), p. 22.

Pottmann, H., Leopoldseder, S., Hofer, M., 2004. Registration without ICP. Comput. Vis. Image Underst. 95, 54–71. https://doi.org/10.1016/j.cviu.2004.04.002.

Rodolà, E., Bulò, S.R., Windheuser, T., Vestner, M., Cremers, D., 2014. Dense non-rigid shape correspondence using random forests. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. IEEE Computer Society, pp. 4177–4184.

Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D., 2017. Partial functional correspondence. Comput. Graph. Forum 36, 222–236. https://doi.org/10.1111/cgf.12797.

Rusinkiewicz, S., Levoy, M., 2001. Efficient variants of the ICP algorithm. In: Proceedings Third International Conference on 3-D Digital Imaging and Modeling, pp. 145–152.

Rusu, R.B., Blodow, N., Beetz, M., 2009. Fast point feature histograms (FPFH) for 3D registration. In: IEEE International Conference on Robotics and Automation, pp. 3212–3217.

Rusu, R.B., Blodow, N., Marton, Z.C., Beetz, M., 2008. Aligning point cloud views using persistent feature histograms. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3384–3391.

Sauvage, B., Hahmann, S., Bonneau, G.P., Elber, G., 2008. Detail preserving deformation of B-spline surfaces with volume constraint. Comput. Aided Geom. Des. 25, 678–696. https://doi.org/10.1016/j.cagd.2007.10.004.

Sorkine, O., 2006. Differential representations for mesh processing. Comput. Graph. Forum 25, 789–807. https://doi.org/10.1111/j.1467-8659.2006.00999.x.

Sorkine, O., Alexa, M., 2007. As-rigid-as-possible surface modeling. In: Proceedings of the Fifth Eurographics Symposium on Geometry Processing, Eurographics Association, pp. 109–116.

Tam, G.K., Cheng, Z.Q., Lai, Y.K., Langbein, F., Liu, Y., Marshall, A.D., Martin, R., Sun, X., Rosin, P.L., 2013. Registration of 3d point clouds and meshes: a survey from rigid to nonrigid. IEEE Trans. Vis. Comput. Graph. 19, 1199–1217. https://doi.org/10.1109/TVCG.2012.310.

Tam, G.K.L., Martin, R.R., Rosin, P.L., Lai, Y.K., 2014. Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry. ACM Trans. Graph. 33, 4. https://doi.org/10.1145/2517967.

Tombari, F., Salti, S., Di Stefano, L., 2010. Unique Signatures of Histograms for Local Surface Description. Springer Berlin Heidelberg, pp. 356–369.

Vestner, M., Lähner, Z., Boyarski, A., Litany, O., Slossberg, R., Remez, T., Rodolà, E., Bronstein, A.M., Bronstein, M.M., Kimmel, R., Cremers, D., 2017. Efficient deformable shape correspondence via kernel matching. In: 2017 International Conference on 3D Vision (3DV), pp. 517–526.

Vlasic, D., Baran, I., Matusik, W., Popović, J., 2008. Articulated mesh animation from multi-view silhouettes. ACM Trans. Graph. 27, 97. https://doi.org/10.1145/1399504.1360696.

Wei, L., Huang, Q., Ceylan, D., Vouga, E., Li, H., 2016. Dense human body correspondence using convolutional networks. In: Proc. CVPR, pp. 1544–1553.

Yang, J., Li, K., Li, K., Lai, Y.K., 2015. Sparse non-rigid registration of 3D shapes. Comput. Graph. Forum 34, 89–99. https://doi.org/10.1111/cgf.12699.

Zeng, A., Song, S., Niessner, M., Fisher, M., Xiao, J., Funkhouser, T., 2017. 3DMatch: learning the matching of local 3D geometry in range scans. In: Proc. CVPR, pp. 1–12.